

Multimedia— An introduction

by R. J. Flynn
W. H. Tetzlaff

Multimedia—the combination of text, animated graphics, video, and sound—presents information in a way that is more interesting and easier to grasp than text alone. It has been used for education at all levels, job training, and games and by the entertainment industry. It is becoming more readily available as the price of personal computers and their accessories declines. Multimedia as a human-computer interface was made possible some half-dozen years ago by the rise of affordable digital technology. Previously, multimedia effects were produced by computer-controlled analog devices, like videocassette recorders, projectors, and tape recorders. Digital technology's exponential decline in price and increase in capacity has enabled it to overtake analog technology. The Internet is the breeding ground for multimedia ideas and the delivery vehicle of multimedia objects to a huge audience. This paper reviews the uses of multimedia, the technologies that support it, and the larger architectural and design issues.

Introduction

Nowadays, multimedia generally indicates a rich sensory interface between humans and computers or computer-like devices—an interface that in most cases gives the user control over the pace and sequence of the information. We all know multimedia when we see and hear it, yet its precise boundaries elude us. For example, *movies on demand*, in which a viewer can select from a large library of videos and then play, stop, or reposition the tape or change the speed is generally considered multimedia.

However, watching the movie on a TV set attached to a videocassette recorder (VCR) with the same abilities to manipulate the play is not considered multimedia. Unfortunately, we have yet to find a definition that satisfies all experts.

Recent multimedia conferences, such as the IEEE International Conference on Multimedia Computing and Systems, ACM Multimedia, and Multimedia Computing and Networking, provide a good start for identifying the components of multimedia. The range of multimedia activity is demonstrated in papers on multimedia authoring (i.e., specification of multimedia sequences), user interfaces, navigation (user choices), effectiveness of multimedia in education, distance learning, video conferencing, interactive television, video on demand, virtual reality, digital libraries, indexing and retrieval, and support of collaborative work. The wide range of technologies is evident in papers on disk scheduling, capacity planning, resource management, optimization, networking, switched Ethernet LANs, Asynchronous Transfer Mode (ATM) networking, quality of service in networks, Moving Picture Expert Group (MPEG**) encoding, compression, caching, buffering, storage hierarchies, video servers, video file systems, machine classification of video scenes, and Internet audio and video.

Multimedia systems need a delivery system to get the multimedia objects to the user. Magnetic and optical disks were the first media for distribution. The Internet, as well as the Transmission Control Protocol/Internet Protocol (TCP/IP) protocol suite or Net BIOS on isolated or campus LANs, became the next vehicles for distribution. The rich text and graphics capabilities of the World Wide Web browsers are being augmented with animations, video, and sound. Internet distribution will be augmented by distribution via satellite, wireless, and cable systems.

©Copyright 1998 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

0018-8646/98/\$5.00 © 1998 IBM

Multimedia uses and applications

Multimedia applications are primarily existing applications that can be made less expensive or more effective through the use of multimedia technology. In addition, new, speculative applications, like movies on demand, can be created with the technology. We present here a few of these applications.

• Home applications

Video on demand

Video on demand (VOD) [1-5], also called *movies on demand*, is a service that provides movies on an individual basis to television sets in people's homes. The movies are stored in a central server and transmitted through a communication network. A *set-top box* (STB) connected to the communication network converts the digital information to analog and inputs it to the TV set. The viewer uses a remote control device to select a movie and manipulate play through start, stop, rewind, and visual fast forward buttons. The capabilities are very similar to renting a video at a store and playing it on a VCR. The service can provide indices to the movies by title, genre, actors, and director. VOD differs from *pay per view* by providing any of the service's movies at any time, instead of requiring that all purchasers of a movie watch its broadcast at the same time. *Enhanced pay per view*, also a broadcast system, shows the same movie at a number of staggered starting times.

Home shopping and information systems

Services to the home that provide video on demand will also provide other, more interactive, home services. Many kinds of goods and services can be sold this way. The services will help the user navigate through the available material to plan vacations, renew driver's licenses, purchase goods, etc.

Networked games

The same infrastructure that supports home shopping could be used to temporarily download video games with graphic-intensive functionality to the STB, and the games could then be played for a given period of time. Groups of people could play a game together, competing as individuals or working together in teams. Action games would require a very fast, or *low-latency*, network.

• Video conferencing

Currently, most video conferencing is done between two specially set-up rooms. In each room, one or more cameras are used, and the images are displayed on one or more monitors. Text, images, and motion video are compressed and sent through telephone lines. Recently, the technology has been expanded to allow more than two

sites to participate. Video conferences can also be connected through LANs or the Internet [6]. In time, video conferences will be possible from the home.

• Education

A wide range of individual educational software employing multimedia is available on CD-ROM. One of the chief advantages of such multimedia applications is that the sequence of material presented is dependent upon the student's responses and requests. Multimedia is also used in the classroom to enhance the educational experience and augment the teacher's work. Multimedia for education has begun to employ servers and networks to provide for larger quantities of information and the ability to change it frequently.

Distance learning

Distance learning is a variation on education in which not all of the students are in the same place during a class. Education takes place through a combination of stored multimedia presentations, live teaching, and participation by the students. Distance learning involves aspects of both teaching with multimedia and video conferencing.

Just-in-time training

Another variation on education, called *just-in-time training*, is much more effective because it is done right when it is needed. In an industry context, this means that workers can receive training on PCs at their own workplaces at the time of need or of their choosing. This generally implies storing the material on a server and playing it through a wide-area network or LAN.

• Digital libraries

Digital libraries [7, 8] are a logical extension of conventional libraries, which house books, pictures, tapes, etc. Material in digital form can be less expensive to store, easier to distribute, and quicker to find [9, 10]. Thus digital technology can save money and provide better capabilities. The Vatican Library [7] has an extraordinary collection of 150 000 manuscripts, including early copies of works by Aristotle, Dante, Euclid, Homer, and Virgil. However, only about 2000 scholars a year are able to physically visit the library in Rome. Thus, the IBM Vatican Library Project, which makes digitized copies of some of the collection available to scholars around the world, is a very valuable service, especially if the copies distributed are of high quality.

• Virtual reality

Virtual reality [11] provides a very realistic effect through sight and sound, while allowing the user to interact with the virtual world. Because of the ability of the user to

interact with the process, realistic visual effects must be created "on the fly."

- *Telemedicine*

Multimedia and telemedicine [12] can improve the delivery of health care in a number of ways. Digital information can be centrally stored, yet simultaneously available at many locations. Physicians can consult with one another using video conference capabilities, where all can see the data and images, thus bringing together experts from a number of places in order to provide better care. Multimedia can also provide targeted education and support for the patient and family.

Multimedia technology

A wide variety of technologies contribute to multimedia. Some of the technologies are going through rapid improvement and deployment because of demand for PCs and workstations. As a result, multimedia benefits from lower-cost, better-performance microprocessors, memory chips, and disk storage. Other technologies are being developed specifically for multimedia systems. We present here the major technologies relevant to multimedia.

- *Compression technology*

Compression is a key multimedia technology because it reduces the number of bits needed to represent a multimedia object. At present, compression of motion video reduces the number of bits by a factor of approximately one hundred. Without this enormous factor, we would have to wait many years for storage, logic, and memory to become inexpensive enough for uncompressed video data to be used. Compression makes the data not only less expensive to store but also easier to move around.

Techniques

The key to compression of analog data is the similarity of pieces of information that are near each other in time or space. For example, the color on a TV screen is usually almost the same at adjacent points, and it is similar in the next frame. One technique for reducing the number of bits needed to represent data is to use a portion of the data as a base and to calculate frame-to-frame differences from the base. Since the changes in arithmetic values of the data samples are relatively small, the differences are much smaller numbers than the base data values themselves and can be represented by binary numbers with fewer bits.

Another compression technique, called discrete cosine transformation (DCT), treats a series of data values (e.g., brightness) as a function and converts such functions from the time domain to the frequency domain. The number of bits needed to represent the data is then reduced by scaling the coefficients in a process called quantization.

DCT can be used to further compress the differences described above. Following DCT compression, sequences of identical values may be replaced by a single value with a count in a technique called run-length encoding.

While expressing data as the difference from a prior data point is effective in reducing the number of bits, it can lead to problems. Since any digital system experiences an occasional loss or garbling of data, the effect of an error could be reflected in all subsequent values for a data point. In many computer systems, we detect a loss and just reread the data. However, with time-sensitive multimedia data, called *isochronous* data, there is not time to read it and send it a second time. Late data causes interruptions in the audio and video, called *jitter*. The video may blink, stop, or show visible artifacts such as rectangular tiling. The audio may pop or pause in random ways, which people find even more unpleasant than the annoying visual effect. Part of the art of creating compression algorithms involves including resilience to occasional data loss. It must be possible for the receiver to observe special codes in the data stream and resynchronize if necessary, as well as to construct a new, accurate image from which future images can be calculated.

Compression standards specify the organization of the string of transmitted bits and indicate how to interpret them, which really makes them decompression standards. The task of compression can be implemented in arbitrary ways, in which the primary objective might be constant compression ratio or constant visual quality. Compression has two forms: *lossless*, in which an exact duplicate of the original data is reconstructed; and *lossy*, which allows greater compression, but in which the reconstructed digital information is not exactly the same. Most standards allow for lossy operation, which permits an enormous number of possibilities for compressing any particular object.

MPEG

The Moving Picture Expert Group has sponsored the creation of video compression standards ISO/IEC 11171, called MPEG-1, and ISO/IEC 13818, called MPEG-2. MPEG [13, 14] is now the predominant high-quality video standard. MPEG "tiles" each video frame into a matrix of separate *macro blocks*, each of which is treated in the same way. MPEG defines three kinds of frames. "I-frames," or intrapicture frames, contain base data and are independent of all other frames. The periodic occurrence of I-frames allows initialization after a loss of data. "P-frames," or predicted frames, are constructed as differences in the data values from those represented by the prior P-frame or I-frame. "B-frames," or bidirectional prediction frames, are constructed from macro blocks from either the prior P- or I-frame or the next P- or I-frame. This allows for good compression of B-frames despite great changes in the visual content. In order for the data to be

decompressed, the frames must be presented to the decoder somewhat out of time order. P-frames are advanced in the stream so that B-frames always follow the P- and I-frames that they depend upon. MPEG also has the ability to consider either motion in the visual field or the panning of the camera. *Motion vectors* allow the specification of the displacement of macro blocks from one frame to the next. VCR-quality MPEG video can be achieved with a data or transmission rate of 1.5 million bits per second (Mb/s). Higher quality, more comparable to broadcast or cable TV, is in the range of 3 to 6 Mb/s.

Low-bit-rate encoding

Low-bit-rate encoding [15], which has its origins in video conferencing, is used to achieve video with a transmission rate in the tens of thousands of bits per second. ITU H.263 is an example of a low-bit-rate standard. The visual quality is not as good as for MPEG, and the frame rate is lower. As with MPEG, DCTs are used in low-bit-rate encoding, but, in contrast to MPEG, resynchronizing frames are not used.

Real-time encoding

Real-time encoding, which is performed as the content is being created, is extremely challenging. It is certainly needed for video conferencing and distribution of live content such as sporting events. High-quality real-time encoding requires special-purpose hardware to assist the compression.

Off-line encoding

Off-line encoding (performed before the data is used) of stored, uncompressed data produces high-quality video. It is compute-intensive, partly because it makes complex decisions to produce good quality despite lossy compression. This is true of both MPEG and low-bit-rate encoding. Parallel processors are typically used for encoding, but even then the encoding is, at present, slower than real time. General-purpose computers are generally used for off-line encoding.

Decoding

A few years ago, all decoding was done by special-purpose cards that were attached to the bus of a PC or workstation. Each card was engineered to decode one of the many standard, and not so standard, compression formats. The faster PCs and workstations that are now available make decoding by program, called *software decompression*, feasible. However, the processor may not be fast enough in some hardware configurations, like STBs, and special-purpose decompression must be used.

• *Video servers*

The role of a multimedia server [16–18] is to store and deliver multimedia objects, which implies having a mapping from the name of an object to the locations where it is stored. Rotating magnetic disk storage is the preferred medium for storing multimedia objects, with solid-state memory used for buffering.

Multimedia servers are based on computer file servers and will benefit from the research being done to improve the performance and reduce the cost of storage, memory, and microprocessors. Research to specifically reduce the cost of storing and playing back video files is important. Research is also appropriate to achieve the unique requirement of isochronous delivery of audio and video files because it is a requirement not found in other computer servers.

Server structure

Multimedia servers are similar to their predecessors: network file servers. However, they satisfy several added requirements. Sometimes many or all of the clients are playing the same material. Multimedia objects must be played continuously at a constant rate once play begins. Multimedia servers must be scalable in order to serve very large numbers of clients. There are only a few ways [16] to create a large, scalable server from many microprocessors, disks, and network adapters.

The requirements lead to several common solutions. The need to play the same material to many clients requires disk throughput far beyond that of a single disk. *Replication* [19, 20] of content, i.e., extra copies on different disks, increases the number of clients that may access it simultaneously. However, this may cause other disks to be underutilized, because their content is not being accessed. An alternative solution is to *stripe* [4, 21, 22] each item across all of the disks. Striping is more resilient to changes in the relative frequency of viewing the various items than replication because it is not necessary to know the distribution in advance in order to achieve balanced use.

Video servers need special processor and disk scheduling to ensure that disk blocks are read and delivered to the network on time. In addition, they must know exactly how much total capacity to deliver video stream exists in order to ensure that when a new playout is started it can be completed without jitter.

Disk organization and scheduling

Video file systems typically stripe content across multiple disk devices, with the granularity of striping (block of contiguous data stored on each disk drive) usually being the amount of data read in a single I/O operation ("seek" followed by "read"). The data to be read in a single I/O operation is placed in sequential locations on disk, to be

sure that the disk arm does not have to be moved within an I/O operation. In order that the overhead (seek time plus disk rotation time) be a small fraction of the total read operation, reads are generally at least 200 kilobytes. These stripe units are then distributed across the available disks.

At any time, a multimedia server generally has a large number of disk read operations queued. Several disk-scheduling algorithms have been developed to maximize disk throughput for traditional applications by ordering the reads according to physical position on the disk in order to reduce the total seek time. However, these algorithms applied to multimedia servers might delay some reads to a point at which jitter would result. The challenge is to develop algorithms that will improve the throughput without causing jitter. Some of the algorithms considered [23] are the following: first in first out, earliest deadline first (EDF), which guarantees no jitter but does not reduce seeks; Scan, which orders by location but does not guarantee service; Scan-EDF, which blends Scan for seek reduction and EDF for service; and Group Sweep Scheduling [24], which controls the amount of out-of-order reading that is done with Scan.

Video servers are expected to be highly available in much the same way that the telephone system is. In order to be sure the content is available, despite disk failures, some form of duplicate or redundant storage [22, 25] of video data is used.

Video file systems

Video file systems [4, 25, 26], which incorporate disk layout, disk scheduling, and service guarantees, are either modified traditional file systems or file systems specially designed for video. Some of the functions, like disk scheduling, are traditionally done by the operating system and not the file system. One of the tensions in video file system design is the desire to be independent of device characteristics, for the sake of serviceability and extendibility, yet to create algorithms that take advantage of device characteristics for the sake of throughput and service.

Storage hierarchy

In classical computer architecture, a small amount of high-speed, high-cost storage is used to store the most frequently referenced data. This provides average data-access time only slightly greater than that of the high-speed storage at an average cost just above the cost of the low-speed storage. In large video servers, it may be economical to store some material in high-cost solid-state storage, other material on disk, and low-frequency material on an archival device [27, 28].

• *Networks*

Telephone networks dedicate a set of resources that forms a complete path from end to end for the duration of the telephone connection. The dedicated path guarantees that the voice data can be delivered from one end to the other end in a smooth and timely way, but the resources remain dedicated even when there is no talking. In contrast, digital packet networks, for communication between computers, use time-shared resources (links, switches, and routers) to send packets through the network. The use of shared resources allows computer networks to be used at high utilization, because even small periods of inactivity can be filled with data from a different user. The high utilization and shared resources create a problem with respect to the timely delivery of video and audio over data networks. Current research centers around reserving resources for time-sensitive data, which will make digital data networks more like telephone voice networks.

Internet

The Internet [29, 30] and intranets, which use the TCP protocol suite [29], are the most important delivery vehicles for multimedia objects. TCP [29] provides communication sessions between applications on hosts, sending streams of bytes for which delivery is always guaranteed by means of acknowledgments and retransmission. User Datagram Protocol (UDP) [29] is a "best-effort" delivery protocol (some messages may be lost) that sends individual messages between hosts. Internet technology is used on single LANs and on connected LANs within an organization, which are sometimes called intranets, and on "backbones" that link different organizations into one single global network. Internet technology allows LANs and backbones of totally different technologies to be joined together into a single, seamless network.

Part of this is achieved through communications processors called routers. Routers can be accessed from two or more networks, passing data back and forth as needed. The routers communicate information on the current network topology among themselves in order to build routing tables within each router. These tables are consulted each time a message arrives, in order to send it to the next appropriate router, eventually resulting in delivery.

Token ring

Token ring [31] is a hardware architecture for passing packets between stations on a LAN. Since a single circular communication path is used for all messages, there must be a way to decide which station is allowed to send at any time. In token ring, a "token," which gives a station the right to transmit data, is passed from station to station. The data rate of a token ring network is 16 Mb/s.

Ethernet

Ethernet [31] LANs use a common wire to transmit data from station to station. Mediation between transmitting stations is done by having stations listen before sending, so that they will not interfere with each other. However, two stations could begin to send at the same time and collide, or one station could start to send significantly later than another but not know it because of propagation delay. In order to detect these other situations, stations continue to listen while they transmit and determine whether their message was possibly garbled by a collision. If there is a collision, a retransmission takes place (by both stations) a short but random time later. Ethernet LANs can transmit data at 10 Mb/s. However, when multiple stations are competing for the LAN, the throughput may be much lower because of collisions and retransmissions.

Switched Ethernet

Switches may be used at a hub to create many small LANs where one large one existed before. This reduces contention and permits higher throughput. In addition, Ethernet is being extended to 100Mb/s throughput. The combination, *switched Ethernet*, is much more appropriate to multimedia than regular Ethernet, because existing Ethernet LANs can support only about six MPEG video streams, even when nothing else is being sent over the LAN.

ATM

Asynchronous Transfer Mode (ATM) [29, 32] is a new packet-network protocol designed for mixing voice, video, and data within the same network. Voice is digitized in telephone networks at 64 Kb/s (kilobits per second), which must be delivered with minimal delay, so very small packet sizes are used. On the other hand, video data and other business data usually benefit from quite large block sizes. An ATM packet consists of 48 octets (the term used in communications for eight bits, called a byte in data processing) of data preceded by five octets of control information. An ATM network consists of a set of communication links interconnected by switches. Communication is preceded by a setup stage in which a path through the network is determined to establish a circuit. Once a circuit is established, 53-octet packets may be streamed from point to point.

ATM networks can be used to implement parts of the Internet by simulating links between routers in separate intranets. This means that the "direct" intranet connections are actually implemented by means of shared ATM links and switches.

ATM, both between LANs and between servers and workstations on a LAN, will support data rates that will allow many users to make use of motion video on a LAN.

• *Data-transmission techniques*

Modems

Modulator/demodulators, or modems, are used to send digital data over analog channels by means of a carrier signal (sine wave) modulated by changing the frequency, phase, amplitude, or some combination of them in order to represent digital data. (The result is still an analog signal.) Modulation is performed at the transmitting end and demodulation at the receiving end. The most common use for modems in a computer environment is to connect two computers over an analog telephone line. Because of the quality of telephone lines, the data rate is commonly limited to 28.8 Kb/s. For transmission of customer analog signals between telephone company central offices, the signals are sampled and converted to "digital form" (actually, still an analog signal) for transmission between offices. Since the customer voice signal is represented by a stream of digital samples at a fixed rate (64 Kb/s), the data rate that can be achieved over analog telephone lines is limited.

ISDN

Integrated Service Digital Network (ISDN) extends the telephone company digital network by sending the digital form of the signal all the way to the customer. ISDN is organized around 64Kb/s transmission speeds, the speed used for digitized voice. An ISDN line was originally intended to simultaneously transmit a digitized voice signal and a 64Kb/s data stream on a single wire. In practice, two channels are used to produce a 128Kb/s line, which is faster than the 28.8Kb/s speed of typical computer modems but not adequate to handle MPEG video.

ADSL

Asymmetric Digital Subscriber Lines (ADSL) [33–35] extend telephone company twisted-pair wiring to yet greater speeds. The lines are asymmetric, with an outbound data rate of 1.5 Mb/s and an inbound rate of 64 Kb/s. This is suitable for video on demand, home shopping, games, and interactive information systems (collectively known as interactive television), because 1.5 Mb/s is fast enough for compressed digital video, while a much slower "back channel" is needed for control. ADSL uses very high-speed modems at each end to achieve these speeds over twisted-pair wire.

ADSL is a critical technology for the Regional Bell Operating Companies (RBOCs), because it allows them to use the existing twisted-pair infrastructure to deliver high data rates to the home.

- *Cable systems*

Cable television systems provide analog broadcast signals on a coaxial cable, instead of through the air, with the attendant freedom to use additional frequencies and thus provide a greater number of channels than over-the-air broadcast. The systems are arranged like a branching tree, with "splitters" at the branch points. They also require amplifiers for the outbound signals, to make up for signal loss in the cable. Most modern cable systems [35–37] use fiber optic cables for the trunk and major branches and use coaxial cable for only the final loop, which services one or two thousand homes. The root of the tree, where the signals originate, is called the *head end*.

Cable modems

Cable modems are used to modulate digital data, at high data rates, into an analog 6-MHz-bandwidth TV-like signal. These modems can transfer 20 to 40 Mb/s in a frequency bandwidth that would have been occupied by a single analog TV signal, allowing multiple compressed digital TV channels to be multiplexed over a single analog channel. The high data rate may also be used to download programs or World Wide Web content or to play compressed video. Cable modems are critical to cable operators, because it enables them to compete with the RBOCs using ADSL.

Set-top box

The STB is an appliance that connects a TV set to a cable system, terrestrial broadcast antenna, or satellite broadcast antenna. The STB in most homes has two functions. First, in response to a viewer's request with the remote-control unit, it shifts the frequency of the selected channel to either channel 3 or 4, for input to the TV set. Second, it is used to restrict access and block channels that are not paid for. Addressable STBs respond to orders that come from the head end to block and unblock channels.

- *Admission control*

Digital multimedia systems that are shared by multiple clients can deliver multimedia data to a limited number of clients. Admission control is the function which ensures that once delivery starts, it will be able to continue with the required quality of service (ability to transfer isochronous data on time) until completion. The maximum number of clients depends upon the particular content being used and other characteristics of the system.

- *Digital watermarks*

Because it is so easy to transmit perfect copies of digital objects, many owners of digital content wish to control unauthorized copying. This is often to ensure that proper royalties have been paid. Digital watermarking [38, 39] consists of making small changes in the digital data that

can later be used to determine the origin of an unauthorized copy. Such small changes in the digital data are intended to be invisible when the content is viewed. This is very similar to the "errors" that mapmakers introduce in order to prove that suspect maps are copies of their maps. In other circumstances, a visible watermark is applied in order to make commercial use of the image impractical.

- *Authoring systems*

Multimedia *authoring systems* are used to edit and arrange multimedia objects and to describe their presentation. The authoring package allows the author to specify which objects may be played next. The viewer dynamically chooses among the alternatives. Metadata created during the authoring process is normally saved as a file. At play time, an "execution package" reads the metadata and uses it as a script for the playout.

Authoring systems, as well as systems for gathering information for multimedia presentations (scanning, classifying, indexing and processing images, audio, and video) are very active research areas. Particularly challenging, and also very useful, are techniques that can be applied to compressed data. Entirely new techniques are required, and the human factors involved in the processing of this new data must be understood.

Multimedia architecture

In this section we show how the multimedia technologies are organized in order to create multimedia systems, which in general consist of suitable organizations of clients, application servers, and storage servers that communicate through a network. Some multimedia systems are confined to a stand-alone computer system with content stored on hard disks or CD-ROMs. Distributed multimedia systems communicate through a network and use many shared resources, making quality of service very difficult to achieve and resource management very complex.

- *Single-user stand-alone systems*

Stand-alone multimedia systems use CD-ROM disks and/or hard disks to hold multimedia objects and the scripting metadata to orchestrate the playout. CD-ROM disks are inexpensive to produce and hold a large amount of digital data; however, the content is static—new content requires creation and physical distribution of new disks for all systems. Decompression is now done by either a special decompression card or a software application that runs on the processor. The technology trend is toward software decompression.

• Multi-user systems

Video over LANs

Stand-alone multimedia systems can be converted to networked multimedia systems by using client-server remote-file-system technology to enable the multimedia application to access data stored on a server as if the data were on a local storage medium. This is very convenient, because the stand-alone multimedia application does not have to be changed. LAN throughput is the major challenge in these systems. Ethernet LANs can support less than 10 Mb/s, and token rings 16 Mb/s. This translates into six to ten 1.5Mb/s MPEG video streams. Admission control is a critical problem. The OS/2* LAN server is one of the few products that support admission control [40]. It uses priorities with token-ring messaging to differentiate between multimedia traffic and lower-priority data traffic. It also limits the multimedia streams to be sure that they do not sum to more than the capacity of the LAN. Without some type of resource reservation and admission control, the only way to give some assurance of continuous video is to operate with small LANs and make sure that the server is on the same LAN as the client. In the future, ATM and fast Ethernet will provide capacity more appropriate to multimedia.

Direct Broadcast Satellite

Direct Broadcast Satellite (DBS), which broadcasts up to 80 channels from a satellite at high power, arrived in 1995 as a major force in the delivery of broadcast video. The high power allows small (18-inch) dishes with line-of-sight to the satellite to capture the signal. MPEG compression is used to get the maximum number of channels out of the bandwidth. The RCA/Hughes service employs two satellites and a backup to provide 160 channels. This large number of channels allows many premium and special-purpose channels as well as the usual free channels. Many more pay-per-view channels can be provided than in conventional cable systems. This allows enhanced pay-per-view, in which the same movie is shown with staggered starting times of half an hour or an hour.

DBS requires a set-top box with much more function than a normal cable STB. The STB contains a demodulator to reconstruct the digital data from the analog satellite broadcast. The MPEG compressed form is decompressed, and a standard TV signal is produced for input to the TV set. The STB uses a telephone modem to periodically verify that the premium channels are still authorized and report on use of the pay-per-view channels so that billing can be done.

Interactive TV and video to the home

Interactive TV and video to the home [2-5] allow viewers to select, interact with, and control video play on a TV set

in real time. The user might be viewing a conventional movie, doing home shopping, or engaging in a network game. The compressed video flowing to the home requires high bandwidth, from 1.5 to 6 Mb/s, while the return path, used for selection and control, requires far lower bandwidth.

The STB used for interactive TV is similar to that used for DBS. The demodulation function depends upon the network used to deliver the digital data. A microprocessor with memory for limited buffering as well as an MPEG decompression chip is needed. The video is converted to a standard TV signal for input to the TV set. The STB has a remote-control unit, which allows the viewer to make choices from a distance. Some means are needed to allow the STB to relay viewer commands back to the server, depending upon the network being used.

Cable systems appear to be broadcast systems, but they can actually be used to deliver different content to each home. Cable systems often use fiber optic cables to send the video to converters that place it on local loops of coaxial cable. If a fiber cable is dedicated to each final loop, which services 500 to 1500 homes, there will be enough bandwidth to deliver an individual signal to many of those houses. The cable can also provide the reverse path to the cable head end. Ethernet-like protocols can be used to share the same channel with the other STBs in the local loop. This topology is attractive to cable companies because it uses the existing cable plant. If the appropriate amplifiers are not present in the cable system for the back channel, a telephone modem can be used to provide the back channel.

As mentioned above, the asymmetric data rates of ADSL are tailored for interactive TV. The use of standard twisted-pair wire, which has been brought to virtually every house, is attractive to the telephone industry. However, the twisted pair is a more noisy medium than coaxial cable, so more expensive modems are needed, and distances are limited. ADSL can be used at higher data rates if the distance is further reduced.

Interactive TV architectures are typically three-tier, in which the client and server tiers interact through an application server. (In three-tier systems, the tier-1 systems are clients, the tier-2 systems are used for application programs, and the tier-3 systems are data servers.) The application tier is used to separate the logic of looking up material in indexes, maintaining the shopping state of a viewer, interacting with credit card servers, and other similar functions from the simple function of delivering multimedia objects.

The key research questions about interactive TV and video-on-demand are not computer science questions at all. Rather, they are the human-factors issues concerning ease of the on-screen interface and, more significantly, the

marketing questions regarding what home viewers will find valuable and compelling.

Internet over cable systems

World Wide Web browsing allows users to see a rich text, video, sound, and graphics interface and allows them to access other information by clicking on text or graphics. Web pages are written in HyperText Markup Language (HTML) and use an application communications protocol called HTTP. The user responses, which select the next page or provide a small amount of text information, are normally quite short. On the other hand, the graphics and pictures require many times the number of bytes to be transmitted to the client. This means that distribution systems that offer asymmetric data rates are appropriate.

Cable TV systems can be used to provide asymmetric Internet access for home computers in ways that are very similar to interactive TV over cable. The data being sent to the client is digitized and broadcast over a prearranged channel over all or part of the cable system. A *cable modem* at the client end tunes to the right channel and demodulates the information being broadcast. It must also filter the information destined for the particular station from the information being sent to other clients. The low-bandwidth reverse channel is the same low-frequency band that is used in interactive TV. As with interactive TV, a telephone modem might be used for the reverse channel. The cable head end is then attached to the Internet using a router. The head end is also likely to offer other services that Internet Service Providers sell, such as permanent mailboxes. This asymmetric connection would not be appropriate for a Web server or some other type of commerce server on the Internet, because servers transmit too much data for the low-speed return path. The cable modem provides the physical link for the TCP/IP stack in the client computer. The client software treats this environment just like a LAN connected to the Internet.

Video servers on a LAN

LAN-based multimedia systems [4, 6, 15] go beyond the simple, client-server, remote file system type of video server, to advanced systems that offer a three-tier architecture with clients, application servers, and multimedia servers. The application servers provide applications that interact with the client and select the video to be shown. On a company intranet, LAN-based multimedia could be used for just-in-time education, on-line documentation of procedures, or video messaging. On the Internet, it could be used for a video product manual, interactive video product support, or Internet commerce. The application server chooses the video to be shown and causes it to be sent to the client.

There are three different ways that the application server can cause playout of the video: By giving the

address of the video server and the name of the content to the client, which would then fetch it from the video server; by communicating with the video server and having it send the data to the client; and by communicating with both to set up the relationship.

The transmission of data to the client may be in *push mode* or *pull mode*. In push mode, the server sends data to the client at the appropriate rate. The network must have quality-of-service guarantees to ensure that the data gets to the client on time. In pull mode, the client requests data from the server, and thus paces the transmission.

The current protocols for Internet use are TCP and UDP. TCP sets up sessions, and the server can push the data to the client. However, the "moving-window" algorithm of TCP, which prevents client buffer overrun, creates acknowledgments that pace the sending of data, thus making it in effect a pull protocol. Another issue in Internet architecture is the role of firewalls, which are used at the gateway between an intranet and the Internet to keep potentially dangerous or malicious Internet traffic from getting onto the intranet. UDP packets are normally never allowed in. TCP sessions are allowed, if they are created from the inside to the outside. A disadvantage of TCP for isochronous data is that error detection and retransmission is automatic and required—whereas it is preferable to discard garbled video data and just continue.

Resource reservation is just beginning to be incorporated on the Internet and intranets. Video will be considered to have higher priority, and the network will have to ensure that there is a limit to the amount of high-priority traffic that can be admitted. All of the routers on the path from the server to the client will have to cooperate in the reservation and the use of priorities.

Video conferencing

Video conferencing [6, 15], which will be used on both intranets and the Internet, uses multiple data types, and serves multiple clients in the same conference. Video cameras can be mounted near a PC display to capture the user's picture. In addition to the live video, these systems include shared white boards and show previously prepared visuals. Some form of mediation is needed to determine which participant is in control. Since the type of multimedia data needed for conferencing requires much lower data rates than most other types of video, low-bit-rate video, using approximately eight frames per second and requiring tens of kilobits per second, will be used with small window sizes for the "talking heads" and most of the other visuals. Scalability of a video conferencing system is important, because if all participants send to all other participants, the traffic goes up as the square of the number of participants. This can be made linear by having all transmissions go through a common server. If the

network has a multicast facility, the server can use that to distribute to the participants.

Summary

Multimedia is obviously a fertile ground for both research and the development of new products, because of the breadth of possible usage, the dependency on a wide range of technologies, and the value of reducing cost by improving the technology. Now that the technology has been developed, however, the marketplace will determine future direction. The technology will be used when clear value is found. For example, multimedia is widely used on PCs using CDs to store the content. The CDs are inexpensive to reproduce, and the players are standard equipment on most PCs purchased today. The acceptance caused a greater demand for players, which, in turn, caused greater production and further reduced prices.

The computer industry is providing demand, and an expanding market, for the key hardware technologies that underlie multimedia. These include solid-state memory, logic, microprocessors, modems, switches, and disk storage. The price declines of 30–60% per year that we have seen for several decades will continue into the foreseeable future. As a result, the application of multimedia, which appears expensive now, will become less expensive and more attractive. An exception to this fast rate of improvement is the cost of data communications. Communications depend both on technology with rapidly decreasing cost and on mundane and basically unchanging tasks such as laying cable with the help of a backhoe or stringing cables from poles. The cost of communication is not likely to decline significantly for quite a while.

We feel that multimedia will spread from low-bit-rate to high-bit-rate, and will begin on established intranets first, move to the Internet, and finally be transmitted on broadband connections (ADSL or cable modems) to the home.

The initial uses will be information dissemination, education, and training on campus LANs. Multimedia will be used in education, government, and business over campus LANs, with low-bit-rate video that will not place excessive stress on the infrastructure. The availability of switched LAN technology and faster LANs will allow increases in both the bit rate per user and the number of users. As the cost of communications decreases, the cost for Internet attachment for servers will decline, and higher-quality video will be used on the Internet. Multimedia will be a compelling interface for commerce and advertising on the Internet. Eventually, cable modems and/or ADSL will provide bandwidth for movies to the home, and the declining computer and switching costs will allow a cost-effective service. The winner between ADSL and cable modems will have as much to do with the ability

of cable companies and RBOCs to raise capital as with the inherent cost and value of the two technologies.

IBM researchers continue to play an active role in developing technology, including MPEG encoding and decoding [14, 41–45], video servers [4, 5, 16, 19, 20, 23–25, 28, 46–51], delivery systems [15, 40, 52, 53], digital libraries [7, 54], applications for indexing and searching for content [9, 10, 55], and collaboration [6]. Researchers are also engaged in many uses of multimedia technology and in building advanced systems [3, 7, 47, 56] with IBM customers.

*Trademark or registered trademark of International Business Machines Corporation.

**Trademark or registered trademark of Moving Picture Expert Group.

References

1. W. Hodge, S. Mabon, and J. T. Powers, "Video Demand: Architecture, Systems and Applications," *SMPTÉ J.* **102**, No. 9, 791–803 (September 1993).
2. W. D. Sincoskie, "System Architecture for a Large Scale Video on Demand Service," *Computer Networks & ISDN Syst.* **22**, 155–162 (1991).
3. T. Sanuki and Y. Asakawa, "Design of a Video-Server Complex for Interactive Television," *IBM J. Res. Develop.* **42**, 199–218 (1998, this issue).
4. R. L. Haskin, "Tiger Shark—A Scalable File System for Multimedia," *IBM J. Res. Develop.* **42**, 185–197 (1998, this issue).
5. M. Kumar, "Video-Server Designs for Supporting Very Large Numbers of Concurrent Users," *IBM J. Res. Develop.* **42**, 219–232 (1998, this issue).
6. C. Bisdikian, S. Brady, Y. N. Doganata, D. A. Foulger, F. Marconcini, M. Mourad, H. L. Operowsky, G. Pacifici, and A. N. Tantawi, "MultiMedia Digital Conferencing: A Web-Enabled Multimedia Teleconferencing System," *IBM J. Res. Develop.* **42**, 281–298 (1998, this issue).
7. F. C. Mintzer, L. E. Boyle, A. N. Cazes, B. S. Christian, S. C. Cox, F. P. Giordano, H. M. Gladney, J. C. Lee, M. L. Kelmanson, A. C. Lirani, K. A. Magerlein, A. M. B. Pavani, and F. Schiattarella, "Toward On-Line, Worldwide Access to Vatican Library Materials," *IBM J. Res. Develop.* **40**, No. 2, 139–162 (March 1996).
8. E. Fox, R. Akscyn, R. Furuta, and J. Leggett, "Digital Libraries," *Commun. ACM* **38**, No. 4, 23–28 (April 1995).
9. R. M. Bolle, B.-L. Yeo, and M. M. Yeung, "Video Query: Research Directions," *IBM J. Res. Develop.* **42**, 233–252 (1998, this issue).
10. V. Castelli, L. D. Bergman, I. Kontoyiannis, C.-S. Li, J. T. Robinson, and J. J. Turek, "Progressive Search and Retrieval in Large Image Archives," *IBM J. Res. Develop.* **42**, 253–268 (1998, this issue).
11. G. Singh, S. Feiner, and D. Thalmann, "Virtual Reality—Software and Technology," *Commun. ACM* **38**, No. 5, 35–36 (May 1996).
12. V. Pebsworth Debold, E. Johnson, A. Cameron, and D. Chuang, "Telemedicine: An Annotated Bibliography: Part III," *Telemedicine J.* **1**, No. 4, 377–394 (Winter 1995).
13. D. LeGall, "MPEG: A Video Compression Standard for Multimedia Applications," *Commun. ACM* **34**, No. 4, 46–58 (April 1991).
14. Joan L. Mitchell, William B. Pennebaker, Chad E. Fogg, and Didier J. LeGall, *MPEG Video Compression Standard*, Chapman & Hall, New York, 1997.

15. M. H. Willebeek-LeMair, K. G. Kumar, and E. C. Snible, "Bamba—Audio and Video Streaming Over the Internet," *IBM J. Res. Develop.* **42**, 269–280 (1998, this issue).
16. W. Tetzlaff and R. Flynn, "Elements of Scalable Video Servers," *Proceedings of the IEEE Computer Society International Conference, COMPCON '95*, pp. 239–248.
17. H. M. Vin and P. V. Rangan, "Designing a MultiUser HDTV Storage Server," *IEEE J. Selected Areas in Commun.* **11**, No. 1, 153–164 (January 1993).
18. J. Gemmell, H. Fin, D. Kandlur, and V. Rangan, "Multimedia Storage Servers: A Tutorial and Survey," *IEEE Computer* **28**, No. 5, 40–51 (May 1995).
19. A. Dan, M. Kienzle, and D. Sitaram, "A Dynamic Segment Replication for Load-Balancing in Video-on-Demand Servers," *Multimedia Syst.* **3**, No. 3, 93–103 (July 1995).
20. J. Wolf, P. Yu, and H. Shachnai, "Disk Load Balancing for Video-on-Demand Systems," *Multimedia Syst.* **5**, No. 6, 358–370 (1997).
21. M. Y. L. Kim, "Synchronously Interleaved Disk Systems with Their Application to the Very Large FFT," Ph.D. thesis, Polytechnic University, New York, 1987.
22. D. A. Patterson, G. Gibson, and R. H. Katz, "A Case for Redundant Arrays of Inexpensive Disks (RAID)," *Proceedings of ACM SIGMOD 88*, pp. 109–116.
23. A. Reddy and J. Wyllie, "I/O Issues in a Multimedia System," *IEEE Computer* **27**, No. 3, 69–74 (March 1994).
24. P. S. Yu, M. Chen, and D. D. Kandlur, "Grouped Sweeping Scheduling for DASD-Based Multimedia," *Multimedia Syst.* **1**, No. 3, 99–109 (1993).
25. M. S. Chen, H. Hsiao, C. Li, and P. Yu, "Using Rotational Mirrored Declustering for Replica Placement in a Disk Array Based Video Server," *Proceedings of the 3rd ACM Conference on Multimedia*, November 1995, pp. 121–130.
26. H. Vin and P. Rangan, "Designing File Systems for Digital Video and Audio," *Proceedings of the 13th ACM Symposium on Operating Systems Principles, Oper. Syst. Rev.* **25**, No. 5, 81–94 (1991).
27. A. Chervenak, "Tertiary Storage: An Evaluation of New Applications," Ph.D. Thesis, University of California at Berkeley, Technical Report UDB/CSD 94/847, December 1994.
28. M. Kienzle, A. Dan, D. Sitaram, and W. Tetzlaff, "Using Tertiary Storage in Video on Demand Servers," *Proceedings of the IEEE Computer Society International Conference, COMPCON '95*, pp. 225–233.
29. D. Comer, *Internet Working with TCP/IP*, Prentice-Hall, Inc., Englewood Cliffs, NJ, 1995.
30. M. Macedonia and D. Brutzman, "Mbone Provides Audio and Video Across the Internet," *IEEE Computer* **27**, No. 4, 30–36 (April 1994).
31. J. Martin, K. Chapman, and J. Leben, *Local Area Networks*, Prentice-Hall, Inc., Englewood Cliffs, NJ, 1994.
32. J. LeBoudec, "The Asynchronous Transfer Mode: A Tutorial," *Computer Networks & ISDN Syst.* **24**, 279–309 (1992).
33. P. J. Keyes, R. C. McConnell, and K. Sistanizadeh, "ADSL: A New Twisted Pair Access to the Information Highway," *IEEE Communications Magazine* **33**, No. 4, 52–59 (April 1995).
34. E. V. Jones, "Asymmetric Digital Subscriber Line (ADSL) Systems," *IEEE Colloquium on High-Speed Access Technology and Services Including Video on Demand*, Digest No. 1994/192, 1994, pp. 1/1–1/8.
35. R. Brown, "Innovative Technology Tests Running Side by Side," *Commun. Eng. Design* **20**, No. 3, 33, 36–37 (March 1994).
36. A. Paff, "Hybrid Fiber/Coax in the Public Telecommunications Infrastructure," *IEEE Communications Magazine* **33**, No. 4, 40–45 (April 1995).
37. B. W. Phillips, "Broadband in the Local Loop," *Telecommun.* **28**, No. 11, 37–42 (November 1994).
38. H. Berghel and L. O'Gorman, "Protecting Ownership Rights Through Digital Watermarking," *IEEE Computer* **29**, No. 7, 101–103 (July 1996).
39. H. Berghel, "Watermarking Cyberspace," *Commun. ACM* **40**, No. 11, 19–24 (November 1997).
40. M. Baugher, "The OS/2 Resource Reservation System," *Proceedings of SPIE Multimedia Computing and Networking Conference*, Boston, February 6–8, 1995, pp. 167–176.
41. X. Li and C. Gonzales, "A Locally Quadratic Model of the Motion Estimation Error Criterion Function and Its Application to Subpixel Interpolations," *IEEE Trans. Circuits & Syst. for Video Technol.* **6**, 118 (February 1996).
42. C. Gonzales and E. Viscito, "Flexibly Scalable Digital Video," *Signal Proc.: Image Commun.* **5**, 5 (February 1993).
43. C. Gonzales and E. Viscito, "Motion Video Adaptive Quantization in the Transform Domain," *IEEE Trans. Circuits & Syst. for Video Technol.* **1**, 374 (December 1991).
44. E. Viscito and C. Gonzales, "A Video Compression Algorithm with Adaptive Bit Allocation and Quantization," *Proc. SPIE* **1605**, 205 (1991).
45. Roy Hoffman, *Data Compression in Digital Systems*, Chapman & Hall, New York, 1997.
46. R. Haskin, "The Shark Continuous Media File Server," *Proceedings of the IEEE Computer Society International Conference, COMPCON '93*, pp. 12–17.
47. M. Kienzle, R. Berbec, G. Bozman, C. Eilert, M. Eshel, and R. Mansell, "Multimedia File Serving with the OS/390 LAN Server," *IBM Syst. J.* **36**, No. 3, 374–392 (1997).
48. A. Dan, S. Dulin, S. Marcotte, and D. Sitaram, "Multimedia Resource Management in the OS/390 LAN Server," *IBM Syst. J.* **36**, No. 3, 393–410 (1997).
49. A. Dan, D. Sitaram, and P. Shahabuddin, "Dynamic Batching Policies for an On-Demand Video Server," *Multimedia Syst.* **4**, No. 3, 112–121 (June 1996).
50. M.-S. Chen, D. Kandlur, and P. Yu, "Storage and Retrieval Methods to Support Fully Interactive Payout in a Disk-Array-Based Video Server," *Multimedia Syst.* **3**, No. 3, 126–135 (July 1995).
51. P. Yu, J. Wolf, and H. Shachnai, "Design and Analysis of a Look-Ahead Scheduling Scheme to Support Pause-Resume for Video-on-Demand Applications," *Multimedia Syst.* **3**, No. 4, 137–150 (September 1995).
52. E. Basturk, A. Birman, G. Delp, R. Guerin, R. Haas, S. Kamat, D. Kandlur, P. Pan, D. Pendarakis, R. Rajan, D. Saha, and D. Williams, "Design and Implementation of a QoS Capable Switch-Router," *Proceedings of the IEEE International Conference on Computer Communications and Networks (IC3N'97)*, September 1997, pp. 276–284.
53. W. Feng, D. Kandlur, D. Saha, and K. Shin, "Understanding TCP Dynamics in an Integrated Services Internet," *Proceedings of the 7th International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV'97)*, May 1997, pp. 279–290.
54. H. M. Gladney, "A Storage Subsystem for Image and Records Management," *IBM Syst. J.* **32**, No. 3, 512–540 (1993).
55. M. Flickner, H. Sawhney, W. Niblack, J. Ashley, Q. Huang, B. Dom, M. Gorkiani, J. Hafner, D. Lee, D. Petkovic, D. Steele, and P. Yanker, "Query by Image and Video Content: The QBIC System" (special issue on content-based retrieval), *IEEE Computer* **28**, 23–32 (September 1995).
56. C. Faloutsos, R. Barber, M. Flickner, J. Hafner, W. Niblack, D. Petkovic, and W. Equitz, "Efficient and Effective Querying by Image Content," *J. Intell. Info. Syst., Integr. Artif. Intell. Database Technol. (Netherlands)* **3**, No. 3/4, 231–262 (July 1994).
57. W. Niblack, R. Barber, W. Equitz, M. Flickner, E. Glasman, D. Petkovic, P. Yanker, C. Faloutsos, and G.

- Taubin, "The QBIC Project: Querying Images by Content Using Color, Texture, and Shape," *Proc. SPIE* **1908**, 173-187 (1993).
58. R. Haskin and F. Stein, "A System for the Delivery of Interactive Television Programming," *Proceedings of the IEEE Computer Society International Conference, COMPCON '95*, March 6-8, 1995.

Received December 12, 1996; accepted for publication August 20, 1997

Robert J. Flynn *Polytechnic University, 36 Saw Mill River Road, Hawthorne, New York 10532 (flynn@poly.edu).* Dr. Flynn is Industry Professor of Computer Science at Polytechnic University. He is Director of Graduate Computer Science and Director of Information Systems Engineering at the Polytechnic's campus in Hawthorne, New York. Dr. Flynn holds a B.S. degree in physics from Manhattan College and an M.S. and Ph.D. in mathematics from Brooklyn Poly (the former name of Polytechnic University). He was a founding member of Poly's computer science efforts, both within the electrical engineering department and as a freestanding department. Dr. Flynn's areas of interest include machine organization and operating systems, with special interests in memory hierarchies, storage systems, systems security, cable modems, and associated networking issues, and evolutionary methods. He is a frequent consultant on these issues as well as an academic in these areas. He is a member of ACM, IEEE, SIAM, AMS, Sigma Xi, and Sigma Pi Sigma. Dr. Flynn has received several distinguished teaching awards, including the designation of Institute Distinguished Teacher.

William H. Tetzlaff *IBM Research Division, P.O. Box 704, Yorktown Heights, New York 10598 (tetzlaff@watson.ibm.com).* Mr. Tetzlaff is a Distinguished Engineer and Senior Manager of Systems Structure in the IBM Research Division. He received his B.S. in engineering from Northwestern University and his M.S. in computer science from Polytechnic University, Hawthorne, New York. He is currently working toward completion of his Ph.D. at Polytechnic, where his subject is scheduling, striping, and scaling issues for video-server file systems. He is a member of the ACM, Upsilon Pi Epsilon, Tau Beta Pi, and Sigma Xi and a Senior Member of the IEEE. Mr. Tetzlaff has received a Corporate Award for his work on System Managed Storage, and he has been elected to the IBM Academy of Technology. He has received five major IBM awards for outstanding contributions and technical innovation. These awards were for the Statistics Generating Package, a page-swapping subsystem for VM, the VM/Performance Reporting Facility, System-Managed Storage for MVS, and a video server.