INTERACTIVE VIDEO

It is estimated that the average American watches an average of four hours of TV per day. If you are one of these TV fans, your viewing habit is about to be revolutionized by the forthcoming interactive TV or interactive video (ITV) system. ITV is a promising interactive digital audiovisual service spawned by technological advances in information retrieval, consumer electronics, data compression, communications, and computing and by the convergence of the communication, computer, and entertainment industries. ITV application and service possibilities are endless. An ITV system will provide movies on demand, courses on demand, news on demand, home shopping, interactive video games, remote learning, as well as conventional television and telephone services. In short, virtually all services and applications of the stored audiovisual streaming types are potential applications of an ITV system. As a result, ITV can be defined as the display of video controlled by the viewer with a certain degree of interactivity to enable both intraprogram and interprogram decisionmaking (1). The terms *video on demand* (VoD) and ITV will be used interchangeably throughout this article.

VoD (2–8) combines information retrieval technologies with TV display and a remote control and enables customers to enjoy the functions of VCR machines with rented video tapes, plus the service quality and convenience of cable TV (or broadcast TV) services. Additional functions may be available, such as choosing the camera angle in a sports program, or allowing the user to customize the development of the story so that different users may watch different versions of a movie. VoD aims to provide individual viewers with the ultimate flexibility of interactive video. Users can select programs from massive remote video archives via an on-screen menu, view at the time they wish without leaving the comfort of their home, and interact with the programs via laser disc or VCR-like functions, such as indexing, fast forward, and rewind. If a VoD system allows users to access any video, at any time, and to perform any VCR-like user interactions, it is a true VoD system; otherwise, it is a near VoD system (9). Pay per view, for example, is an application of near VoD services.

Figure 1 depicts a VoD system from the user's perspective. The user interacts with the VoD system with a small intelligent set-top box via a remote control. (The user may also interact with the system via a keyboard, mouse, and even natural language input in the future.) The user selects the video program from a remote server through an on-screen menu provided by the service provider and interacts with the program via VCR-like functions. The transport network delivers digital video programs to the customers from the remote video archives. This network may be a telephone network, cable TV system, satellite system, wireless network, local area network, or even the Internet.

There are four major players in the provision of VoD services: service provider, network provider, content provider, and the subscriber. The content provider provides the video programs. The subscriber generates requests for service to the service provider, which will obtain the necessary material from the program providers and deliver it to the user on the

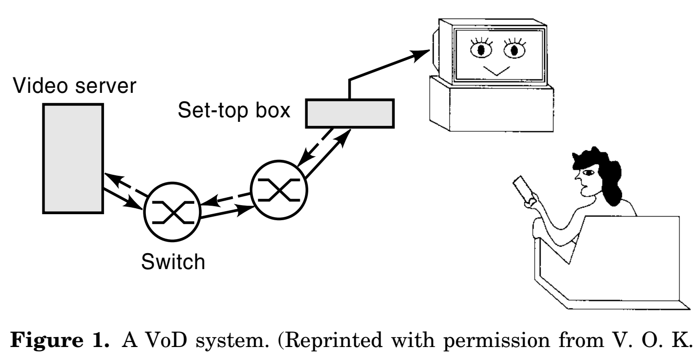

Figure 1. A VoD system. (Reprinted with permission from V. O. K. Li and W. J. Liao, Distributed Multimedia Systems, *Proc. IEEE,* **85** (7): 1063–1108, 1997. © 1997, IEEE.)

facilities of the network provider. Thus the service provider acts as a content broker. It is possible that the network, program, and service providers are the same organization, but, in general, they will be distinct. In fact, anyone with marketable materials can offer services to other users through the service provider. The relationships between the four players are illustrated in Fig. 2.

Digital videos are typically very large in volume and require real-time and continuous delivery for meaningful presentation. With entertainment-oriented VoD services, high-fidelity video and audio are required. This will, in turn, be translated into high transmission rate and huge storage volume requirements. For example, VHS quality video requires MPEG-1 compressed video streams at a transmission rate of 1.5 Mbps, while broadcast TV quality video requires MPEG-2 compressed video streams at a transmission rate of up to 6 Mbps. Without compression, existing computer platforms, storage devices, and networks are unable to satisfy the requirements of the massive storage space, high data transfer rate, and huge transmission bandwidth typical of multimedia data. The characteristics of VoD are summarized as follows:

1. VoD is a write-once-read-many application.
2. VoD systems deal with stored, compressed video rather than live video, as in videoconferencing services.
3. VoD systems demand quality of service (delay, loss, delay variation) guarantees from the system.
4. VoD systems require asymmetric transmission in the sense that the downstream (from the video server to the user) network bandwidth requirements are much higher than the upstream (from the user to the video server) requirements. The downstream transmission delivers video programs from the server, whereas the upstream transmission carries the control signals from the user.
5. VoD systems support user interactivity.

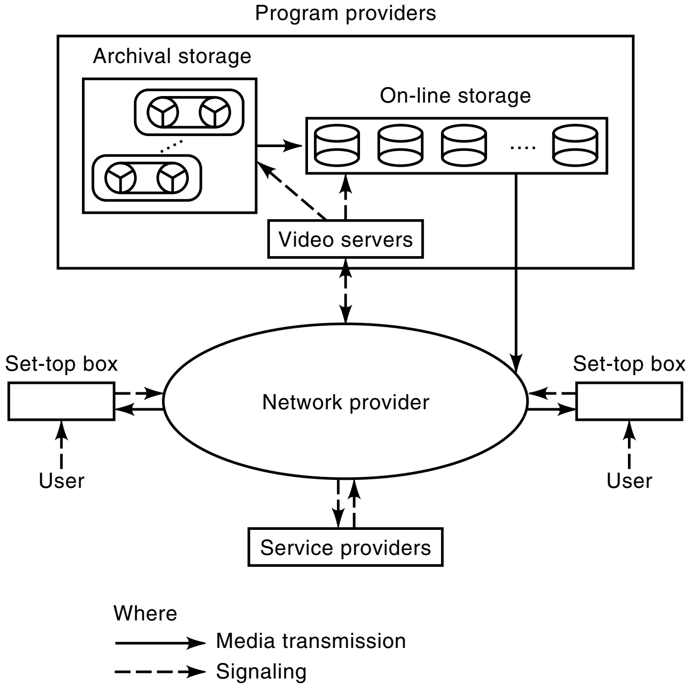

Figure 2. The relationships among the four key players in VoD services.

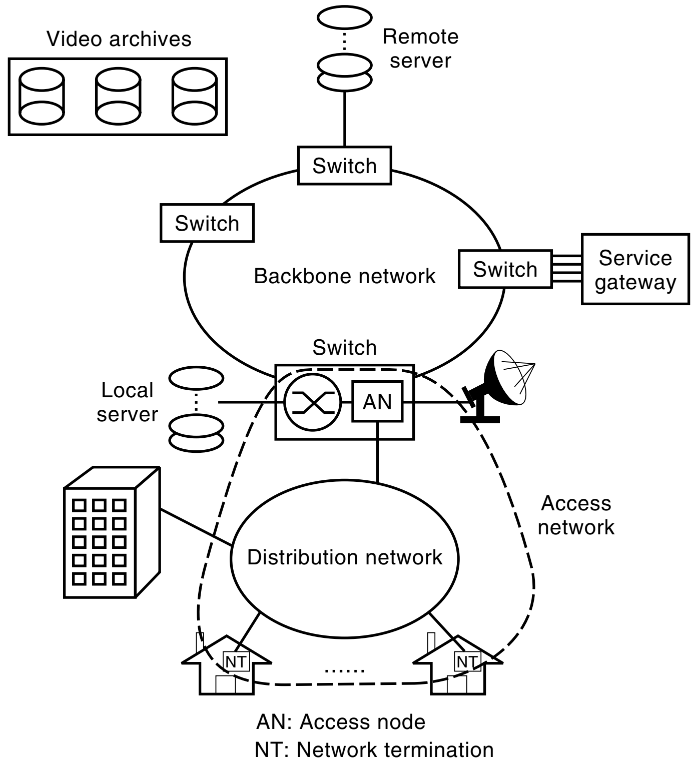

Figure 3. VoD system architecture.

The rest of this article will identify the important issues in the realization of VoD services, including the system architecture, storage management, delivery protocols, standardization efforts, and future developments.

VIDEO-ON-DEMAND SYSTEM INFRASTRUCTURE

A VoD system consists of the video server, the transport network, the service gateway, and the set-top box. Note that the service gateway can either be an independent unit or integrated with the network component. For the transport network, the focus will primarily be on a residential VoD system. Therefore, the transport network will be further decomposed into two parts: backbone network and access network. Figure 3 shows the typical architecture of a VOD system, and Fig. 4 indicates the relationships among the four components.

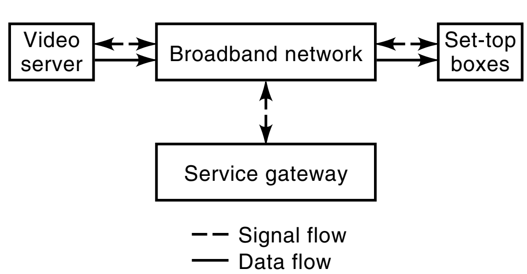

Figure 4. The four components of a VoD architecture. (Reprinted with permission from V. O. K. Li and W. J. Liao, Distributed Multimedia Systems, *Proc. IEEE,* **85** (7): 1063–1108, 1997. © 1997, IEEE.)

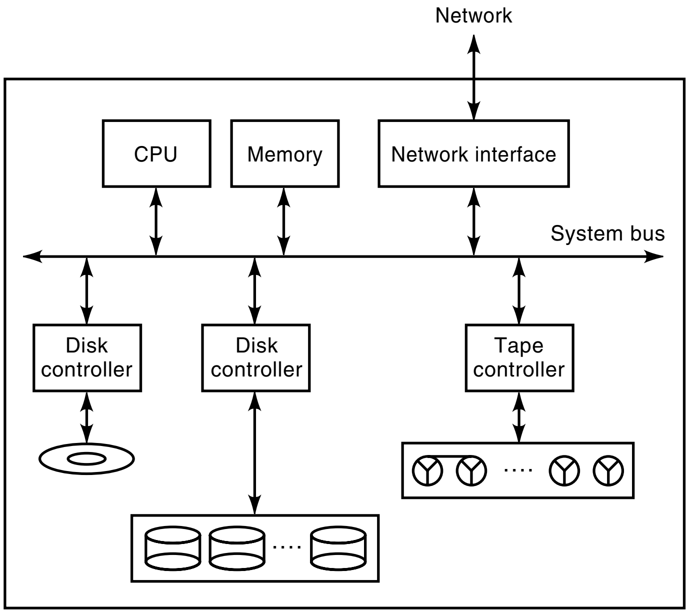

Figure 5. A video server architecture. (Reprinted with permission from V. O. K. Li and W. J. Liao, Distributed Multimedia Systems, *Proc. IEEE,* **85** (7): 1063–1108, 1997. © 1997, IEEE.)

Video Server

A video server, as shown in Fig. 5, consists of massive storage and media controllers to store a large number of movies and serves a considerable number of simultaneous video requests to the same or to different videos on demand. The storage media usually consists of magnetic disks (on-line storage), optical jukeboxes (near-line storage), and magnetic tapes (off-line storage) and is usually organized hierarchically for cost-effectiveness. Under this configuration, popular videos are stored on disks, and less popular ones in tape devices with terabyte capacities (retrieved as necessary to the disk drives for processing). Another coarse-grained storage hierarchy is on the server itself. The video server may be placed at the local switch or regional switch of the network provider or may be placed at remote information archives. Servers located remotely allow several service areas to share the same video server at the expense of transmission costs. Servers located locally allow the video programs to be in close proximity to the subscribers and therefore save transmission cost at the expense of coordination complexity and storage costs. Li and colleagues (10) proposed a performance model that may be used to evaluate the requirements of network bandwidth and server storage, and hence the tradeoff between communication and storage costs, for various placement alternatives. Finally, the video program migration within the storage hierarchy and the server location hierarchy should be carefully planned so that the performance requirement of real-time video delivery can be supported.

Transport Network

The transport networks deliver video programs from the video server to the customers with a guaranteed level of quality. The network must have very high bit rate to satisfy the real-time delivery constraints of video traffic. The transport network of a residential VoD system typically consists of two major parts: the backbone network with high-speed switches, and the local access network. The backbone network provides switched connections between remote video servers at geographically dispersed locations and the end users, via access networks. The backbone network provides reliable transfer of information between remote servers and the end users, addressing, establishment and termination of a connection, and network management for network configuration, performance and fault monitoring, and billing and accounting purposes. The trend is toward a Synchronous Optical Network (SONET) backbone with Asynchronous Transfer Mode (ATM) switching and multiplexing, because of their low error rate, high data transfer rate, bandwidth on demand, and seamless services.

The access network consists of the access node, the network termination, and the distribution network. The primary functions performed by the access network are (1) transmission, multiplexing, concentration, and broadcasting of information flows between the backbone network and the customer premises; (2) relevant control and management for the services; and (3) the transport of other services such as conventional telephone and television services.

The access nodes are located at the local switches and glue the backbone network to the distribution networks. Depending on the system implementation, an access node may be the head-end in cable TV (CATV) networks, the central office in telephone networks, or the base station in mobile systems (11). The access nodes may be equipped with satellite dishes to receive analog broadcast TV programs, have a local media server to store digital programs, and serve as interfaces to the backbone network to access information from remote video servers. In short, the access node performs the adaptation between the backbone network and the distribution network and contains centralized functions responsible for processing information flows in preparation for delivery through the distribution network. The major functions performed by access nodes include modulating (for passband systems) or multiplexing (for baseband systems) incoming signals to the desired outgoing formats, and switching the streams for downstream transmissions.

The network termination (NT) terminates the medium and the transmission technology used in the distribution network and performs the signal adaptation to a different medium with a different transmission technology in the home network. For example, in the case of a satellite broadcast access network, the signal will be terminated in the network termination and put on the coaxial cable or twisted pairs in the customer premises. Depending on the functions performed inside the NT, two alternatives can be used: passive NT and active NT. A passive NT passively adapts to the home network, such as a passive splitter, whereas an active NT actively adapts, converts, and changes the medium and transmission technology from the access network to the home network.

The distribution network delivers video programs from the access node for the last few miles to the customers. The types of such networks range from telecommunication networks to satellite broadcast systems. This part of the network infrastructure demands the most investment from the service providers for service deployment. One of the most active debates in residential broadband deployment is the choice of distribution network architectures. Due to the existing network infrastructure, operational policies, and other historical reasons,

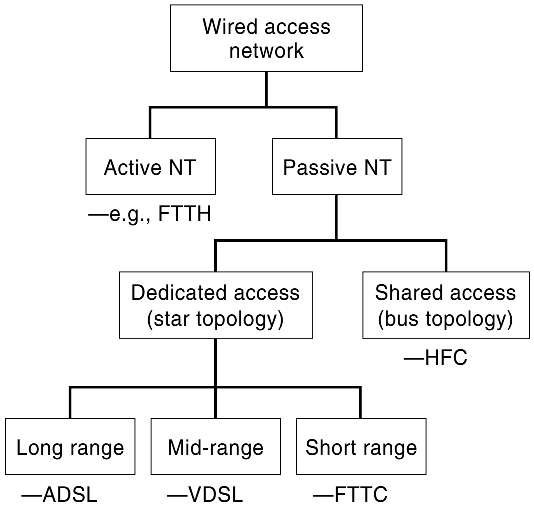

Figure 6. Wired access network classification. (Reprinted with permission from (12). © DAVIC 1997.)

the telephone and cable TV industries have proposed different technologies in order to capitalize on their existing network. This is the so-called last-mile debate.

Access network technologies can be classified by the transmission medium used. It can be either wired or wireless. Figure 6 shows the classification of wired networks, and Fig. 7 shows the classification of wireless networks (12). Five alternatives are presented next, with the first three being wired and the last two wireless.

Hybrid Fiber Coax. Hybrid fiber coax (HFC) is proposed by the CATV industry. It attempts to minimize initial investment by taking advantage of the existing CATV cable plants and combines high-bandwidth, low-loss fiber optics with low-cost coaxial cables, thereby allowing more channels with better picture quality to the customer (13). Figure 8 depicts a typical structure of an HFC system. This technology uses fiber trunks from the head-end to fiber nodes where optoelectronic conversion is performed, and continues with amplified tree-and-branch coaxial feeder cable to the homes in the neighborhood. The coaxial cable is shared by several homes and is fed through a passive NT to the customer premise. Each fiber node serves typically about 500 to 2000 subscribers in the neighborhood clusters, and each branch serves about 125 to 500 customers via coaxial cables.

Figure 9 shows a typical spectrum allocation of HFC. The spectrum allocation is determined by the service provider (14,15). The downstream (server to customer) frequency band

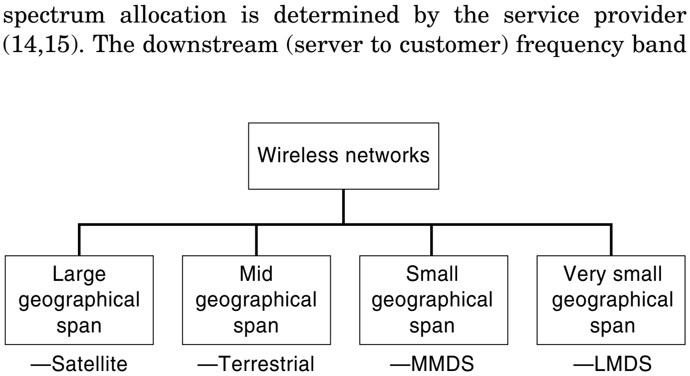

Figure 7. Wireless access network classification. (Reprinted with permission from (12). © DAVIC 1997.)

of 50 MHz to 550 MHz is allocated for analog TV broadcast and 550 MHz to 750 MHz for digital interactive services and downstream telephone services, while the upstream (customer to server) frequency band of 5 MHz to 45 MHz is allocated for return messages and upstream telephone and data services.

The cable TV network is designed for the broadcast of cable programs to the customers. To support two-way asymmetric communications for interactive digital services, there are two alternatives for the upstream transmissions: phone return and cable return. With phone return, the downstream transmission is by way of the regular cable plant, with the upstream interactive control signals via twisted-pair phone lines. The field trial conducted by US West and Continental Cablevision in the Boston area adopted phone return. With cable return, the interactive signals are transmitted via the same cable used in the downstream transmission, usually at 5 MHz to 45 MHz. Because cable bandwidth is shared by multiple users, a special multiple access control (MAC) protocol is required in the upstream direction to prevent collisions between users. In addition, reverse amplifiers are required for the upstream communications. Moreover, the subscribers share the same cable and therefore the available capacity for both the downstream and upstream bandwidth. It is necessary to have privacy and security safeguards. The IEEE 802.14 standard defines the physical and MAC layer protocols for HFC networks.

Asymmetric Digital Subscriber Loop. The telecommunication industry proposed asymmetric digital subscriber loop (ADSL) as their broadband access technology. It takes advantage of the high market penetration of telephone installation. This technology places ADSL modems at each end of the copper loop to create high-speed access link on existing telephone lines and enables multi-megabit per second transport of digital information over standard twisted-pair copper cables (16,17). In other words, ADSL offers a point-to-point connection between an ADSL modem located in the central office and the other ADSL model located in the customer premises. Figure 10 shows the ADSL network architecture.

The *asymmetric* in ADSL refers to the asymmetric bandwidth characteristic, with a high bit rate channel downstream toward the customer and a low bit rate channel upstream toward the network. This fits in nicely with the requirements of such client-server applications as VoD applications or World Wide Web (WWW) access, where the clients typically receive much more information (e.g., video programs) from the server than they are able to generate (e.g., control signals). ADSL allows downstream data rates ranging from 2 Mbps for 5 k/m loops, up to 7 Mbps for shorter loops, and typically 64 kbps to 640 kbps upstream. ADSL is typically used to cross long distances ranging from 1500 m to 5000 m. A variation called VDSL (very high bit rate digital subscriber line) operates downstream at OC1 (51.84 Mbps) data rate with a shorter range of 300 m to 1500 m. The most important feature of ADSL is that it can offer high-speed digital services on the existing twisted-pair phone lines, without interfering with the traditional analog telephone service.

Fiber to the Curb. Fiber to the curb (FTTC) takes advantage of the extremely high bandwidth of fiber optics and switched digital services. This technology uses fiber-optical connection from the telephone central office to the optical network units

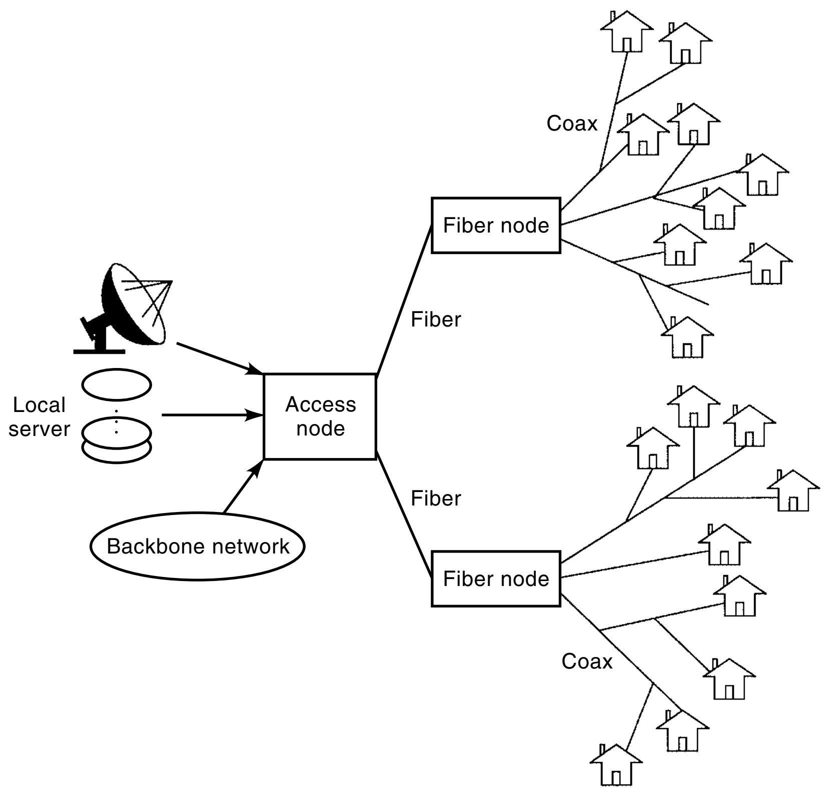

Figure 8. HFC network architecture. (Reprinted with permission from V. O. K. Li and W. J. Liao, Distributed Multimedia Systems, *Proc. IEEE,* **85** (7): 1063–1108, 1997. © 1997, IEEE.)

(ONU) (e.g., at the curbside or at the shared equipment within a building). From the ONU to the customer premises (i.e., the last drop, typically 300 m in length), twisted-pair or coaxial cable is used (18), as shown in Fig. 11. Special modulation techniques are required in the twisted-pair or coaxial cable to allow high transmission rates.

An ONU typically supports between 8 and 24 homes and is located somewhere between several tens to a few hundred meters from the customer premises. A variation called FTTB (fiber to the building) is used in cities with most of the population residing in high-rise buildings. Fibers are connected to the basements of such high-rises, with the rest of the paths to the individual flats over twisted-pairs. Another variation called FTTH (fiber to the home) provides point-to-point fiber connection from the central office directly to homes. The high cost and immature all-optical technologies, however, have delayed the deployment of FTTH.

Wireless Cable. The term *wireless* cable is a misnomer. It refers to the broadcast of cable TV type programming directly to receivers on the outside of customer premises via line-of-sight microwave transmissions (here *line-of-sight* means there is no obstruction between the transmitter and the receiver). Some telephone companies see this as a quick way to deliver video services to their customers. A large number of channels will be available, and some of the channels may be used to offer near VoD or true VoD services. Two technologies are available: multichannel multipoint distribution system (MMDS) and local multipoint distribution system (LMDS). Uni-directional MMDS can be treated as a wireless CATV network for video program broadcast. Bi-directional MMDS can provide two-way video, voice, and data transmissions. MMDS operates at around 2.5 GHz with a total bandwidth of 200 MHz (equivalent to 33 analog 6 MHz channels) and has been used by schools and archdioceses for about 20 years. With digitization, however, the bandwidth may be used to transport up to 150 digital channels. In MMDS, the access node performs the signal adaptation and transmits the signals to an antenna. This antenna may either be located in the access node or elsewhere, like the top of a hill. The antenna

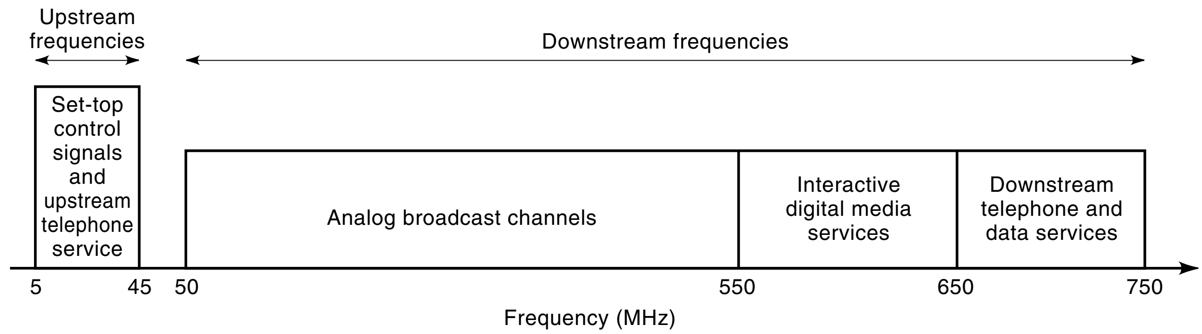

Figure 9. A typical spectrum allocation for a multiservice cable network.

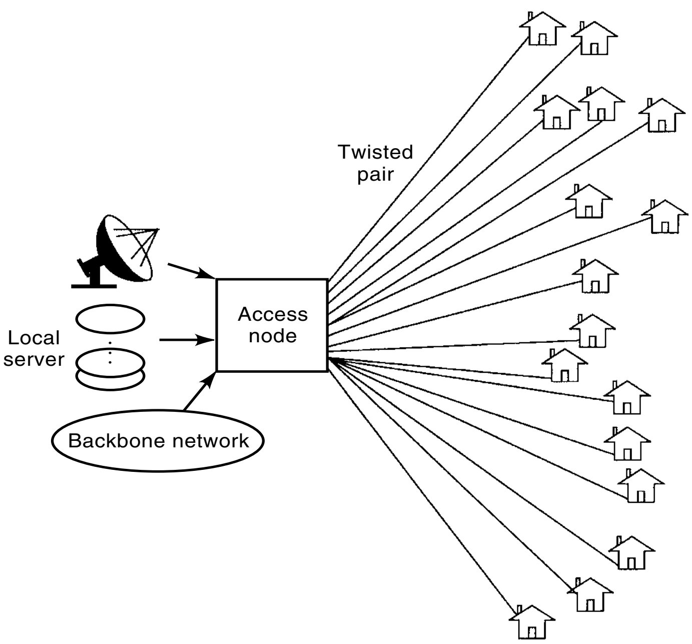

Figure 10. ADSL network architecture. (Reprinted with permission from V. O. K. Li and W. J. Liao, Distributed Multimedia Systems, *Proc. IEEE,* **85** (7): 1063–1108, 1997. © 1997, IEEE.)

then broadcasts the signals to the subscribers, with the typical broadcast radius in the order of 50 km. The customer premises must be equipped with an MMDS antenna, an RF (radiofrequency) transceiver, and a set-top box. The TV and set-top box then adapt the signals to match their internal format.

Just like MMDS, LMDS transmission may be one-way or two-way but operates at around 28 GHz with a much wider bandwidth of 1.3 GHz. Therefore, much more channels will be available.

Direct Broadcast Satellite. In a direct broadcast satellite (DBS) system (e.g. DirecTV™), digital audiovisual materials are sent by the service provider from ground stations via the uplink transmitter to satellites above the earth's surface. Each satellite has many transponders. Each transponder receives a video program on one frequency and retransmits it on another. The satellites broadcast the signals over a wide geographic region. Each user has a small antenna dish (18 inches in diameter for DirectTV) that is pointed in the direction of one of the satellites and receives video signals directly from it. As in wireless cable, a large number of channels will be available, and some of the channels may be used to offer near VoD or true VoD services. The main characteristic of DBS is its distributive (or broadcast) nature in the downstream direction. The information can be received by anyone with an appropriate satellite dish. A separate network must be employed to allow interactivity.

Note that these access networks may also be used to offer high-speed Internet access to the customers. Due to the strong demand for Internet services, some service providers have postponed the deployment of advanced interactive video services in order to focus on web services. It is expected that as more customers acquire high-speed Internet connections, the high-speed network infrastructure required for VoD services will be gradually developed.

Set-Top Box

A set-top box, along with the television monitor and the infrared controller (i.e., remote control), serves as the bridge between the subscribers and the system. With on-screen images and cursorlike devices, viewers may browse through video selections and interact with the video. The major functions of a set-top box include receiving the incoming video streams; demodulating, demultiplexing, and decoding the signals; performing the necessary signal conversion, such as digital-to-analog transformation for playback on the TV monitor; and sending outgoing control messages.

The key components of the set-top box, as shown in Fig. 12, include the central controller, line transceiver, back channel, demodulator, decoder, video buffer, display controller, and peripheral interface (8). The central controller consists of one or more microprocessors and memory system [e.g., read-only memory (ROM) and dynamic random-access memory

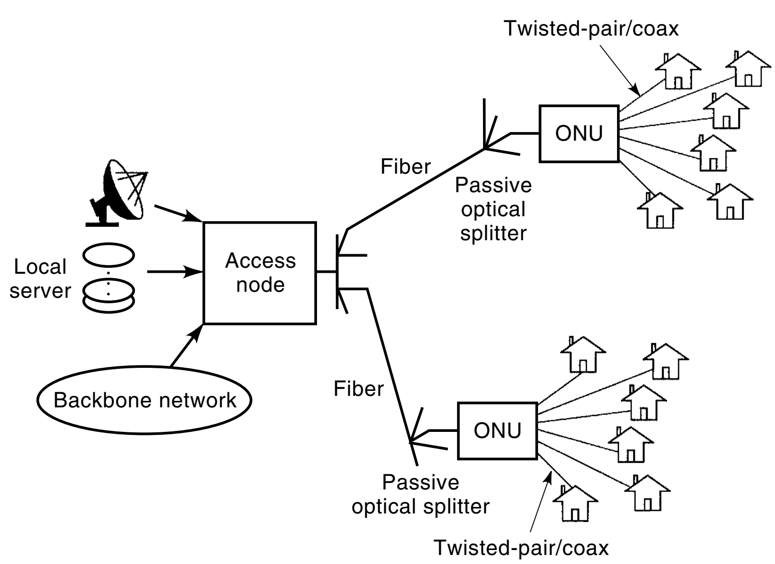

Figure 11. FTTC network architecture. (Reprinted with permission from V. O. K. Li and W. J. Liao, Distributed Multimedia Systems, *Proc. IEEE,* **85** (7): 1063–1108, 1997. © 1997, IEEE.)

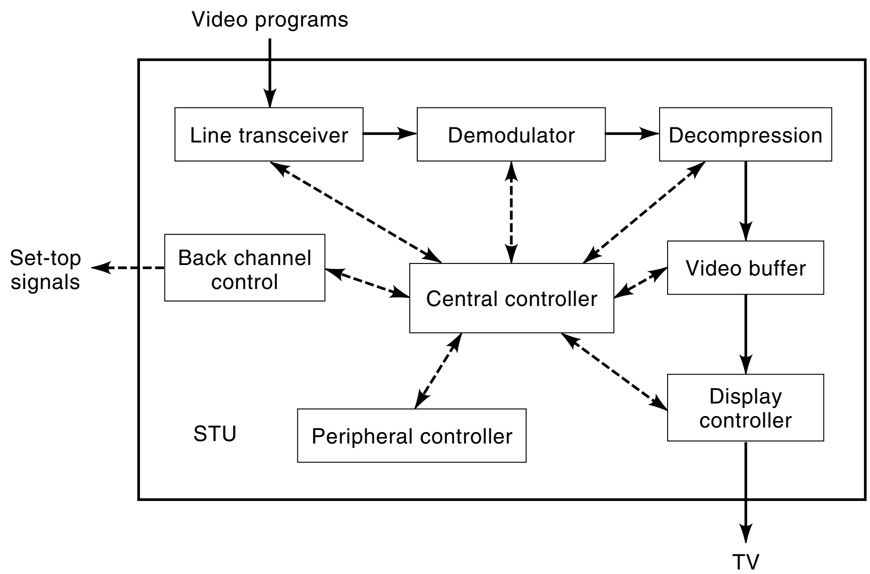

Figure 12. A set-top box architecture.

(DRAM)] and runs a real-time operating system (OS) to manage resources, activities, and components in the set-top box. The line transceiver and back channel, along with the demodulator, interface to the access network. The line transceiver is the entry point from which the incoming signals are received from the server. The back channel receives the returned control messages from the user and transmits them to the server. The incoming signals are demodulated to baseband, decompressed to reproduce the original video, queued in the video buffer, and then displayed on the video monitor. The peripheral controller allows the user to connect multimedia peripherals to the set-top box. Peripherals may include a VCR, telephone line, cable modem, CD-ROM drive, video camera, high-fidelity audio system, scanner, color printer, storage device, and remote control.

The set-top box must accommodate the heterogeneity of technologies and formats from various access networks and service providers and provides interactive controls to services (19). The ability of the set-top box to adapt to the diversity of access networks, service providers, applications, and user interfaces distinguishes it from the cable box currently used in franchised CATV that is dedicated to one cable company.

Service Gateway

This component may be integrated with an access node or may be a separate element in the network. The main functions performed by the service gateway include

1. Directory services to provide menu browsing and program scheduling
2. Mapping function from service ID to corresponding location and program provider
3. Control, coordinate, and signal multimedia session establishment, maintenance, and disconnection
4. System management, including operation management, fault management, configuration, resource management, and performance management
5. Subscriber profile maintenance and billing
6. Secured communication to prevent theft of service or unauthorized access, including authentication, encryption, and signal scrambling

STORAGE MANAGEMENT

The multimedia server (here the focus is on the continuous media server) must have (1) massive storage for huge amounts of multimedia information, (2) real-time storage and retrieval performance to satisfy the continuous retrieval of media stream at specified quality of service (QoS) guarantees, and (3) high data throughput to serve a large number of simultaneous user requests.

A buffer is allocated for each stream at the user site to overcome the mismatch between the data transfer rate and the playback rate of the media stream. Figure 13 shows the relationship between the frame retrieval, consumption (i.e., playback), and buffering. To avoid starvation during playback (i.e., there are no data available in the buffer to be played), the disk head must transfer another batch of data to the buffer before the previous batch has been consumed. This constraint implies that there is an upper bound on the number of concurrent requests a disk can handle. To ensure playback continuity, the simplest approach is to employ a dedicated disk head for each stream. The disk transfer rate is typically greater than the playback rate of an individual video stream.

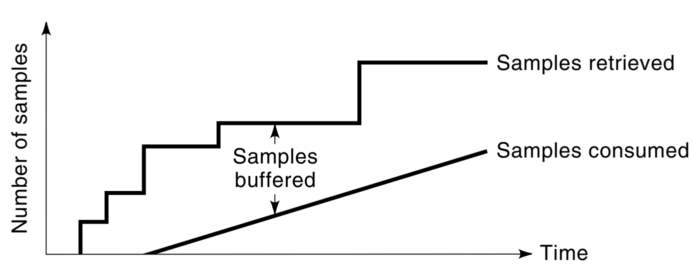

Figure 13. Relationships between data retrieval, consumption, and buffering.

This approach, therefore, wastes system resources, is not cost-effective, and limits the number of user requests to the number of disk heads. Another promising approach allows a disk head to be time-shared among multiple streams, employs real-time storage and retrieval techniques to satisfy the real-time constraint, and adopts admission control mechanisms to ensure that the QoS guarantees of the existing streams will not be degraded by accepted new requests. There are three major issues to be considered in providing such real-time support of continuous media: frame grouping, data placement, and real-time disk scheduling. Frame grouping reduces the access latency per frame by allowing a group of frames to be retrieved as an atomic unit. Data placement exploits the characteristic of continuous media and employs a specific disk layout structure to reduce seek and latency overhead and to gain timely access (20). Disk scheduling places no specific requirement on data organization but develops real-time disk scheduling with data buffers for real-time retrieval.

Frame Grouping

From a storage and retrieval viewpoint, it is advantageous to group video frames into segments. A video segment is an atomic unit of retrieval from disk. Since the disk retrieval overhead (seek and rotational latency) does not depend on the size of the video segment, a large segment improves disk retrieval efficiency. However, a large segment requires large buffers. Thus, there exists a tradeoff between disk retrieval efficiency and buffer requirement. Depending on the application, a certain segment length (defined as the number of frames per segment) may be preferred over others. When interframe coding is employed and when segment skipping is resorted to during fast playback, the problem of suitably grouping video frames into segments and that of skipping video segments during fast playback is a challenging one. Reference 21 describes grouping and retrieval schemes for Motion Picture Experts Group (MPEG) frames. These schemes describe how MPEG frames need to be grouped into segments of any desired segment length and which segments need to be skipped to achieve any desired fast playback rate. Several properties and closed form equations are provided.

Data Placement

Digitization of continuous media yields a sequence of frames. In the storage system, the continuous server groups the frames into segments and stores them on a disk. The goal of data placement is to optimize the disk layout structure by exploiting the characteristics of continuous media, such as continuity, write-once-read-many, and sequential retrievals, so that periodic and timely playback constraints can be satisfied. Example approaches include contiguous and constrained placement. Contiguous placement stores the successive segments of a media stream consecutively on a disk. This placement strategy takes advantage of sequential stream retrieval and therefore requires only one seek time to locate a particular stream. Increasing the segment size further improves the retrieval efficiency. Contiguous placement, however, suffers from fragmentation due to insertion and deletion. Thus, it is not suitable for read-write applications but works well with read-only ones. Constrained placement allows the successive segments of a media stream to be scattered so long as the intersegment separation is bounded by the constraint of playback continuity. The storage pattern, denoted (M, G), where M is the size of a data block and G is the intersegment separation, of this placement strategy is therefore determined by the playback rate of a stored stream and must satisfy the condition that the playback duration of a segment should exceed the time to skip over a gap and to retrieve the next segment (22,23). That is,

$$D \geq \frac{M + G}{T}$$

where D is the playback duration and T is the disk transfer rate. This placement strategy requires the disk to be sparsely occupied to begin with, so data can be stored in accordance to its storage pattern. Otherwise, an elaborate merging algorithm in which multiple media streams are interleaved should be applied to improve disk storage efficiency.

Disk Scheduling

Traditional disk scheduling is concerned with maximizing throughput by minimizing seek time and latency, or with fairness in usage by avoiding starvation. Such criteria are inappropriate with multimedia real-time constraints. Multimedia disk scheduling needs real-time guarantees. Example approaches include scan-earliest-deadline-first (Scan-EDF) (24), grouped sweeping scheduling (GSS) (25), quasi-static retrieval scheme (26), and Quasi-static Interactive Video Retrieval Protocol (QuIVeR) (27).

Scan-EDF combines Scan with EDF and provides the seek optimization of Scan and the real-time performance of EDF. Data are stored on concentric tracks on a disk, which rotates at a high speed. The disk read/write head is at the end of a disk arm, which moves back and forth just above the surface of the disk, in a direction perpendicular to the tracks. Basically, the disk head moves to the target track, waits for the data to rotate around until they are underneath the disk head, and reads or writes the data. Scan is one of the traditional disk scheduling algorithms. It serves requests in the same order as the data are encountered as the disk arm moves forward. The disk arm retraces back as soon as no pending requests are waiting in the forward direction. Data are transferred in both the forward and the retracing directions. In a variation called C-Scan, no data is transferred during the retracing step. In Scan-EDF, the requests are normally served in earliest deadline first order. When there are several requests with the same deadline, these requests are served with Scan scheduling, in which the specific one located nearest to the disk arm in the disk scan direction is served first. The optimization of this algorithm depends highly on the number of requests with the same deadline, and therefore its efficiency relies on how often seek optimization can be applied.

GSS is a variation of Scan. It divides the set of n requests being served into g groups. Groups are scheduled in a fixed order, and individual requests within a group are served with Scan scheduling. If g equals 1, GSS degenerates to Scan scheduling, and if g equals n, GSS becomes round-robin scheduling. Optimization of this algorithm depends on deriving groups to balance the tradeoff between the round duration and the latency of successive retrieval (28).

The video retrieval scheme in Ref. 29 uses the scan algorithm for video retrieval of all segments in a service round. A novel shared buffer architecture at the server is proposed for storing segments retrieved from disk. Video retrieved from disk is stored at designated locations in the buffer, while transfer of video on to the network takes place in a serial fashion from the start to the end of the buffer. The retrieval schemes in Refs. 26 and 27 divide a service round into a certain number of groups, and the scan algorithm is used for video retrieval within each group. The quasi-static video retrieval scheme uses the EDF algorithm to assign segments to be retrieved to groups. The QuIVeR protocols describe several simple and efficient ways of assigning segment retrievals to groups. By suitably constraining allowable user playback rates, the protocols achieve a small number of groups within a service round with a large number of segment retrievals within each group. Thus, in addition to achieving high disk retrieval efficiency, these protocols have an extremely small buffer requirement. They dispense with the server buffer and require a set-top box buffer of merely two video segments per user irrespective of the user playback rate.

Disk Array

Each storage device has a limit on the storage capacity and throughput bandwidth. For example, a disk drive with 45 Gbyte of storage and 68 Mbps of data transfer rate is only able to store about 40 MPEG-1 encoded movies and serves less than 20 video sessions with VHS quality. A disk array is a collection of independent disks using currently available technology to form a logical high-capacity and high-throughput disk. The redundant arrays of inexpensive disk (RAID) technique (30) is one such attempt to increase I/O (input/output) bandwidth. Important design issues associated with disk-array-based continuous servers include data placement across multiple disks, load balancing, and fault tolerance. For data placement across multiple disks, three different techniques can be used:

1. *Disk Striping.* Assume there are n disks in a disk array. The data blocks in a file are striped in a round-robin manner over n disks. To play back the file, each disk in the array operates synchronously (i.e., each disk head is located at the same position, and transfer data together), and all together, the n disks function as a single disk with n times the transfer rate and storage capacity. For example, if an audio clip X is striped onto 10 disks, the frames $X_0, X_{10}, X_{20}, \ldots$ are stored on disk 1, $X_1, X_{11}, X_{21}, \ldots$ are stored on disk 2, and so forth. To play back audio X, in the first time period, disk 1 reads X_0, disk 2 reads $X_1, \ldots$ and disk 10 reads X_9. All these reading activities happen at the same time. Therefore, the large logical data block transferred is composed of n smaller physical data blocks. This approach works well for applications with large bandwidth requirements, since it exploits the parallelism of multiple disks.
2. *Disk Declustering.* Assume the number of disks is the same and distribute the data blocks in a round-robin manner, as in data striping. To play back the file, each disk in the array operates asynchronously and retrieves data blocks in the same round-robin manner. For example, in the preceding example, disk 1 reads frame X_0 during the first time period, disk 2 reads X_1 during the second time period, and so forth until all blocks of X are read. This approach exploits the concurrency of multiple disks.
3. *Hybrid Combination.* This method combines the two schemes described previously (i.e., it scatters data blocks into multiple, independent and synchronized modules by data declustering, with each module consisting of several synchronized disks according to the data striping scheme). These independent modules can even be distributed in remote servers over the network.

Dynamic load balancing attempts to reduce the load disparities among multiple disks due to different popularities of objects (e.g., a new movie may be accessed much more often than a documentary), user access patterns (e.g., for video-on-demand services, subscribers tend to watch programs at night rather than during the day), and device utilization. The goal is to avoid load imbalance and to have optimal data placement across the storage hierarchy (e.g., a “hot” video is stored in a disk drive, while an unpopular video is archived in tertiary storage) and across multiple disks (e.g., popular movies may be replicated in multiple disks, while unpopular movies are stored in singleton) to achieve the maximum utilization of both space and bandwidth in the storage system. Important techniques include time staggering (31), rate staggering (32), and rotational mirrored-and-chained declustering (33). A fault-tolerant design prevents playback interruption due to disk failures and improves the reliability and availability of the storage system. In a continuous media server, disk failure does not cause data loss because data are archived in tertiary storage. Instead, it interrupts the retrieval progress and results in discontinuity in playback. Solutions include mirroring and parity approaches, in which a part of disk space is allocated to store redundant information.

SERVICE DELIVERY

To be commercially viable, VoD services must be priced competitively with existing video rental services (the local video rental store). As a result, system design issues center around the improvement of system capacity in terms of the number of customers the system can support. There are many ways to increase the system capacity. For instance, one may design more efficient and intelligent video servers (such as by grouping video frames or by scheduling user requests to decrease the movements of disk heads), enhance the modulation scheme in the hybrid fiber coax (HFC) access network to pack more users into a 6 MHz channel, and smooth the variable bit rate traffic of the digitized stored compressed video stream in the transport network. Another promising approach takes advantages of resource sharing and multicasting by batching multiple users in a video stream (34–37). Note that a video stream here refers to the total resource requirement along the data path from the server to the set-top box, including the I/O bandwidth in the server, buffers, and transmission bandwidth in the network. This method has the potential to improve system capacity dramatically at the expense of complicating the provision of user interactions.

Since the deployment of VoD service requires a major investment in the system and network infrastructures, a major uncertainty faced by potential service providers is whether their customers are willing to pay enough for VoD services so they can get a reasonable return on their initial investments. Without reducing the service delivery cost, some predict that VoD is a dead market. Therefore, improving system capacity is the key to the successful commercial deployment of VoD services. To allow true VoD services, one solution is to have a dedicated video stream for each customer. This is expensive. In addition, it is wasteful of system resources because if multiple users are viewing the same video, the system has to deliver multiple identical copies at the same time. To reduce this cost, batching and multicasting may be used. This allows multiple users accessing the same video to be served by the same video stream. Batching increases the system capacity in terms of the number of customers the system can support. However, it does complicate the provision of user interactions. In staggered VoD (34), for example, multiple identical copies (streams) of the same video program are broadcast, staggered in times (e.g., every 5 min). A user will be served by one of the streams. User interactions will be simulated by jumping among different streams. However, not all user interactions can be simulated in this fashion, and even for those that can be simulated, the effect is not exactly what the user demands. For example, fast playback cannot be issued because there is no stream in the fast playback mode. Also, pause duration can be allowed only for the integer multiples of staggered intervals, such as 5, 10, 15 min, but not for 7 min, because there is no stream which is offset in time by 7 min from the original stream. Even though staggered VoD only permits near VoD services, it is easy to implement and enjoys nice deployment advantages. Reference 35 added more complexities in the set-top buffer at the customer premises to allow limited user interactivity.

In Ref. 36, the authors proposed look-ahead scheduling with set-aside buffer scheme to simulate VoD services. Only pause and resume operations are considered. The advantage of this approach is that it exploits batching; the disadvantage is that it only allows near VoD services because only pause and resume operations are supported.

In Ref. 10, the authors realized that user interaction for an individual viewer must be made on a dedicated channel. To exploit batching, users are originally served in a batch. When users in the batch initiate user interactions, a new stream will be created for each interactive user, who will hold onto this stream until disconnection. This will work only if very few users are expected to issue interactive operations. Otherwise, the system may start in a batch mode but will degrade to a nonsharing mode as more and more users split off into their own streams.

The split-and-merge (SAM) protocol (37) provides true VoD services while fully utilizing the strength of the batching concept. Split-and-merge refers to the split and merge operations incurred when each user initiates user interactions. It starts with serving customers in a batch. When a user in the batch initiates a user interaction, the interactive user is split out of the original batch and temporarily assigned to a new video stream. Because the interactive user has a dedicated video stream, any type of interaction may be performed. As soon as the user interaction ends, the system will merge this interactive user back to the nearest ongoing video stream. Thus the system exploits batching and simultaneously allows the full spectrum of user interactions.

STANDARDIZATION EFFORTS

The Digital Audio-Visual Council (DAVIC), based in Geneva, Switzerland, was established in June 1994. It is a consortium of the manufacturing industry (computer, consumer electronics, telecommunication equipment), service industry (broadcasting, telecommunications, CATV), government agencies, and research organizations. It aims to standardize of the tools to provide digital audiovisual applications and services for end-to-end systems. DAVIC issued specifications for open interfaces, protocols and architectures that maximize interoperability across countries and applications/services (38). Such specifications are released in versions: DAVIC 1.0, DAVIC 1.1, DAVIC 1.2, DAVIC 1.3, and so on. Each future version will extend different grades of functionality defined previously or add more functionality to existing ones.

In DAVIC 1.0, a set of tools to specify a common interface to diverse access networks and a convergent view of end-to-end signaling were defined. Basic applications supported by DAVIC 1.0 include TV distribution, near video on demand, video on demand, and simple forms of home shopping. DAVIC 1.1 added tools to support basic Internet compatibility, multichannel multipoint distribution system (MMDS) and local multipoint distribution system (LMDS) access, network-independent set-top units (STU), and STUs that behave as virtual machines. In DAVIC 1.2, JAVA API (application program interface), Dolby AC3 high-quality sound, and 1080-line resolution high-quality video were included. The latest version is DAVIC 1.3. It added comprehensive service and network management, multiple broadcast servers, mobile reception, scaleable audio, content and meta-data packaging, and a novel definition of contour. The rule of thumb for DAVIC specifications is to pick existing standards for the most part and to define a new standard if the particular technology does not exist. For example, DAVIC selected MPEG-2 for video, MPEG-1 for audio, digital storage media-command and control (DSM-CC) for session control and user control, Q.2931 for connection control, ATM for the core network, and HFC, FTTB, FTTC, FTTH, ADSL, VDSL, MMDS, LMDS, and so on for the access networks.

DAVIC aims to facilitate the introduction of a wide range of digital audiovisual applications of broadcast and interactive types. Currently 19 core DAVIC applications have been identified and prioritized. Such applications include movies on demand, teleshopping, broadcast, near VoD, delayed broadcast, game, telework, karaoke on demand, Internet access, news on demand, TV listings, distance learning, videophone, home banking, telemedicine, content production, transaction services, and videoconferencing. These applications may have their own specific functionality and may have some in common with others. DAVIC identifies some core functions that are function elements used by more than one application, providing the basis for system operation, integrity, and development. Figure 14 depicts those core functional groups currently defined in DAVIC 1.3, including bit transport, session, access control, navigation and program selection, application launch, media synchronization links, application control, presentation control, usage data, and user profile. The interrelationship of the core functional groups and running applications is also illustrated in the figure.

DAVIC Reference Model

A typical system addressed by DAVIC consists of five components—namely, the content provider system (CPS), the service provider system (SPS), the service customer system

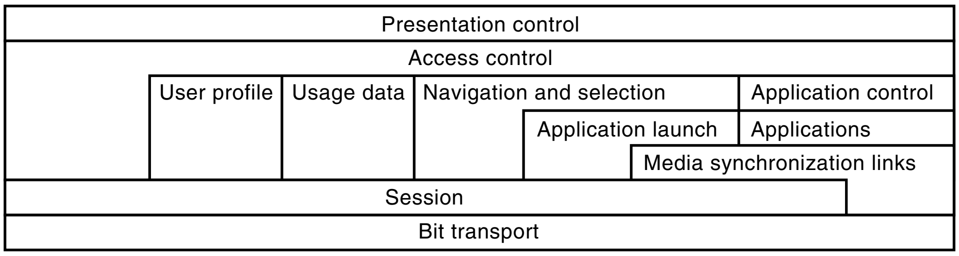

Figure 14. Interrelationships of the core functional groups and applications. (Reprinted with permission from (38). © DAVIC 1997.)

(SCS), and the CPS-SPS and SPS-SCS delivery systems. Figure 15 depicts the reference model of a DAVIC system.

The model defines reference points in the system. A reference point represents a set of interfaces between any two subsystems through which information flows from one to the other. These points are of particular interest in the system, and those that are accessible have normative values defined. A system conforms to DAVIC specifications if its external reference points do. Thus, to ensure interoperability for a given set of services, interacting systems must obey the functional requirements and the interface requirements specified at given reference points. The information flows through the subsystems can either be content information or control information. Five information flows have been defined: S_1 to S_5.

1. S_1 is content information flow that corresponds to the major service information, such as audio, video, or data transferred from a service provider to a customer.
2. S_2 is control information flow that corresponds to user control information, such as the command to play or stop a movie, to change S_1 information flow.
3. S_3 is control information flow that corresponds to session control information. Examples include establishing, modifying, or terminating a session, reporting exceptions, and negotiating resource requirements.
4. S_4 is control information flow that corresponds to connection management information. Examples include establishing or releasing connections, communicating addresses and port information, and exchanging routing information.
5. S_5 is management information flow that corresponds to billing and other management functions.

Figure 16 shows an architectural view of a DAVIC system. Multiple SPSs may be used to provide DAVIC services. SCS may request a DAVIC service from a given SPS. An SPS may either own media servers to provide materials directly or act as an agent to obtain information from a CPS through the SPS-CPS delivery system, and then transports information to the SCS via the SPS-SCS delivery system. Figure 17 depicts the system model of SPS. Each content provider subdomain element in a given content provider provides contents and applications for the users. The service gateway controls access to each content provider. The rest of the elements provide such functions as connection management, session control, and client profile management. Figure 18 depicts the system

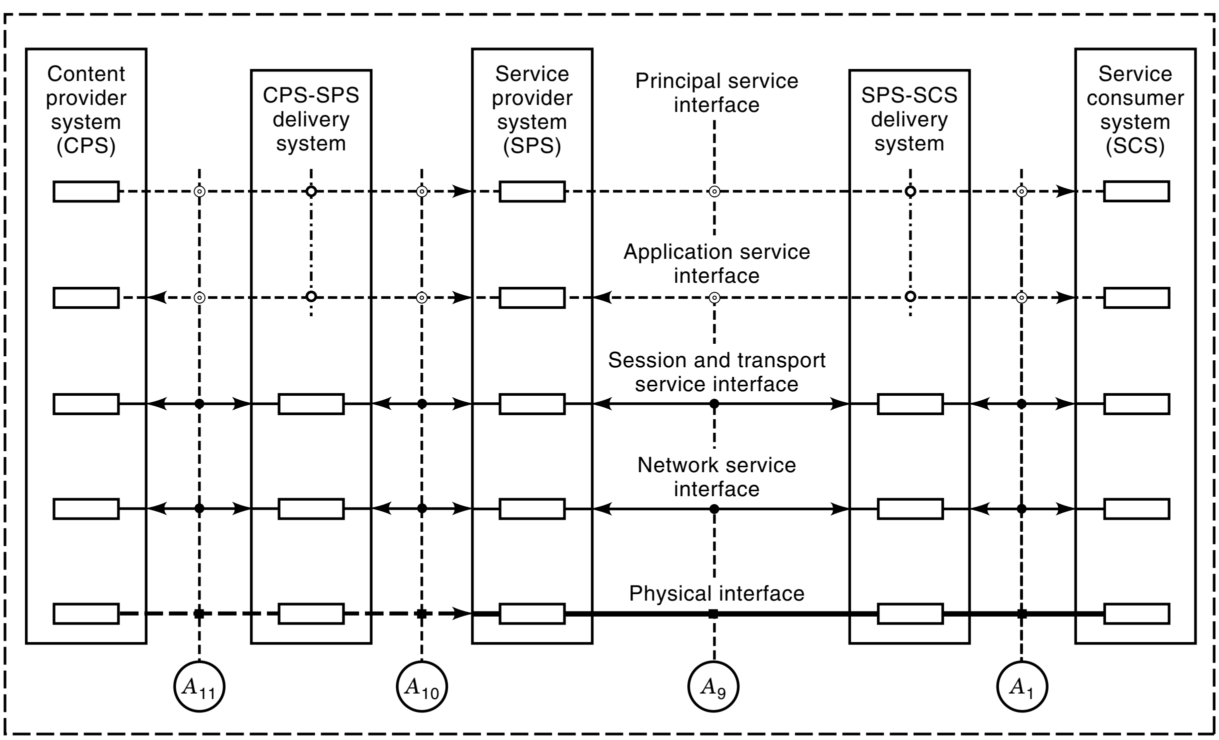

Figure 15. The reference model of DAVIC. (Reprinted with permission from (39). © DAVIC 1997.)

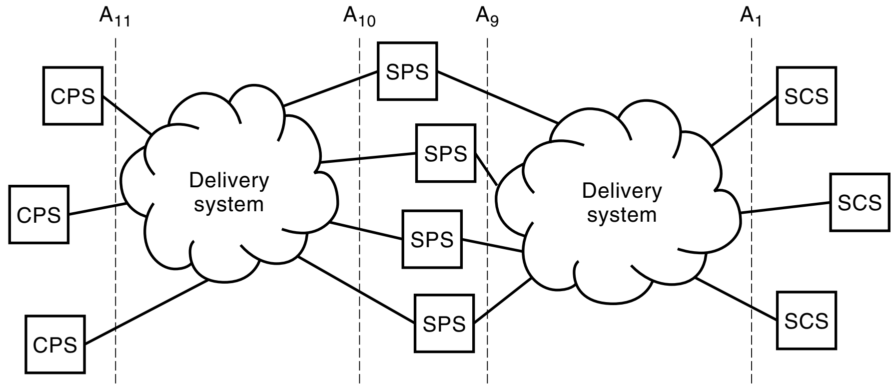

Figure 16. Typical architecture of a DAVIC system. (Reprinted with permission from (12). © DAVIC 1997.)

model of the delivery system. In DAVIC 1.3, only the reference points A_1 and A_9 are considered in the delivery system. SL_0, SL_1, SL_2 provide all controls for the services offered by the delivery system. SL_3 provides such network control functions as network configuration, connection establishment and termination, and routing. The core network plays the role of the backbone network described in an earlier section, and the access network serves the same purpose as that described earlier. The management entity provides operation and maintenance functions for the network, using S_5 information flow to communicate with other management entities. Figure 19 depicts the system model of SCS. It consists of two primary parts: STU and a network interface unit (NIU). The STU contains the core functions of the SCS that allows user interactions with the system. The product entity accepts and presents content information to the customer; the application entity handles interactive messages. The environment entity provides environmental requirements such as negotiated service quality on which the DAVIC application operates. The connectivity entity is responsible for the reliable transmission of content flow information to and from the STU. NIU adapts network-dependent content information flow from the delivery system to network-independent interface to the STU.

FUTURE DEVELOPMENTS

Originally, ITV was intended to provide high-fidelity, entertainment-oriented interactive TV services to residential customers. Such large-scale VoD systems, supporting tens and hundreds of thousands of users geographically distributed over a wide area and providing services such as movies on demand, were considered the basic building blocks of distributed multimedia services. Cable TV operators, telephone carriers, satellite companies, wireless access providers, and other service providers are considered to be the potential competitors for this market. Numerous experiments and field trials have been announced around the world (40). However, primarily because of the significant investment in communication and video server resources required for such large-scale

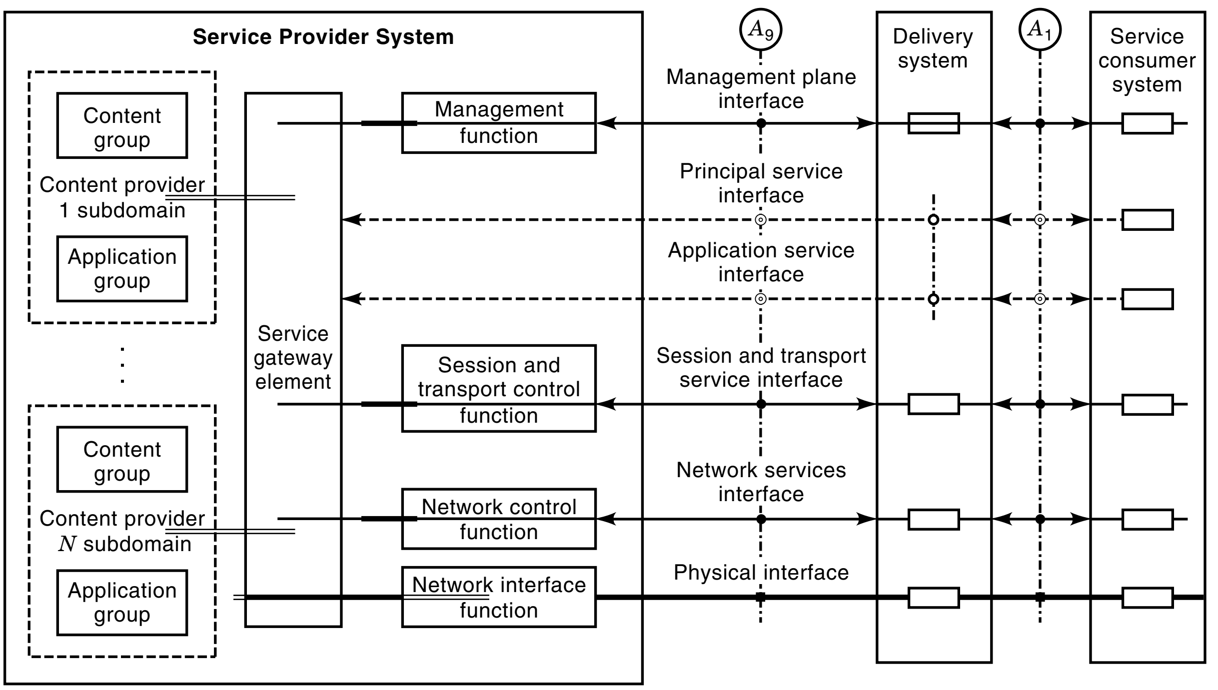

Figure 17. System model of a service provider system. (Reprinted with permission from (39). © DAVIC 1997.)

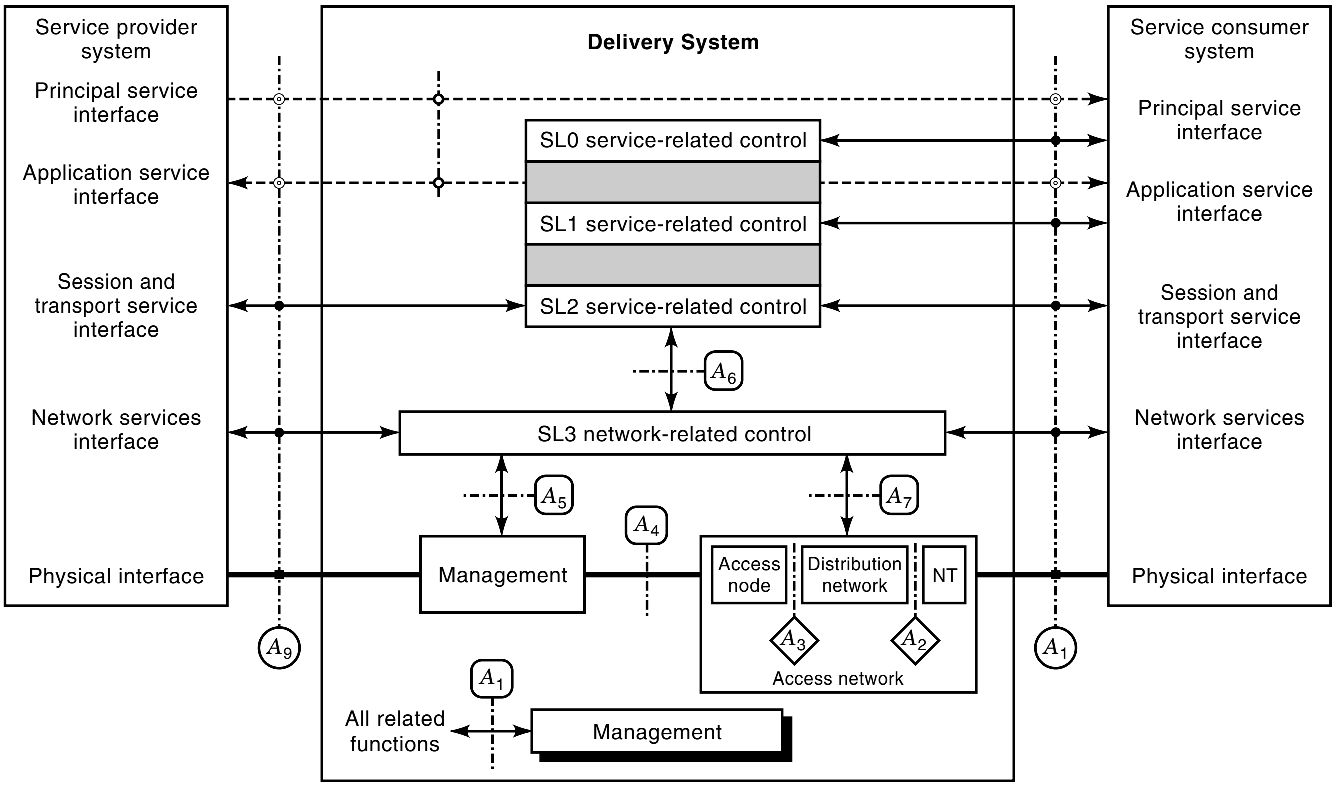

Figure 18. System model of a delivery system. (Reprinted with permission from (39). © DAVIC 1997.)

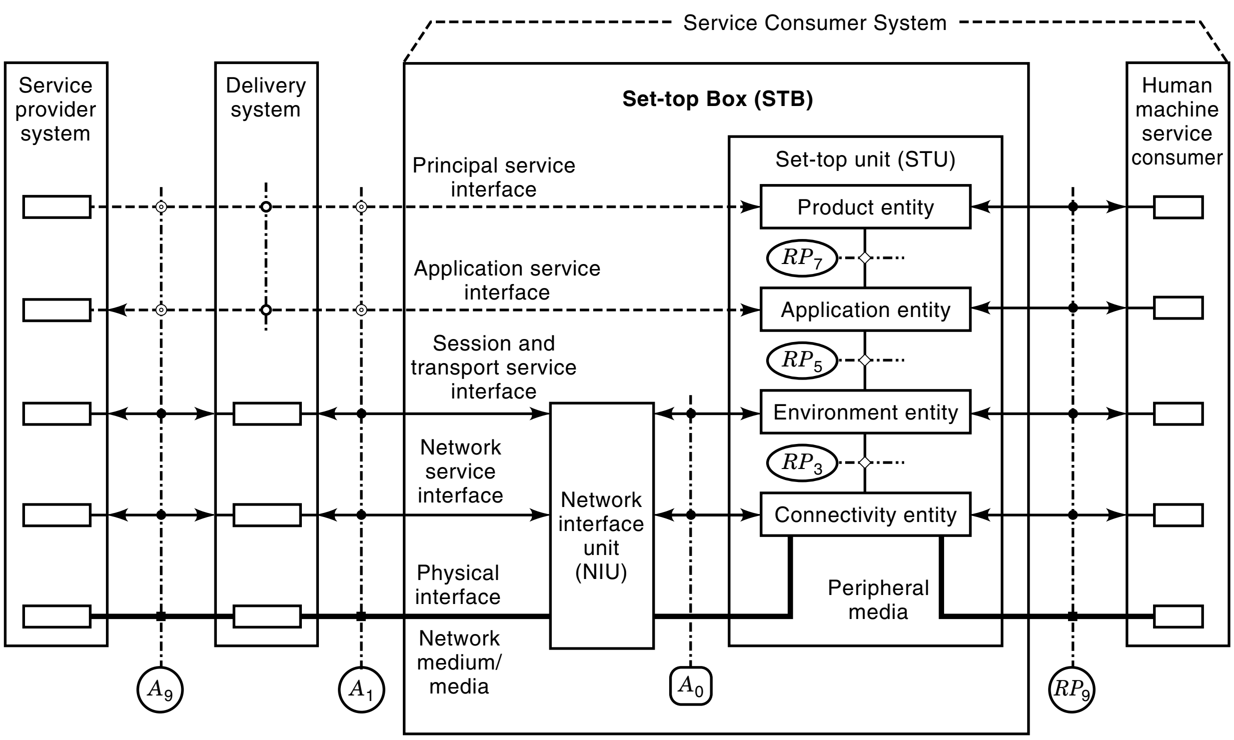

Figure 19. System model of a service consumer system. (Reprinted with permission from (39). © DAVIC 1997.)

systems and the uncertainty of customer demand, many organizations have scaled down or postponed their deployments. However, a few organizations have pressed on. Hongkong Telecom is one of them.

First Commercial Residential ITV Deployment

On 23 March 1998, Hongkong Telecom (HKT) officially launched the world's first interactive TV (iTV), making Hong Kong the first major city where truly interactive television services are commercially available. HKT iTV offers video on demand, music on demand, home shopping, broadband Internet, and home banking. Each customer is provided with a remote control and a digital smart box that is connected to the TV and the phone line. The monthly subscription charge is about 30 U.S. dollars (USD30), and the cost of a movie is up to USD3.55, with a short cartoon costing USD1.45. The VOD trial program started in 1994, with commercial service initially scheduled for July 1996. It was delayed to allow HKT to make use of new and emerging technologies, resulting in households being connected to the Internet through a television-top smart box. Seven thousand households in Hong Kong participated in the trial program, which was free of charge until the commercial launch. Hongkong Telecom expects to sign up 200,000 of the 1.6 million households in Hong Kong within the first year of the launch. The system uses NEC video servers and Fujitsu ATM switches capable of a capacity of 80 Gbps with ADSL access technology, equivalent to about 18,000 MPEG-2 motion picture streams.

Other Developments

Smaller-scale VoD systems that do not require huge investments in infrastructure and that serve tens and hundreds of users in a small area have been developed and deployed. Examples include movies-on-demand systems in hotels, hospitals, airplanes, and cruise ships. In addition, interactive VoD has found other applications besides entertainment, such as in education and training, and many organizations are developing enterprise VoD systems. With the explosive growth of the Internet, the availability of Web-enabled software tools and programs and the development of related enabling protocols such as real time transport protocol (RTP) (41) and real time streaming protocol (RTSP) (42), another delivery system is emerging to furnish VoD services. The Next Generation Internet (NGI) is expected to provide advanced multimedia services (including video conferencing, ITV, and virtual reality) over an integrated network. This integrated network, expected to provide much higher bandwidth compared to today's Internet, will consist of a network of networks, including the existing cable, telephone, satellite networks. Therefore, ITV is expected to fulfill its promise to provide the ultimate flexibility in television and video services.

BIBLIOGRAPHY

1. W. W. Hodge, *Interactive Television: A Comprehensive Guide for Multimedia Technologies,* New York: McGraw-Hill Series on Visual Technology, 1995.
2. Y. H. Chang et al., An open-systems approach to video on demand, *IEEE Commun. Mag.,* **32** (5): 68–80, 1994.
3. D. Deloddere, W. Verbiest, and H. Verhille, Interactive video on demand, *IEEE Commun. Mag.,* **32** (5): 82–88, 1994.
4. Special Issue on Access to Broadband Services, *IEEE Commun. Mag.,* **33** (8): 1995.
5. Special Issue on Digital Interactive Broadband Video Dial Tone Networks, *IEEE Network Mag.,* **9** (5): 1995.
6. T. D. C. Little and D. Venkatesh, Prospects for interactive video-on-demand, *IEEE Multimedia,* **1** (3): 14–24, 1994.
7. W. D. Sincoskie, System architecture for a large scale video on demand service, *Computer Networks and ISDN Systems,* **22**: 155–162, 1991.
8. B. Furht et al., Design issues for interactive television systems, *IEEE Computer,* **28** (5): 25–39, 1995.
9. V. O. K. Li and W. J. Liao, Interactive video on demand systems, *1996 Multimedia Technol. Appl. Conf.,* Kaohsiung, Taiwan, 1996, pp. 9–25.
10. V. O. K. Li et al., Performance model of interactive video-on-demand systems, *IEEE J. Select. Areas Commun.,* **14** (6): 1099–1109, 1996.
11. B. Furht, Multimedia systems: An overview, *IEEE Multimedia,* **1** (1): 47–59, 1994.
12. The Digital Audio-Visual Council (DAVIC), "Part 4—Delivery System Architecture and Interfaces," DAVIC 1.3 Specifications, Geneva, Switzerland, [Online], Available 1997. http://www.davic.org
13. S. Ramanathan and R. Gueslla, Toward management systems for emerging hybrid fiber-coax access networks, *IEEE Network Mag.,* **9** (5): 58–68, 1995.
14. S. Dixit and P. Skelly, MPEG-2 over ATM for video dial tone networks: Issues and strategies, *IEEE Network Mag.,* **9** (5): 30–40, 1995.
15. D. Large, Creating a network for interactivity, *IEEE Spectrum,* **32** (4): 58–63, 1995.
16. W. Y. Chen and D. L. Waring, Applicability of ADSL to support video dial tone in the copper loop, *IEEE Commun. Mag.,* **32** (5): 102–109, 1994.
17. P. J. Kyees, R. C. McConnell, and K. Sistanizadeh, ADSL: A new twisted-pair access to the information highway, *IEEE Commun. Mag.,* **33** (4): 52–59, 1995.
18. Special Issue on Fiber-Optic Subscriber Loops, *IEEE Commun. Mag.,* **32** (2): 1994.
19. V. Loen and E. Miller, Subscriber terminal units for video dial tone systems, *IEEE Network Mag.,* **9** (5): 48–57, 1995.
20. R. Steinmetz, Analyzing the multimedia operating system, *IEEE Multimedia,* **2** (1): 68–84, 1995.
21. S. Sengodan and V. O. K. Li, A generalized grouping and retrieval scheme for stored MPEG video, *Proc. IEEE ICC,* Montreal, Canada, 1997, pp. 1674–1678.
22. P. V. Rangan, H. M. Vin, and S. Ramanathan, Designing an on-demand multimedia service, *IEEE Commun. Mag.,* **30** (7): 56–65, 1992.
23. H. M. Vin and P. V. Rangan, Designing a multiuser HDTV storage server, *IEEE J. Select. Areas Commun.,* **11** (1): 153–164, 1993.
24. A. L. N. Reddy and J. Wyllie, Disk scheduling in a multimedia I/O system, *Proc. ACM Multimedia '93,* Anaheim, CA, 1993, pp. 225–233.
25. M. S. Chen, D. D. Kandlur, and P. S. Yu, Optimization of the grouped sweeping scheduling (GSS) with heterogeneous multimedia streams, *Proc. ACM Multimedia '93,* Anaheim, CA, 1993, pp. 57–64.
26. S. Sengodan and V. O. K. Li, A quasi-static retrieval scheme for interactive disk-based VOD servers, *Computer Commun.,* **20** (12): 1031–1041, 1997.
27. S. Sengodan and V. O. K. Li, QuIVeR: A class of interactive video retrieval protocols, *Proc. IEEE Conf. Multimedia Comput. Syst. (ICMCS),* Ottawa, Canada, 1997, pp. 186–193.

28. D. J. Gemmell et al., Multimedia storage servers: A tutorial, *IEEE Comput.,* **28** (5): 40–49, 1995.
29. S. Sengodan and V. O. K. Li, A shared buffer architecture for interactive video servers, *Proc. IEEE INFOCOM'97,* Kobe, Japan, 1997, pp. 1343–1350.
30. D. A. Patterson, G. Gibson, and R. H. Katz, "A case for redundant array of inexpensive disks (RAID)," *Proc. ACM SIGMOD,* Chicago, IL, 1988, pp. 109–116.
31. S. Berson et al., Staggered striping in multimedia information systems, *Proc. ACM SIGMOD,* Minneapolis, MN, 1994, pp. 79–90.
32. M. S. Chen, D. D. Kandlur, and P. S. Yu, Using rates staggering to store scalable video data in a disk-array-based video server, *Proc. IS&T/SPIE Symp. Electronic Imaging—Conf. Multimedia Comput. Netw.,* SPIE **22417**, 1995, pp. 338–345.
33. M. S. Chen et al., Using rotational mirrored declustering for replica placement in a disk-array-based video server, *Proc. ACM Multimedia'95,* 1995, pp. 121–130.
34. R. O. Banker et al., *Method of providing video on demand with VCR like functions,* US Patent 5,357,276, 1994.
35. K. C. Almeroth and M. H. Ammar, The use of multicast delivery to provide a scalable and interactive video-on-demand service, *IEEE J. Select. Areas Commun.,* **14** (6): 1110–1122, 1996.
36. P. S. Yu, J. L. Wolf, and H. Shachnai, Design and analysis of a look-ahead scheduling scheme to support pause-resume for video-on-demand application, ACM *Multimedia Syst.,* **3** (4): 137–149, 1995.
37. W. J. Liao and V. O. K. Li, The split-and-merge (SAM) protocol for interactive video on demand, *IEEE Multimedia,* **4** (4): 51–62, 1997.
38. The Digital Audio-Visual Council (DAVIC), Part 1—Description of DAVIC functionalities, DAVIC 1.3 specification, Geneva, Switzerland, [Online], 1997. Available http://www.davic.org
39. The Digital Audio-Visual Council (DAVIC), Part 2—System Reference Models and Scenarios, DAVIC 1.3 specification, Geneva, Switzerland, [Online], 1997. Available http://www.davic.org
40. T. S. Perry, The trials and travails of interactive TV, *IEEE Spectrum,* **33** (4): 22–28, 1996.
41. H. Schulzrinne, *RTP profile for audio and video conferences with minimal control,* RFC 1890, IEFT, 1996.
42. H. Schulzrinne, A. Rao, and R. Lanphier, *Real time streaming protocol (RTSP),* Internet draft, IETF, 1998.

VICTOR O. K. LI
The University of Hong Kong
WANJIUN LIAO
National Taiwan University

INTERACTIVE VOICE RESPONSE SYSTEMS. See VOICE MAIL.

INTERCONNECT DELAY. See INTEGRATED CIRCUIT SIGNAL DELAY.

INTERCONNECTION HIERARCHY. See SYSTEM INTERCONNECTS.