

Evolution and challenges in multimedia

by A. Dan
S. I. Feldman
D. N. Serpanos

This paper outlines the challenges and issues in developing multimedia-based application programs and solutions. It provides a brief history of multimedia-based applications and a critical perspective of the research issues that should be addressed in the future.

Introduction

The integration of multiple media, e.g., voice, video, image, and data, provides an effective means of communication to the users of various services. Because of the advances in computer and communication technologies, creating sophisticated multimedia user interfaces is no longer limited to special-purpose applications. The popularity of the World Wide Web, where most applications currently utilize images and data, is adequate testimony to this fact. The increasing power of electronic circuitry in workstations, personal computers, and consumer electronics, in conjunction with the decreasing cost of high-bandwidth and low-latency communication, have created a large momentum to develop sophisticated multimedia applications as well as to provide new types of services to businesses and homes. These capabilities involve improving user interfaces to many traditional applications as well as creating futuristic applications such as *virtual environments* and *augmented reality* [1]. The basic idea behind the latter types of applications is to immerse a user in an imaginary, computer-generated *virtual world* or to augment the real world (i.e., via augmentation of human perception by supplying information not ordinarily detectable by human senses) around a user with superimposed computer

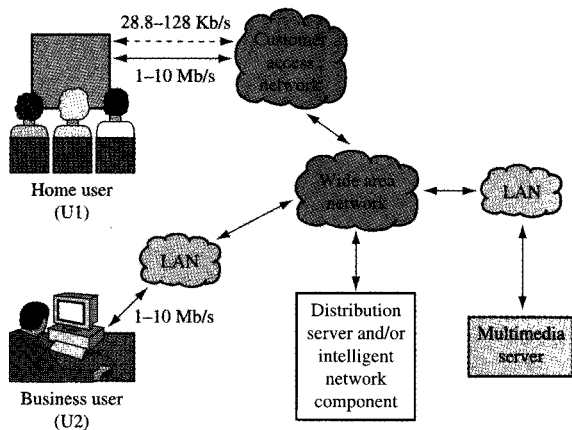
graphics projected onto the walls or onto a head-mounted display. In another example (referred to as *telepresence*), a multitude of stationary cameras mounted at a remote environment are used to acquire both photometric and depth information. A virtual environment is then constructed in real time and redisplayed in accordance with the local participant's head position and orientation. This allows local users to interact with other remote individuals (say, with a medical consultant) as if they were actually within the same space.

The evolution of multimedia applications can be traced through three major stages. First, even prior to the deployment of delivery-network infrastructure, stand-alone applications (e.g., video arcade games) and CD-ROM-based applications had successfully integrated multiple media, mostly in the form of games, entertainment, and educational materials.

Next, high expectations of technological breakthroughs to make available lower-cost delivery bandwidth (as needed for sophisticated multimedia applications) created a tremendous excitement in the area of multimedia. The potential for the convergence of multiple services (e.g., TV, movie, and telephone) had stirred up the marketplace. Almost every day, newspaper headlines announced new field trials and potential mergers of corporations. The focus in multimedia shifted to the creation of large-scale video servers and delivery infrastructure that would be capable of delivering thousands of simultaneous high-quality video streams to homes and businesses. However, the economics of the marketplace has run counter to these high expectations. For example, in movie-on-demand applications, the cost of storage of a large video library and the cost of delivery

©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 application environment.

bandwidth for a two-hour video per customer were found to be prohibitively expensive in comparison with traditional movie rentals. In addition, the infrastructure for delivering high-quality video was mostly unavailable, particularly to the home. This issue is referred to as the *last-mile bandwidth problem*. Compounding this problem, most available multimedia contents were created as movies. Video-on-demand applications lacked the depth of content (complexity of presentation) and provided only limited interactivity (e.g., VCR control operations) as compared to Web-based applications. New forms of content have to be created to provide marketable services to users (e.g., distance learning, travel, advertisement, electronic commerce). Furthermore, sophisticated tools must be available for easy creation of such multimedia contents.

The popularity of the Web, which can use very low-bandwidth networks in the last mile, and the lack of deployment of large-scale video servers have laid the groundwork for the current stage of multimedia. Applications using high-bandwidth multimedia streams will be deployed in enterprise (high-bandwidth-network) environments and more slowly (as higher-bandwidth infrastructure is deployed) for residential consumers. Such services will likely follow the near-video-on-demand (NVOD) model [2] in which a large number of clients can share a set of broadcast channels, rather than the video-on-demand (VOD) model, which requires an individual stream per client. For example, the Digital Satellite System (DSS) broadcasts (via satellite) a large number of entertainment streams to its paid subscribers. In time, these services will evolve to include NVOD-like features;

e.g., VCR control features will be provided by the mechanism of copying a broadcast stream to the local disk of a client or of switching to a different broadcast channel with later playback of the same material. (Of course, any local copying must take into account the protection of intellectual property rights.)

The primary focus of current research activities, however, is on the creation of multimedia content that can be delivered inexpensively by means of the existing low-bandwidth networking infrastructure as well as on the development of platform-independent applications that can adapt to heterogeneity in environments arising from the differences in capabilities of system components and end-user devices. The new forms of content fall into two categories: low-bandwidth, real-time conferencing, and delivery of composite multimedia documents consisting of images, text, and possibly short audio and video clips. The quality of the presentation will change (i.e., dynamically adapt) with the availability of delivery bandwidth. The business and research focus has shifted from just video or audio quality to information delivery. Video-conferencing applications, for example, are replacing special video-conferencing rooms (which have high operational cost) used in many business environments, and bringing the desired functionality to the desks.

Challenges in multimedia-system design

Efficient multimedia systems are not just simple extensions of conventional computing systems. Delivery of *time-sensitive* information, such as voice and video, and handling of large volumes of data require special considerations to produce successful widespread applications. Consequently, all system components, e.g., storage, operating systems, and networks, must support these additional requirements. This results in very different architectural and design decisions in all subsystems and their services. More specifically, some of the important issues that have to be addressed when one is designing a multimedia system are

- Storage organization and management.
- Available physical bandwidth in the delivery path to the users.
- Quality-of-service (QoS) management (real-time delivery and adaptability to the environment).
- Information management (indexing and retrieval).
- User satisfaction.
- Security (especially management of content rights).

A generalized environment supporting multimedia applications to homes or to business environments is shown in **Figure 1**, in which U1 represents the user at home having access to an application or service through an access network, while a business user (U2) is connected

through a local area network, i.e., a customer-premises network. A multimedia server delivers services (i.e., multimedia data streams) over a series of (possibly) heterogeneous networks to the clients. (Further processing can take place within the network, e.g., multicasting, transformation of content.) As the figure illustrates, the need to provide flawless service to users places significant requirements on a large number of components, which have to work well together. In order to cope with the complexity of several heterogeneous systems involved in the end-to-end service, one can systematically analyze the components in order to evaluate various options and the design challenges. The design issues are broadly classified into the following categories:

1. Server subsystems that store, manage, and retrieve the multimedia data stream(s) upon a user request.
2. Network subsystems that transport, deliver, adapt (and perhaps transform) the data streams isochronously (i.e., at a specified rate, without observable delay) to the clients.
3. Client subsystems that receive and/or prefetch the data streams and manage the presentation of data.
4. Application programs that deal with relationships among data frames and media segments, and manage user navigation and retrieval of this data.

Multimedia technology has to address all of the components in coordination, since each component has evolved individually up to now. The following sections describe some of the major challenges to be overcome for modern multimedia applications to become ubiquitous and as influential as we predict. The discussions are not intended to be exhaustive, but to refer to significant research activities that are described in this issue of the *IBM Journal of Research and Development* or are under way at IBM and elsewhere. For a broader introduction to and survey of the multimedia area, the readers should refer to a companion paper in this issue [3].

• *Server-design issues*

Multimedia servers store and manage multimedia documents and deliver data streams in an isochronous manner, in response to delivery requests from users. They may also process the stored information before delivery to users. In the future, multimedia servers may range from low-cost, PC-based simple servers that deliver a few streams, to scalable large (either centralized or distributed [4, 5]) servers that will provide thousands of streams. The content may range from long, sequentially accessed videos to composite documents consisting of a mixture of small multiple-media segments (e.g., video, audio, image, text).

The multimedia objects that are large and are played back sequentially require relatively large storage space and

playback bandwidth. Therefore, considerable initial research has been directed toward a) efficient retrieval of data from disks, in order to amortize seek time and to avoid *jitter* [6], b) *striping* of data across storage devices for load balancing [7], and c) scheduling of system resources to avoid jitter [8–10]. Retrieval of large data blocks is now a well-accepted practice in most commercial servers [7, 8]. The use of *double buffering* to improve bandwidth utilization without incurring jitter can be cost-effective [11]. While simple striping is effective for load balancing across a set of homogeneous storage devices for long sequential accesses, complex *block-allocation* policies that take into account bandwidth and storage capacities of the devices are necessary. An alternative approach is to partition the storage devices into sets of homogeneous striping groups. A higher-level policy assigns media objects to these devices, matching the storage and bandwidth requirements of the documents to the availability of those in the striping groups [12]. Another related problem is *dynamic load balancing* across servers or striping groups [13]. For example, in [14] a policy of temporarily replicating “hot” data segments (i.e., frequently accessed parts of the media objects) by copying retrieved data in playback streams is used to dynamically balance the load across striping groups. Such a dynamic policy is complementary to the initial data-placement policy.

The low-cost servers that are designed to serve a few dozen streams may be simple, PC-based servers and may not employ various resource-optimization techniques (other than retrieval of large data blocks). However, large-scale servers that can serve thousands of streams exploit the economies of scale and can be quite complex [8]. For example, in a video-on-demand server, multiple requests for the same popular video that arrive within a short time period can be *batched* for service by a single playback stream. Various video-scheduling policies take into account user satisfaction while improving effective server capacity [15]. Another optimization technique known as *bridging* speeds up the following streams while slowing down the preceding streams for the same video in order to achieve an eventual merging of these streams, and hence to avoid multiple playback streams [16]. The main memory buffer of the server can be used to cache retrieved data for later playback. A technique known as *interval caching* uses buffers to cache the short intervals between successive playback streams of the same video [17]. This caching can also be used for balancing changing loads across storage devices [17].

To guarantee QoS and avoid jitter, multimedia servers employ admission control; i.e., a new request is granted only if enough server resources are available. This implies that a video server has to know in advance the number of streams it can serve without jitter. The capacity of a video server can be expressed in statistical terms, e.g., the

Table 1 Network layers.

Multimedia application	(OSI 6-7)
Transport	(OSI 4-5-6)
Network	
Data-link control	
Physical	

number of streams it can serve with a certain level of jitter [11]. For a server complex comprising many different components, the capacity is expressed not by a single number but by means of a graph model with individual capacities of the components [8]. The admission control relies on finding a delivery path from the storage devices to the selected network interface and reserving appropriate capacities on all components on this path for normal playback. The admission control may also rely on resource-optimization techniques, e.g., batching and caching. VCR control operations (e.g., pause, resume, fast-forward) introduce new challenges for efficient use of resources. Resources may be released and reacquired upon resumption. However, it may not be possible to guarantee that resources are available upon resumption. For example, a new playback stream must be created upon resumption of a client stream that was originally served via batching. In large-scale servers, a small amount of capacity may be set aside as contingency capacity for dealing with VCR control operations [18].

In addition to addressing the above resource optimization issues, the next generation of multimedia servers must address the challenges in retrieval of *composite multimedia documents*. A composite multimedia document consists of many different parts (e.g., video and audio clips, image, text, and data), with very different storage and bandwidth requirements, that are to be presented in a synchronized fashion. Therefore, the bandwidth requirement varies over time. Fine-granularity, time-varying advanced reservation of resources along with prefetching of data can be used to improve resource utilization [19]. The issues must be addressed both at the clients and at the servers. The prefetch schedule for various media objects should also take into account the bandwidth fragmentation in a server where different components may be located in different storage devices.

- *Network challenges*

Network subsystems create delivery paths from the multimedia servers to the individual clients, and transport the media objects. To avoid long latency and the need for large client buffers, some media objects (e.g., large video and audio files) are streamed isochronously to the clients from the servers. [Note, however, that not all media delivery requires isochronous streaming, as is the case when the

user browses noncontinuous media (e.g., text, image).]

Table 1 illustrates the Open System Interface (OSI) layers and their services that should be addressed for end-to-end multimedia applications. Starting from the lowest layers, one finds challenges that are directly or indirectly related to providing higher quality of service to the user.

Networks, covering the lowest four layers, have to meet several challenges in the direction of multimedia services. New technologies are required to bring high bandwidth to the user. In business environments there is typically a LAN attachment for each user station, which provides shared or point-to-point bandwidth of the order of 10 Mb/s or better. However, the network to the home is currently limited to a few tens of Kb/s (or to 128 Kb/s in the case of ISDN), because of the current telephone network infrastructure. The evolution of technologies such as ADSL [20, 21] and HFC (CATV) [20, 21] will bring bandwidth of the order of Mb/s to the home, which is a requirement for the successful deployment of real-time multimedia to the home.

Raw bandwidth is not enough for effective delivery of multimedia services, especially when this bandwidth is shared among several systems or applications; it is necessary to provide mechanisms that achieve delivery of voice and video information isochronously, so that the corresponding network connections allow users to enjoy flawless delivery of media information. The general concept of QoS, which has been introduced to broadband ISDN, is used in multimedia applications. Although the term QoS is general, it implies all of the properties that are required for a flawless end-to-end service, i.e., delivery of information isochronously, so that the end user experiences a continuous flow of audio and/or video without interruption or noticeable problems of degraded quality.

Network technology and the communication protocol stack(s) used for multimedia information transmission clearly play an important role in achieving the above goals, i.e., to provide high bandwidth to the user with the appropriate QoS. The wide spectrum of technologies in all areas of networking—local, metropolitan, and wide-area—require the development of interconnecting systems (adapters, bridges, routers, switches, etc.), which support high-speed communication among heterogeneous or similar networks with low latency and low cost. The challenges in this area are numerous: minimization of data-transmission and data-reception overhead in order to preserve link bandwidth for applications [22, 23], low-latency protocol processing [22, 24, 25], and high-speed routing [26, 27], to mention a few. These problems originate from data networking; the stricter latency requirements of video and audio, in conjunction with the differing standards that may be adopted at the various networks, impose even stricter requirements. For

noncontinuous, media-based applications, latency is often far more important than jitter, since it may affect the interactive response time. For certain networking technologies (e.g., cable modem with asymmetrical bandwidth, using TCP/IP), this latency can be quite high.

Depending on the level and protocol stack used for the service, different problems may arise; for example, transmission over the Internet or intranets with the use of the currently widespread TCP/IP protocol suite does not guarantee in-time delivery of time-sensitive information, while MPEG** transmission with the use of the AAL-5 protocol over an ATM network requires that the network have bounded jitter [21].

Considering that one of the main targets of multimedia applications is provision of services to users in the same fashion as telephone service, e.g., in kiosks and public multimedia terminals, it is imperative that the network infrastructure support security to appropriate levels, in order to protect customers and ensure privacy. Tamper-free systems, efficient authentication mechanisms, and efficient and reliable transmission of encrypted information are a necessity in such environments. In environments where information is transmitted encrypted and/or encoded (for compression), special attention has to be paid to the recovery of lost information in the network.

- *Client-design issues*

The design of the client will, as always, be dominated by questions of cost, functionality, and comfort. There are several possible forms for the client device: an upgraded television set (with internal or external "set-top" capability), an optimized personal computer (a network computer with video capability), or a full-fledged personal computer. If the control device is shared by all access devices in the household as part of a residence gateway, higher functionality may be possible than if the circuitry must be replicated in each of the many TVs in a residence. The pressure to restrict the amount of RAM and video-processing power will continue for some years, but limiting buffer sizes and local computing will affect the load on the network. (Small buffers and low processing capability cannot handle highly compressed video, i.e., compression across multiple frames. Also, prefetching and/or double buffering are required to maintain the quality of the displayed image, i.e., jitter reduction, smoothing.) Many functions can be implemented most easily if there is significant local secondary storage, but servicing of moving parts, such as disks, in a home-video environment could be a detriment to sales.

The restrictions on local capabilities will certainly affect which functions are performed well locally and which depend on server or network support. Classic VCR control operations can be implemented in many different ways

[17, 18]; the tradeoffs depend crucially on the split of capability throughout the system. For example, the server can respond to the commands or switch from VOD to NVOD mode, an intermediate node can store entire streams, or the client can manage a cache.

There are a variety of physical methods for enabling user interaction, ranging from standard infrared remote controls to upgraded remote controls to traditional computer pointing devices (e.g., a mouse, TouchPoint). Occasionally, alphabetic or numeric inputs are required, at which time a real keyboard or keypad is desirable. In addition, audio and video inputs (for video conferencing) will be increasingly common, resulting in significant implications for upstream bandwidth and control.

Apart from the above hardware considerations, decisions about the software interfaces, control paradigms, and operating system, such as which functions will be static and which will be loaded with the content or upon need, will have significant impact on cost and usability of the client. A fixed set of capabilities is easy to manage but restricts future upgrades and improvements in types of content that can be viewed. As multimedia applications evolve from simple replay to the inclusion of complex search [28-30] and simultaneous streams [19, 31-33] and further to personalized environments, the client devices (e.g., set-top box, camera, microphone) are likely to be difficult to use. This will result in challenges with regard to optimizing resources and avoiding misbehavior and mischief, just as with current personal computers.

- *Application-design issues*

Future multimedia applications will utilize valuable and enjoyable types of information, including multiple audio and video streams and time-, user-, and data-dependent behaviors. Multiple applications that operate on the same data will have to manage information, analyze images, search for appropriate items, and mesh live, pre-recorded, and interactive streams. The software framework must make it easy for application developers to create software that satisfies these and other needs.

Creating advanced content will become increasingly challenging. It has always been difficult to produce videos that are pleasant to watch. Advanced tools (e.g., for authoring multimedia content [31, 32]) will be required to make it possible to combine multiple streams in attractive ways that have maximum impact. However, addressing human perceptions and psychological effects of a composition will always remain the realm of creative artists.

Searching image, audio, and video data is far more difficult than doing the same for text, but the rewards can be correspondingly higher. The adage "A picture is worth a thousand words" suggests not only the higher value that people put on multimedia data, but also the higher cost of

image and video data. Approaches to finding desired information include creation and analysis of text annotations [28, 29], examining reduced versions of the video (thumbnails or key frames) [28], and full-scale analysis of high-resolution data [30]. Practical considerations force us to maximize the value received from a small amount of computation. The requirements become particularly onerous in interactive and real-time situations such as push technology [34], browsing, and interactive navigation in large sets of multimedia data [30].

Presenting the information is also difficult, especially when system and network capabilities are limited. "Best-effort" networks (with no guarantee of on-time delivery) create challenges for synchronizing audio and video substreams and require adaptation to feedback about user behavior as well as network and server performance. The adaptation may take into account QoS measurements by the operating system and other system components [35], heterogeneous media [36], and heterogeneous environments [37]. Managing controlled collaborations and teleconferences with many audio, video, and data streams is even more challenging [38, 39]. In the future, people will want even more complicated composite moving media, with complex constraints on location and timing of the various parts [31–33].

Interpretations of media data can also vary from user to user. Thus, there is a need for a framework to accommodate views of different users on the same data and to provide mechanisms for creating and updating views. The BRAHMA [29] framework proposes an approach for annotating media data (i.e., creating views) via user-defined attribute-value pairs for hierarchical organization and retrieval of media data.

Finally, virtual views can be created from a set of image data. Conventional *virtual reality* and 3D applications, including those based on virtual reality markup language (VRML), allow continuous movement of the viewpoint but often suffer from poor performance, crude 3D models, large data files, low-quality rendered images of synthetic scenes, and the need for expensive hardware. PanoramIX [40] alleviates these problems by employing image-based rendering algorithms. Although the viewpoint is restricted to discrete points in space, the "look-around performance" is much better, the resulting images are very realistic, the data files are relatively small, and no special hardware is needed. In addition, PanoramIX provides a composite-media environment for embedding and hyperlinking disparate media types, such as streaming audio, video, 1D–3D sprites, and synthetic 3D objects [40].

Summary and conclusions

Up to now, the lack of infrastructure for delivery of high-quality video, the weakness of supporting technologies for

easy content creation, the absence of experience with the benefits of futuristic multimedia, and a variety of economic factors have impeded the widespread deployment of multimedia applications. Experience with video on the Web and availability of the Internet for disseminating low- and medium-bandwidth products are breaking down these barriers. People now accept the enormous value of widespread use of multimedia technologies to support societally desirable, financially profitable, and personally rewarding services such as distance-independent learning, low-cost video telephone service, improved customer-relationship and help services, new forms of entertainment and interactive games, and new types of electronic commerce (such as tourism, real-estate, and merchandise sales). There are still many technology barriers, but the papers in this issue of the *IBM Journal of Research and Development* demonstrate that we are making significant progress at overcoming them.

Acknowledgments

We thank Marc Willebeek-LeMair, Nelson Manohar, Jai Menon, William Pence, Junehwa Song, and anonymous referees for their helpful comments and feedback.

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

References

1. R. Azuma, "A Survey of Augmented Reality," *Presence: Teleoperators & Virtual Environments* **6**, No. 4, 355–385 (August 1997).
2. Kevin C. Almeroth and Mostafa H. Ammar, "The Use of Multicast Delivery to Provide a Scalable and Interactive Video-on-Demand Service," *IEEE J. Selected Areas in Commun.* **14**, No. 6, 1110–1122 (August 1996).
3. R. J. Flynn and W. H. Tetzlaff, "Multimedia—An Introduction," *IBM J. Res. Develop.* **42**, No. 2, 165–176 (1998, this issue).
4. D. Serpanos and A. Bouloutas, "Centralized vs. Distributed Multimedia Servers," *Research Report RC-20411*, IBM Thomas J. Watson Research Center, Yorktown Heights, NY, March 1996.
5. D. M. Dias, B. Kish, R. Mukherjee, and R. Tewari, "A Scalable and Highly Available Web Server," *Proceedings of IEEE COMPCON: Technologies for the Information SuperHighway*, February 1996, pp. 85–92.
6. P. S. Yu, M. S. Chen, and D. Kandlur, "Grouped Sweeping Scheduling for DASD Based Multimedia Storage Management," *ACM Multimedia Syst.* **1**, 99–109 (1993).
7. M. G. Kienzle, R. R. Berbec, G. P. Bozman, C. K. Eilert, M. Eshel, and R. Mansell, "Multimedia File Serving with the OS/390 Lan Server," *IBM Syst. J.* **36**, No. 3, 374–392 (1997).
8. A. Dan, S. Dulin, S. Marcotte, and D. Sitaram, "Multimedia Resource Management in OS/390 LAN Server," *IBM Syst. J.* **36**, No. 3, 393–410 (1997).
9. D. P. Anderson, "Metascheduling for Continuous Media," *ACM Trans. Computer Syst.* **11**, No. 3, 226–252 (1993).
10. K. K. Ramakrishnan, L. Vaitzbilt, C. Gray, U. Vahalia, D. Ting, P. Tzelnic, S. Glaser, and W. Duso, "Operating

- System Support for a Video-on-Demand File Service," *ACM Multimedia Syst.* 3, No. 2, 53-65 (1995).
11. A. Dan, D. Dias, R. Mukherjee, D. Sitaram, and R. Tewari, "Buffering and Caching in Large Scale Video Servers," *Proceedings of IEEE COMPCON*, San Francisco, 1995, pp. 217-224.
 12. A. Dan and D. Sitaram, "An Online Video Placement Policy Based on Bandwidth to Space Ratio (BSR)," *Proceedings of ACM SIGMOD'95*, San Jose, May 1995, pp. 376-385.
 13. D. Serpanos, L. Georgiadis, and A. Bouloutas, "MMPacking: A Load and Storage Balancing Algorithm for Distributed Multimedia Servers," *IEEE Trans. Circuits & Syst. for Video Technol.* 8, No. 1, 13-17 (February 1998).
 14. A. Dan, M. Kienzle, and D. Sitaram, "Dynamic Policy of Segment Replication for Load-Balancing in Video-on-Demand Servers," *ACM Multimedia Syst.* 3, No. 3, 93-103 (July 1995).
 15. A. Dan, D. Sitaram, and P. Shahabuddin, "Dynamic Batching Policies for an On-Demand Video Server," *ACM Multimedia Syst.* 4, No. 3, 112-121 (June 1996).
 16. L. Golubchik, J. C. S. Lui, and R. Muntz, "Reducing I/O Demand in Video-On-Demand Storage Servers," *Proceedings of ACM SIGMETRICS*, 1995, pp. 25-36.
 17. A. Dan and D. Sitaram, "Multimedia Caching Strategies for Heterogeneous Application and Server Environment," *Multimedia Tools & Appl.* (special issue: The State of the Art in Multimedia) 4, No. 3 (1997).
 18. A. Dan, P. Shahabuddin, D. Sitaram, and D. Towsley, "Channel Allocation Under Batching and VCR Control in Video-On-Demand Servers," *J. Parallel & Distr. Computing* 30, No. 2, 168-179 (November 1995).
 19. J. Song, A. Dan, and D. Sitaram, "Efficient Retrieval of Composite Multimedia Objects in the JINSIL Distributed System," *Proceedings of ACM SIGMETRICS*, Seattle, June 1997, pp. 260-271.
 20. *IEEE Communications Magazine* (special issue: Reshaping Communications Through Technology Convergence) 33, No. 4 (April 1995).
 21. *IEEE Network Magazine* (special issue: Digital Interactive Broadband Video Dial Tone Networks) 9, No. 5 (September/October 1995).
 22. L. Peterson and B. S. Davie, *Computer Networks: A System Approach*, Morgan Kaufmann Publishers, San Francisco, 1996.
 23. H. Meleis and D. Serpanos, "Designing Communication Subsystems for High Speed Networks," *IEEE Network Magazine* 6, No. 4, 40-46 (July 1992).
 24. C. Partridge, *Gigabit Networking* (Addison-Wesley Professional Computing Series), Addison-Wesley Publishing Co., Reading, MA, 1993.
 25. D. Serpanos, "Protocol Processing in High-Speed Communication Subsystems," *Proceedings of High Speed Networks '93, the 5th International Conference on Data Communication Systems and Their Performance*, Raleigh, NC, October 1993, pp. 351-360.
 26. A. Brodnik, S. Carlsson, M. Degermark, and S. Pink, "Small Forwarding Tables for Fast Routing Look Up," *Proceedings of ACM SIGCOMM*, September 1997, pp. 3-14.
 27. D. Serpanos, "Multiprocessing in Multiprotocol Routers," *J. Microcomputer Appl.* 17, 99-112 (1994).
 28. R. M. Bolle, B.-L. Yeo, and M. M. Yeung, "Video Query: Research Directions," *IBM J. Res. Develop.* 42, No. 2, 233-252 (1998, this issue).
 29. A. Dan, D. Sitaram, and J. Song, "BRAHMA: Browsing and Retrieval Architecture for Hierarchical Multimedia Annotation," *Proceedings of Multimedia Information Systems*, West Point, NY, September 1996, pp. 25-29.
 30. 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, No. 2, 253-268 (1998, this issue).
 31. M. Kim and J. Song, "Multimedia Documents with Elastic Time," *Proceedings of the 3rd ACM International Multimedia Conference*, San Francisco, 1995, pp. 143-154.
 32. M. Cecelia-Buchanan and P. T. Zellweger, "Scheduling Multimedia Documents Using Temporal Constraints," *Proceedings of the 3rd International Workshop on NOSDAV*, November 1992, La Jolla, CA, pp. 237-249.
 33. N. R. Manohar and A. Prakash, "The Session Capture and Replay Paradigm for Asynchronous Collaboration," *Proceedings of the European Conference on Computer Supported Cooperative Work*, Stockholm, September 1995, pp. 149-164.
 34. PointCast Home Page, <http://www.pointcast.com>.
 35. N. R. Manohar and A. Prakash, "Dealing with Synchronization and Timing Variability in the Playback of Interactive Session Recordings," *Proceedings of the 3rd ACM International Multimedia Conference*, San Francisco, 1995, pp. 45-56.
 36. N. R. Manohar, M. Willebeek-LeMair, and A. Prakash, "Applying Statistical Process Control to the Adaptive Rate Control Problem," *Proceedings of IS&T/SPIE 10th Annual Symposium, Multimedia Computing and Networking Conference*, San Jose, 1998, pp. 42-56.
 37. M. Naghshineh and M. Willebeek-LeMair, "End-to-End QoS Provisioning in Multimedia Wireless/Mobile Networks Using an Adaptive Framework," *IEEE Communications Magazine*, pp. 72-81 (November 1997).
 38. M. H. Willebeek-LeMair, K. G. Kumar, and E. C. Snible, "Bamba—Audio and Video Streaming Over the Internet," *IBM J. Res. Develop.* 42, No. 2, 269-280 (1998, this issue).
 39. 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, No. 2, 281-298 (1998, this issue).
 40. W. L. Luken, J. S. Lipscomb, K.-S. Cheng, and J. P. Menon, "PanoramIX: Photorealistic Multimedia 3D Scenery," *Research Report RC-21145*, IBM Thomas J. Watson Research Center, Yorktown Heights, NY, 1998.

Received January 26, 1998; accepted for publication February 2, 1998

Asit Dan *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 704, Yorktown Heights, New York 10598 (asit@us.ibm.com)*. Dr. Dan received a B. Tech. degree in computer science and engineering from the Indian Institute of Technology, Kharagpur, India, and the M.S. and Ph.D. degrees from the University of Massachusetts, Amherst, in computer science and computer engineering, respectively. His doctoral dissertation on performance analysis of data sharing environments received an Honorable Mention in the 1991 ACM Doctoral Dissertation Competition and was published by the MIT Press. He has published extensively on the design and analysis of video servers as well as transaction-processing architectures, and served as the tutorial chair for SIGMETRICS '97 and various other program committees. He is a major contributor to various IBM products in the above areas, holds several top-rated patents, and has received two IBM Outstanding Innovation Awards and five IBM Invention Achievement Awards. Currently, Dr. Dan is leading a project on the development of COYOTE middleware for supporting long-running transactional applications across autonomous business organizations.

Stuart I. Feldman *IBM Research Division, Thomas J. Watson Research Center, P.O. Box 704, Yorktown Heights, New York 10598 (sif@watson.ibm.com)*. Dr. Feldman received an A.B. degree in astrophysical sciences from Princeton University and a Ph.D. in applied mathematics from the Massachusetts Institute of Technology. He spent ten years in computer science research at Bell Laboratories as a member of the original UNIX research team. He spent the next 11 years at Bellcore, where he held several research management positions in software engineering and computing systems. He joined IBM in 1995 at the Thomas J. Watson Research Center as Department Group Manager of Network Applications Research. He has recently been named director of the newly established IBM Institute for Advanced Commerce in Hawthorne, New York, which brings together leaders in business and academia to research the impact of emerging technologies on the future of business and commerce. Dr. Feldman is an ACM Fellow, an IEEE Fellow, and a member of Phi Beta Kappa and Sigma Xi.

Dimitrios N. Serpanos *Department of Computer Science, University of Crete, P.O. Box 1470, GR-71110 Heraklion, Crete, Greece (serpanos@csd.ucl.gr)*. Prof. Serpanos received a Diploma in computer engineering and information sciences from the University of Patras, Greece, in 1985 and a Ph.D. in computer science from Princeton University in 1990. He joined the IBM Research Division at the Thomas J. Watson Research Center in 1990 as a Research Staff Member, working on high-bandwidth systems and interactive media solutions. In 1996 he joined the faculty of the Department of Computer Science of the University of Crete in Greece; he also works at the Institute of Computer Science of the Foundation for Research and Technology-Hellas (ICS-FORTH). His research interests include high-speed interconnections, high-speed communication systems, and multimedia systems. Prof. Serpanos has received awards from IBM, holds two U.S. patents, and has published several papers. He is a Senior Member of the Institute of Electrical and Electronics Engineers (IEEE) and a member of the Association for Computing Machinery (ACM), the New York Academy of Sciences, and the Technical Chamber of Greece.