Distributed Multimedia Systems

Chapter Overview

In this chapter, we outline the issues concerning Distributed Multimedia Systems (DMS). We give an overview of DMS: main features, resource management, networking and multimedia operating system. Next, we identify the applications (interactive television, telecooperation and hypermedia) and survey the important enabling technologies. These topics will continue to be of great interest in the near future.

4.1 Introduction

A DMS is an integrated communication, computing and information system that enables the processing, management, delivery and presentation of synchronized multimedia information that the quality of service guarantees [4.1, 4.2]. It integrates and manages the information communication and computing subsystems to realize multimedia applications. Such a system enhances human communications by exploiting both visual and aural senses and provides the ultimate flexibility in work and entertainment by allowing you to collaborate with remote participants, view movies on demand and access online digital libraries from the desktop. DMS will create an electronic world. Technological advances in computers, high-speed networks, data compression and consumer electronics—coupled with the availability of multimedia resource mechanism, and manipulation functions; the development of the relevant standards and the convergence of the computer, telecommunications, and digital TV industries—are accelerating the realization of such systems. An example of DMS is a number of multimedia PCs and/or workstations interconnected with continuous media servers usings the Internet that allow users to

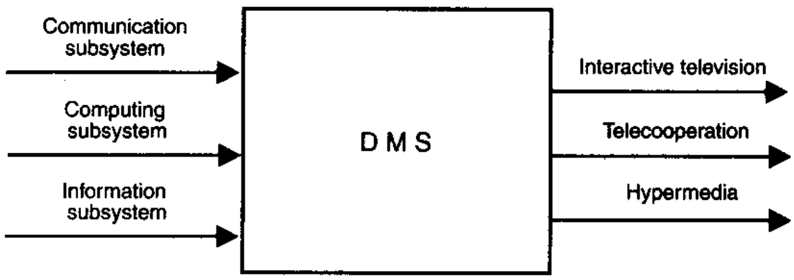

Figure 4.1 Block scheme of a summarized DMS.

retrieve, browse and manipulate video or audio. Besides constraints on bit error rates, packet-loss probabilities and delivery delays required in a point-to-point information delivery system, additional constraints are introduced in a DMS, such as the synchronization among multiple media streams from distributed sources to achieve a meaningful presentation. Figure 4.1 summarizes a DMS.

The inputs of the system consist of the factors that drive a DMS from concepts to reality, and the outputs consist of a wide range of distributed multimedia applications. The system inputs are a collection of the enabling technologies of the communication subsystem (for transmission), the computing subsystem (for processing) and information subsystem (for storage). The communication subsystem consists of the transmission medium and transport protocols. It connects the users with distributed multimedia resources and delivers multimedia materials with QoS guarantees, such as real-time delivery for video or audio data and error-free delivery for text data. The computing subsystem consists of a multimedia platform, operating system (OS), presentation and authoring tools and multimedia manipulation software. It allows users to manipulate the multimedia data. An authoring tool is specialized software that allows a producer or designer to design and assemble multimedia elements for a multimedia presentation. The information subsystem consists of the multimedia servers, information archives and multimedia database systems.

The outputs of the system can be broadly classified into three different types of distributed multimedia applications: Interactive Television (ITV), telecooperation and hypermedia. ITV allows subscribers to access video programs and interact with them. Services include home shopping, interactive video games (which can be classified as hypermedia applications), financial transactions, Video on Demand (VoD), news on demand and so forth. Telecooperation overcomes time and location restrictions and allows remote participants to join a group activity. Services include remote learning, telecommuting, teleservicing, teleoperation, multimedia emails, videophone, desktop conferencing, electronic meeting rooms, joint editing and group drawing. A hypermedia document is a multimedia document with links to other multimedia documents and it allows users to browse multimedia information in a consequential manner. Services include digital libraries, electronic encyclopedias, multimedia magazines, multimedia documents, information kiosks, computer-aided learning tools and the Web.

4.2 Main Features of a DMS

The main features of a DMS can be summarized as follows [4.3]:

- ***Technology integration***—Integrates information, communication and computing systems to form a unified digital processing environment.
- ***Multimedia integration***—Accommodates discrete data as well as continuous data in an integrated environment.
- ***Real-time performance***—Requires the storage systems processing systems and transmission systems to have real-time performance. Hence, huge storage volume, high network transmission rate and high CPU processing rate are required.
- ***Systemwide QoS support***—Supports diverse QoS requirements on an end-to-end basis along the data path from the sender, through the transport network and to the receiver.
- ***Interactivity***—Requires duplex communication between the user and the system and allows each user to control the information.
- ***Multimedia synchronization support***—Presents the playback continuity of media frames within a single continuous media stream, and temporal relationships among multiple related data objects.
- ***Standardization support***—Allows interoperability despite heterogeneity in the information content, presentation format, user interfaces, network protocols and consumer electronics.

According to the different requirements imposed upon the information, communication and computing subsystems, distributed multimedia applications may be broadly classified into ITV, telecooperation and hypermedia. ITV requires a very high transmission rate and QoS guarantees. It demands point-to-point, switched connections as well as good customer services and excellent management for information sharing, billing and security. Telecooperation, such as videophone and desktop conferencing, allows lower picture quality and therefore has a lower bandwidth requirement. It requires powerful multimedia database systems rather than just continuous media servers. Sharing information among groups is the key to effective collaboration. Hypermedia systems may be treated as an application of database systems because they provide flexible access to multimedia information and a novel method to structure and manage data. Hypermedia applications require point-to-point and switched services.

4.3 Resource Management of DMS

A DMS integrates and manages the information, communication and computing subsystems. The resource management ensures end-to-end QoS guarantees. Such guarantees have the following characteristics:

- Systemwide resource management and admission control to ensure desired QoS level.

- Quantitative specification (packet loss probability and delay jitter) rather than qualitative description as in the Internet TCP/IP. This gives the flexibility to accommodate a wide range of applications with diverse QoS requirements [4.4].
- Dynamic management, which means that QoS is dynamically adjusted rather than statistically maintained throughout the lifetime of the connection.

In the context of DMSs, QoS is defined as the quantitative description of whether the services provided by the system satisfy the application needs and is expressed as a set of parameter-value pairs [4.5]. Packet-loss probability (10^{-3}) and packet delay (10^{-6}s) are examples of such parameter-value pairs. These QoS parameters are negotiated between the users and the service providers. QoS requirements can be mapped into desired resources in system components. Then, they can be managed by integrated resource managers to maintain the QoS commitment according to the negotiated service contracts. The system supports three levels of QoS commitment: deterministic (guarantees that the performance is the same as the negotiated service contract), statistical (guarantees the performance with some probability) and best effort (does not offer service guarantees). The goal of systemwide resource management is the coordination among system components to achieve end-to-end QoS guarantees. The major functions include the following:

- Negotiate, control and manage the service contracts of the users.
- Reserve, allocate, manage, adopt and release system resources according to the negotiated values.

After the service contract has been negotiated, it will be preserved throughout the lifetime of the connection. It is also possible, through proper notification and renegotiations, to tune the QoS level dynamically [4.6]. Admission control protects and maintains the performance of existing users in the system. The principle of admission control is that new requests can be accepted so long as the performance guarantees of existing connections are not violated [4.7].

4.4 Networking

The network transports multimedia traffic to satisfy QoS guarantees on an end-to-end basis. The transmission media may be wired (coaxial cables or fiber optics) or wireless (radio channels or satellite channels). The transport protocols may provide connection-oriented or connectionless services and best effort, statistical or deterministic performance guarantees. The network can be a LAN, a Metropolitan Area Network (MAN) or a WAN. A LAN (Ethernet, token ring or token bus) may cover an area within a building, a campus or an organization. A MAN (for example, a Fiber Distributed Data Interface ([FDDI]), may cover a metropolitan area, such as a small city. A WAN (TCP/IP, ATM) is a nationwide or an international network.

Multimedia traffic has diverse characteristics and various QoS requirements. Discrete media traffic (for example, file transfer or image retrieval) requires error-free services, but is tolerant of delay. Continuous media traffic (for example, video or audio playback) requires real-time, high-speed transmission and is connection oriented. It is sensitive to delay and delay jitter or the packet delay variations between consecutive packets. On the other hand, it is tolerant of occasional packet losses. In addition, the network has to support application-specific requirements. For example, video conferencing needs multicast service for group distribution, but ITV requires switched point-to-point services and asymmetric bandwidth for the downstream (video server to user) and the upstream (user to video server) directions.

The Internet has rapidly evolved into a significant network infrastructure for communications. It runs the IP with its best effort delivery service and enjoys a large user base. Another promising technique, ATM, is rapidly appearing in the market. It allows bandwidth on demand and guaranteed QoS and is expected to be the best candidate for high-quality media delivery. In what follows, we will examine the Internet effort on the support of multimedia transmission, the properties of ATM that are especially suited for distributed multimedia applications and the issues of integrating ATM networks with the emerging integrated services Internet.

4.4.1 IP Networking

IP networking is a booming market for telecom and datacom service providers and equipment vendors. Service providers are quickly defining and bringing to market differentiated IP services, including voice transport, VPNs, application/policy prioritization, multimedia and transport LANs.

Internet refers to the global network to which a large percentage of existing networks are now interconnected by routers or gateways, running the TCP/IP protocol suite. As the key to success for the Internet, IP provides datagram delivery (best effort and connectionless) service and leaves the reliability issues (delay, out-of-order delivery, packet loss and misdelivery) to the end systems. It uses a global addressing scheme for a vast range of services. In addition to data, each datagram carries routing information in the header to independently forward to the destination.

Integrated management architecture leverages common data and human resources across applications such as email, and voice transport across an IP-based infrastructure demands low latency and jitter. The Round-Trip Time (RTT), which is the time required by a network to travel from the source to the destination and back, including the time to process the message and generate a reply, should be less than 250 ms. The servicing of these packets is done based on a priority scheduling scheme.

The primitive service model of the Internet provides only point-to-point and best-effort services. Such services are suited for traditional applications, such as file transfer and remote login. It performs well for real-time media traffic (for example video and audio) only under lightly loaded networks. UDP and Real-Time Transport Protocol (RTP) are typically used to transfer a Voice over IP (VoIP) packet. UDP is a connectionless transport layer protocol in the TCP/IP protocol stack. It is a simple protocol that exchanges datagrams without acknowledgment or guarantee of delivery, requiring that error processing and retransmission function be handled by other protocols. RTP is a protocol designed to provide end-to-end network transport functions for applications transmitting real-time data, such as voice, video or simulation data across multicast or multicast network services. RTP provides services such as stamping and delivery and monitoring of real-time applications.

The support of IP multicast, resource reservation and higher level RTPs are the major development efforts on the Internet toward multimedia applications. We will examine these three issues.

IP Multicast

Multicast refers to the ability to send a single packet to multiple destinations. Applications that require multicast include teleconferencing, email multicast, remote learning, and group communication. Traditional protocols like TCP or UDP provide only unicast transmission. The provision of multicast with unicast services requires the delivery of replicated unicast packets to each recipient. To avoid sending multiple unicast copies, thereby increasing network use, multicast transport service is required. With IP multicast, each data source sends a packet to a multicast address and lets the network forward a copy of the packet to each group of hosts. Figure 4.2 shows unicast and multicast services for multipoint communication. As it can be noted, there is a difference between the two transport services.

Multicast routing protocols provide a mechanism that enables routers to propagate multicast datagrams in order to minimize the number of excess copies transmitted to any particular subnetwork. In addition, they must be flexible enough to allow participants to join and leave the system without affecting others. The most important thing is to determine which hosts on which subnetwork will participate in the group. Two important multicast routing protocols are link state and distance vector multicast. They are extensions from the corresponding unicast routing algorithms.

In a link state protocol, such as Open Shortest Path First (OSPF), each router monitors the state (or metric) of its adjacent links and sends this information to all other routers whenever the state changes. The metric may be the reliability, delay, cast, and so forth of the link. Each router has a picture of the entire network. If router A wants to multicast to a group of users, it needs to identify the subset of routers. This is accomplished by adding the identifications of the groups that have members on a particular link to the state of the link. Then it computes the shortest distance tree routed at A and spanning this subset of routers based on the link-state metrics. The multicast is then performed on this shortest distance tree.

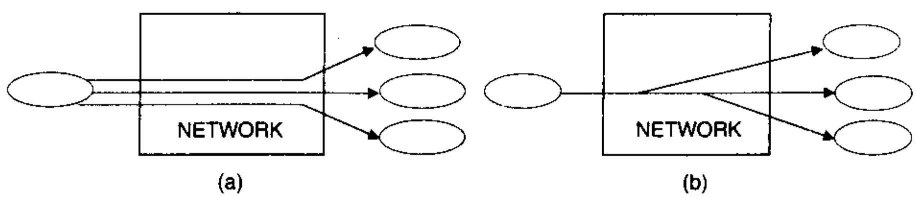

Figure 4.2 Unicast (a) and multicast (b) services for multipoint communication.

In a distance vector multicast protocol, each router maintains a routing table with entries (destination, cost, and next node) indicating for each destination the least cost to reach it and the next node on the shortest path to it. For each destination that it can reach, a node sends the destination cost to its directly connected neighbors, informing them of the cost of reaching the destination if this node is used as a relay. Each node will update its routing table if it finds that a neighbor can provide a shorter path to the destination. To support multicasting, each node will also send the identification of the groups that can be reached using this node to each of its directly connected neighbors. This information can be included in the routing table, and each router then knows across which links it should forward multicast packets for each group. To avoid sending duplicate multicast packets when multiple routers are connected to a router or link, we can designate the router as the parent router or link. The identification of this parent is a function of the source of the multicast. The router with the shortest path to the service is selected as the parent. Only the parent will forward a multicast packet.

Resource Reservation Protocol (RSVP)

This protocol attempts to provide guaranteed QoS for heterogeneous receivers across the Internet with multipoint-to-multipoint communications. Figure 4.3 summarizes the RSVP architecture. To reserve resources at a router, the RSVP block communicates with the admission control and policy control modules. Admission control determines if these are sufficient resources to satisfy the QoS of the new request while guaranteeing the QoS requirements for existing connections. Policy control determines if the user has the administrative permission for resource reservation. If both check out, the RSVP block sets parameters in the packet classifier, which determine the QoS for each packet, and the packet scheduler, which orders packet transmissions.

RSVP is receiver initiated. Two types of messages are used to reserve resources, that is, PATH and RESV. Each data source sends a PATH message, which contains a flow specification (for example, bandwidth) to the destination multicast address. When a router receives a PATH message, it records the relevant information (for example, IP multicast address, flow specification, source identification, and so forth). As a result, not only are the receivers informed of the

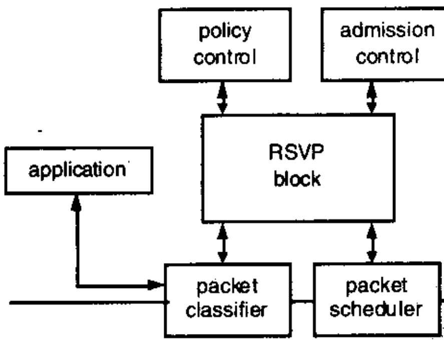

Figure 4.3 RSVP architecture.

flow specification of the data traffic, but the intermediate nodes also obtain the PATH state information. Based on information obtained from PATH messages and from higher layer protocols, each receiver can determine its QoS requirements and can initiate a RESV message to reserve specific resources along the reserved route of the PATH message. Multiple receivers may send RESV messages toward the same data source. RSVP scales well and accommodates heterogeneous receiver requirements. Resources will be received along the route from the source toward the receiver. For reliability and robustness, RSVP periodically updates the soft states at intermediate routers to allow for users joining and learning a group [4.8].

Three reservation styles have been defined: fixed filter, wildcard filter and shared explicit. Fixed filter allows explicit sender selection and distinct reservation and is suitable for applications, such as video conferencing, in which each videostream has its own RSVP flow. Wildcard sender selection includes any sender. Explicit sender selection refers to the reservation made only for senders explicitly listed in the RSVP message. Distinct reservation means that each sender has its own RSVP flow, and shared reservation uses the same RSVP flow for multiple senders. Wildcard filter allows wildcard sender selection and shared reservation. It is suitable for applications, such as audio conferencing, in which participants typically take turns speaking. Shared explicit allows explicit sender selection and shared reservation.

RTP

One of the IPs that can be used in conjunction with reservation models at the network layer is RTP [4.9] . RTP is an end-to-end protocol for the transport of real-time data. An important application type supported by RTP is multiparty conferencing because of its support for synchronization, framing, encryption, timing and service identification. RTP has its companion RTP Control Protocol (RTCP), which is used to interchange QoS and failure information between the QoS monitor applications in the end systems.

RTP does not define any kind of QoS itself and does not provide reordering or retransmission of lost packets. However, it provides a sequence number that enables the application using RTP to initiate such steps. RTP is directly used on top of UDP/IP. The RTP stack provides the information necessary to make educated guesses about the behavior of the datastream based on the RTP's knowledge of the data format. In addition to the base RTP specification, a number of companion documents exist that provide encapsulations for various continuous media formats, such as Motion Joint Photographic Experts Group (M-JPEG) or MPEG. Hence, RTP itself provides no real QoS support. It relies on other appropriate protocols and mechanisms.

4.4.2 Integrated Management Architecture for IP-Based Networks

Ensuring profitability from multimedia services requires a comprehensive service management architecture that enables service providers to plan carefully, provide quickly, operate efficiently and bill accurately for these services [4.10]. Figure 4.4 illustrates the typical structure of an IP network including core, edge and access subnetworks. The backbone technology is based on Multiprotocol Label Switching (MPLS), such as core and edge-label switch routers. MPLS networks integrate IP routing protocols, which allow efficient support of services such as IP

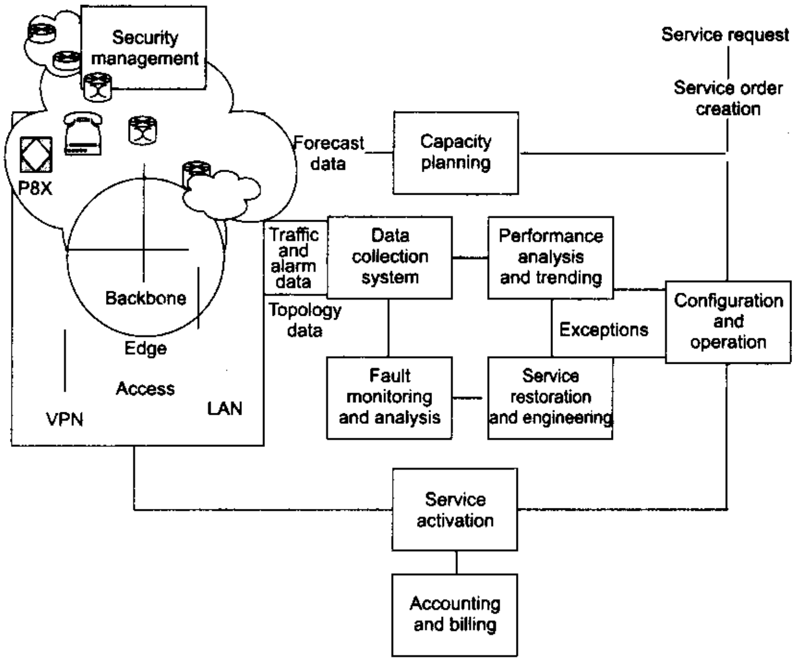

Figure 4.4 Typical structure of an IP network [4.10]. ©2000 IEEE.

multicast, IP class of service and IP VPNs. Customer sites are connected to the edge network, which is typically owned and controlled by a service provider. The extended functions include service restoration, traffic engineering, data collection, service activation and network planning.

In what follows, we will discuss an extended framework of the FCAPS functions for IP-based networks. The FCAPS acronym is used to refer to integrated management of various types of networks. It refers to the Open System Interconnection (OSI) for five functional areas: performance management, fault management, configuration management, security management and accounting and billing management [4.11].

Capacity planning provides a long-term view of network demands and requirements. It computes network element growth rates and generates a long-term capacity expansion plan. The network administrator who wants to do hypothetical studies typically carries out the capacity planning function to determine the required capacity as well as optimal equipment locations based on forecast or expected demands.

The data collection function collects and forwards data on a regular basis to the appropriate module. For example, alarms and related fault-statistics data are forwarded to the fault module to provide comprehensive diagnostic capabilities. Traffic-statistics data is forwarded to the performance module for data analysis. Traffic-statistics as well as network topology data may be forwarded to the planning module [4.12]. Estimated factors can then be used to estimate the future forecasting traffic loads. Finally, traffic statistics can also be used as the basis to observe traffic loads and to estimate the load for use in network engineering.

Performance Management

This is the process of converting IP traffic measurements into meaningful performance measures. It can be divided into real-time and long-term management. The real-time performance management process is a mechanism to guarantee that enough bandwidth is reserved for time-sensitive IP voice traffic, while other applications sharing the same link get their share without interfacing with the mission traffic. Another example is constant monitoring of high-priority customer services as well as customers who have been complaining about the performance of their services. Long-term performance management supports studies that monitor the ability of the existing IP networks to meet service objectives. The purpose is to identify situations where corrective planning is necessary. This is needed when objectives are not being satisfied and, where possible, to provide early warning of potential service degradation so that a corrective plan can be formulated before service is affected.

Traffic measurements are collected, validated by data-collection systems and then stored in batch mode in a database. Examples of IP performance traffic measurements include the following:

- Number of packets received per interface
- Number of packets transmitted per interface
- Number of packets dropped due to mild and severe congestion per interface (wild and severe congestion states are defined by the network administrator for each service)
- Number of packets dropped due to protocol errors
- Amount of time a network element is in a mild or a severe congestion state
- Number of times a network element enters a mild or severe congestion state

The performance management process then converts the validated measurements into meaningful network element loads (use, packet loss ratio, delay, jitter, and so forth). Next, it calculates statistics to characterize the load for traffic engineering purposes (for example, average peak values or average busy season). The process then computes network element performance measures (route-delay, end-to-end packet loss, average and peak packet loss, and so forth) based on the characteristic engineering loads. Finally, the performance management process compares the calculated performance results for the short and long terms with the service objectives to identify service or performance.

Fault Management

This process is similar to the real-time performance process except that it uses the collected alarms and fault statistics to detect and correct problems by pointing and correlating faults through the system. It simplifies the service provider's ability to monitor customer services by providing the status of the subscribed services.

One of the most challenging functions of the management process for IP networks is traffic engineering. It represents the action that the network should consider in order to relieve a potential servicing problem before the service is affected. This may include rerouting, load balancing and congestion control. Traffic engineering is also on for network dimensioning and planning as well as for capacity expansions. It is an optimization process that involves a set of algorithms that determine the required network resources (capacity) to meet a specific set of performance objectives.

The development of appropriate models for traffic engineering depends primarily on clear understanding of quality and grade-of-service requirements and the statistical characteristics of the traffic. Several traffic models and network dimensioning methods for packet networks have been proposed. Generally speaking, the models can be divided into two categories: those that exhibit long-range dependence (the fractional Brownian motion model on the on/off model with heavy-tailed distributions for the on/off duration) and Markovian models that exhibit only short-range dependence (on/off models with exponential on/off distributions, Markov-modulated Poisson process or Gaussian autoregressive models, which typically have exponentially decaying correlation functions). The on/off model has been proposed to model VoIP calls with alternating active periods and silent periods. The parameters of the on/off models can be estimated from actual traffic traces or by using typical default values. Finally, traffic engineering methods depend on the function of the network element. For example, traffic techniques for IP edge routers include packet classification, admission control and configuration management. Congestion management and congestion avoidance are typical considerations of backbone routers or switches.

Configuration Management

It deals with the physical and geographical interconnections of various IP network elements, such as routers, switches, multiplexers and lines. It includes the procedure for initializing, operating, setting and modifying the set of parameters that control the day-to-day operation of the networks. Configuration management also deals with service provisioning, user profile management and collection of operational data, which is the basis for recognizing changes in the state of the network. The main functions of configuration management are creation, deletion and modification of network elements and network resources. This includes the action of setting up an IP network or extending an already existing network, setting various parameters, defining threshold values, allocating names to managed IP objects and taking out existing network elements.

Security Management

This process includes authentication, authorization and other essential secure communication issues. Authentication establishes the identity of both the sender and the receiver of information. Integrity checking of confidential information is often done if the identity of the sending or receiving party is not properly established. Authorization establishes what a user is allowed to do after the user is identified. Authorization usually follows any authentication procedures. Issues related to authentication and authorization include the robustness of the methods used in verifying an entity's identity, the establishment of trusted domains to define authorization boundaries and the requirement of namespace uniqueness.

Accounting and Billing Management

This process deals with the generation and processing functions of end-user usage information [4.13]. This includes measuring the subscribers and possibly the network resources for auditing purposes and managing call detail information generated during the associated call processing. The records created in the application servers are of growing importance in IP networks. Such records are contents and services delivered by the network. Billing data collection and systems between the IP architecture and the billing platforms may aggregate usage-related data and usage detail records. The access usage detail can then be transferred to a billing system to render invoices to the subscribers that use IP services. Fraud detection and subscriber-related profile information, such as authorization to charge, are also a function of accounting and billing management.

4.4.3 ATM

The development of ATM is motivated by the merger of the computer networks and the telecommunications approaches toward multimedia communication. ATM technology was selected to support BISDN. Two groups work on the standardization of ATM: the telecommunications sector of the ITU-T, which is the international standardization organization for telecommunication, and the ATM Forum, which is a consortium of industrial and research organizations. Although originally designed for WANs, it is also suited for LANs and MANs. Therefore, it is a perfect vehicle for their seamless integration. ATM uses small, fixed-size cells (53 bytes, 5 for header and 48 for payload) and simple communication protocols to reduce per-cell processing and to speed up switching.

The ATM performs the following functions:

- Asynchronously multiplexes small packets called cells going from a number of information services to various destinations in a constantly flowing train if there is a seat for the specified destination with the required QoS
- Switches cells during transport if necessary
- Lets cells jumps off the train at the destination

The ATM uses a connection-oriented operation. It establishes a sequence of switches so that a connection is made from the source to the destination. Such a connection is called a Virtual Circuit Connection (VCC). The switches can be established to perform simplex, duplex, multicast and broadcast communications. A Virtual Connection (VC) is a connection between a switching node and the next node. Thus, a VCC consists of services of VCs. There are two kinds of VCs [4.14]:

- A Permanent VC (PVC) for a leased line
- A Switched VC (SVC) for a dynamically established connection

To simplify the management of VCs, a number of VCs with the same starting and ending nodes is grouped together as a virtual path (VP). To identify a VP as a VC, a number is used as the identifier and is labeled VP Identifier/VC Identifier (VPI/VCI).

Example 4.1 In Figure 4.5 VC1, VC2 and VC3 are grouped into VP1, VC4 and VC5 are grouped into VP2, and VP3 contains only one, namely, VC6.

A connection or call is assigned a VC. The end node uses the VC1 of the cell to direct it to the corresponding terminal. At the transit nodes, the VP1 of the cell provides enough information to direct this cell to the corresponding path in the network.

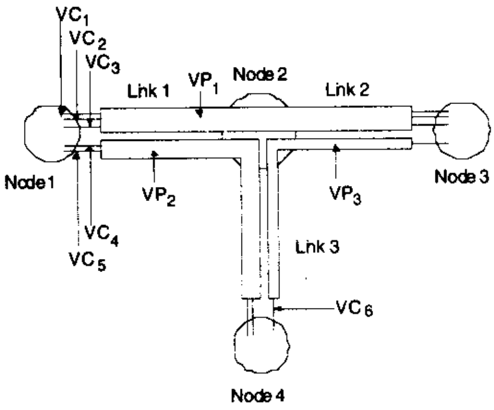

(a)

I.D.	Route
VP_1	Link 2
VP_2	Link 3
VP_3	Link 2

(b)

VPI, VCI	User Pay

(c)

Figure 4.5 The ATM VP concept: a) VP, VCs, links and nodes. b) Routing table at node 2. c) VP, VC identifiers in a cell [4.2]. ©1997 IEEE.

The ATM network has a layered structure allowing multimedia traffic to be mixed in the network. It includes the upper layer, the ATM Adaption Layer (AAL), the ATM layer and the physical layer. Figure 4.6 shows the relationship between the ATM layer structure and the OSI protocol stacks.

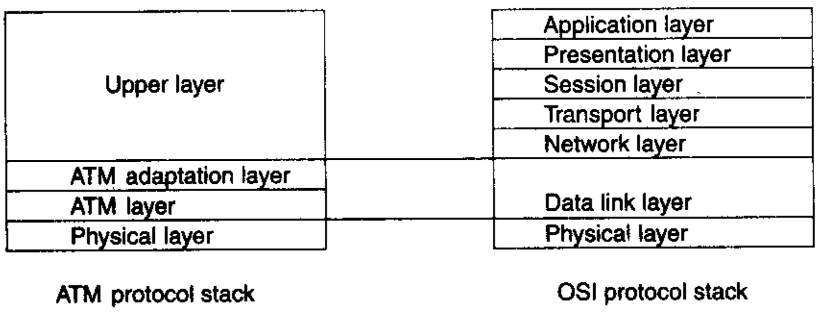

Figure 4.6 ATM protocol stack versus OSI protocol stacks [4.2]. ©1997 IEEE.

The upper layer includes the higher layer protocol, such as TCP, RTP and Xpress Transport Protocol [4.15].

The AAL layer adapts the upper layer protocols to the ATM format. It inserts or extracts user information as 48-byte payloads. The AAL layer consists of two sublayers: the Convergence Sublayer (CS) and the Segmentation and Reassemble (SAR) sublayer. The CS converges different types of user traffic and encapsulates/decapsulates data flow to and from the SAR sublayer. The SAR sublayer in the sender segments data into a series of 48-byte cells, and, in the receiver, it reassembles cell sequences back to the original data.

The ATM layer adds or strips the 5-byte header to or from the payload. For the sender, it takes 48-byte data from the AAL and adds the 5-byte header that contains routing information to ensure that the cell is sent on the right connection. For the receiver, it strips the 5-byte header and passes the 48-byte payload to the corresponding AAL.

The physical layer defines the electrical characteristics and network interfaces and places ATM cells into the transmission medium.

Figure 4.7 shows the rate of the ATM layered structure in the ATM network.

The ATM is suitable for multimedia communication because it provides a guaranteed QoS. ATM does not prevent cells from being lost, yet it guarantees that the cell order is always maintained in a connection. QoS is conceptually negotiated between three entities: the calling party (initiator of the connection), the network and the called party. The calling party requires a connection to the called party with a SETUP message, in which it provides its QoS requirements to the network and to the called party.

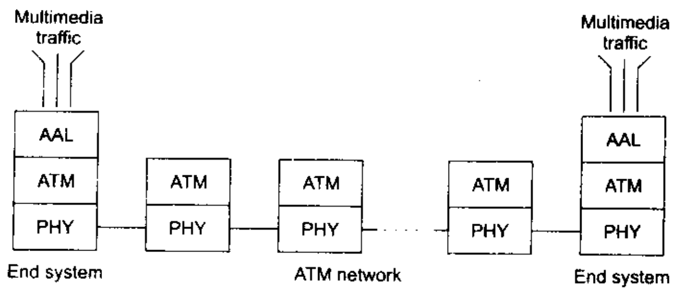

Figure 4.7 An ATM network [4.2]. ©1997 IEEE.

The QoS parameters supported for ATM connections differ slightly from those that are considered in networks that are based on variable length packets. The ATM parameters are the following:

- **Sustainable rate**—The minimum number of cells per second that must be supported by the network for the entire length of a connection
- **Peak rate**—The number of cells that must be expected at each node in the network in rapid succession (one burst)

- **Maximum burst length**—The length of an interval in which at most one burst must be expected by a network node
- **Cell loss ratio**—The maximum rate of lost or corrupted cells that an application can accept for a connection
- **Maximum end-to-end delay**—The restriction on the sum of all waiting times that each cell can stand in the queues between the sender and the receiver of the cell
- **Maximum cell delay variation**—The maximum difference in end-to-end transmission time that two cells of a connection can experience

To adapt to different characteristics of the traffic, ATM provides five types of adaptation. Type 1 is for circuit emulation at a Constant Bit Rate (CBR) service for isochronous data. Type 2 is for VBR connection-oriented service for isochronous data. Type 3 is for connection-oriented data service. Type 4 supports the services such as connectionless data communications. This provides the features equivalent to AAL Type 3 features. Type 5 is for LAN emulation and all other possible traffic. The ATM Forum also defines Available Bit Rate (ABR), which guarantees a minimum rate, but the delay may vary. The Unspecified Bit Rate (UBR) is defined, too. It is similar to ABR, but does not guarantee a minimum rate, and cells may be lost due to congestion.

4.4.4 Integration of IP and ATM

The strength of IP is its large installed base. The Internet is a packet-switched network and basically provides connectionless and best-effort transport services. The major problem is that the best-effort delivery of messages does not support the various QoS required in a true integrated service network. In addition, the existing Internet has very limited bandwidth, although there are plans to upgrade it to higher bandwidth. The development of IP multicast, resource reservation protocols and real-time transport protocols allows the Internet to provide limited forms of multimedia services. Even with the upgraded bandwidth, the Internet will be inadequate for many multimedia applications, for example, VoD. On the other hand, the strength of ATM lies in the possibility of setting up connections with a class of service that matches the requirements of the application involved. The planned bit rates and bit error rates also meet the requirements of real-time transmission of audio and videostreams. ATM works best in an environment where everyone uses ATM. The global interworking protocol will be the one that supports QoS guarantees and may include both ATM and IP technologies.

Although the Internet community is sponsoring ATM as the promising high-speed bearer technology of subnetworks, the vision of the Internet is still focusing on the protocol suite to be executed excessively among hosts at network edges and routers in subnetwork boundaries. The perspective is to consider ATM as a high-speed data pipe technology and to guarantee QoS by means of RSVP [4.16, 4.17, 4.18].

Cooperation between RSVP and ATM signaling protocols to guarantee QoS is an attractive perspective for building novel broadband networks to carry Internet traffic across public infrastructures and across enterprise Internets. The underlying subnetwork layer protocols are

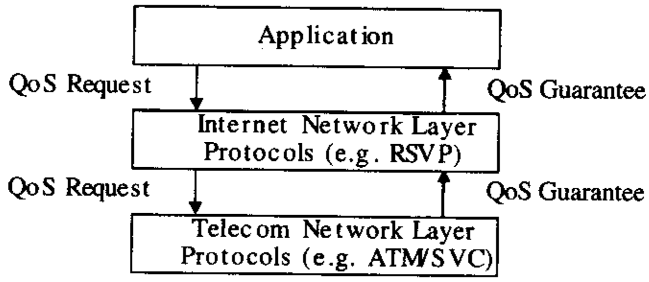

Figure 4.8 Application support for QoS [4.18]. ©1997 IEEE.

cooperating with the Internet network layers protocols to provide QoS. Application support for QoS is shown in Figure 4.8.

A number of proposals have been introduced to integrate IP and ATM protocols. Examples are IP switching [4.19] and tag switching [4.20] schemes. The former addresses a radical substitution of conventional routes and switching technologies to build high-speed IP backbones. The latter suggests a smoother overlay approach to be implemented across conventional routers and switching technology.

IP switching nodes dynamically shift between store-and-forward and cut-through switching [4.19]. IP switches base their operation on the traffic flow concept. A traffic flow is a sequence of IP packets sent from a particular service to a particular destination, sharing the same protocol type, the same type of service, and other characteristics as determined by examing information in the packet header. Prior to cut-through, an IP switch acts like a router with store-and-forward routing of IP datagrams.

Tag switching enables the forwarding of IP packets directly across a network of tag-switching-compliant routers [4.20]. A tag-switching network is made of two types of nodes: tag-edge routers and tag switches. The former are routers at the edge of the tag-switching network that apply tags and perform Internet network layer functions. The information can be carried in data units in two ways:

- As a part of the ATM layer header
- As a part of the Internet network layer header

Tag switching is independent of the routing protocols implied.

4.4.5 Real-Time Multimedia over ATM (RMOA)

The communication industry is currently engaged in extensive work on two flavors of voice over packet transport: Voice Over ATM (VoATM) and VoIP. VoATM is likely to be successful in delivering toll quality with special gateways preparing native Public Switched Telephone Network (PSTN) voice for ATM transport while using the QoS features of ATM technology. VoIP is becoming increasingly important for enterprises that can more easily provide resources to deliver quality to the voice traffic within their corporate networks.

Because ATM is becoming increasingly ubiquitous in the core of service provider networks, the ATM Forum developed an efficient and scalable means to transport native H.323 VoIP traffic over ATM that is not possible with the existing IP over ATM solutions. This effort was carried out within the RMOA working group and defined a new type of gateway called H.323-H.323 gateway [4.21].

ITU-T has created the H.323 standard describing the system components call model and signaling procedures to be used by entities engaged in multimedia communications across a network such as the Internet.

Example 4.2 The H.323 standard does not impose any QoS requirements on the network used to carry the H.323 media streams. We will focus on the scenario depicted in Figure 4.9. Here, an Internet backbone has an ATM core where the devices (routers or switches) at the edges of the core (devices X and Y) are capable of receiving H.323 RTP/UDP/IP packets and sending them across VCs in the ATM core.

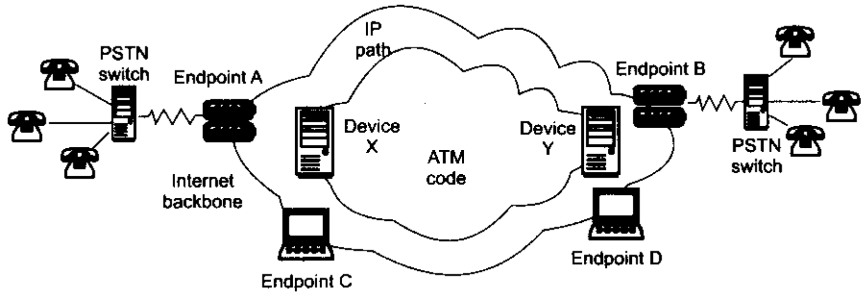

Figure 4.9 The ATM core in an Internet backbone [4.21]. ©2000 IEEE.

The H.323 endpoints in Figure 4.9 are IP hosts, and the H.323 media traffic traverses an access IP path to reach devices X and Y at the edge of ATM core. Two types of H.323 endpoints are illustrated: gateways (endpoints A and B) and terminals (endpoints C and D). H.323 terminals are capable of initiating and receiving H.323 calls, and H.323 gateways do the same on behalf of the other non-H.323 terminals, such as PSTN phones. This particular type of H.323 gateway provides a form of VoIP service.

Example 4.3 Figure 4.10 elaborates on the scenario in Figure 4.9, focusing on an H.323 call between endpoints A and B. It indicates the access IP paths used by the traffic originating at the end points and introduces another standard H.323 entity, the gatekeeper. Here, we have three gatekeepers used in the signaling architecture to break the end-to-end call into three legs. The different legs result from control and media termination forced devices Gateway (GW) X and GW Y at the edge of the ATM core.

An H.323 gatekeeper provides the services of Registration, Admission and Status (RAS) through the RAS channel, which also includes address resolution. Endpoint A will use the RAS

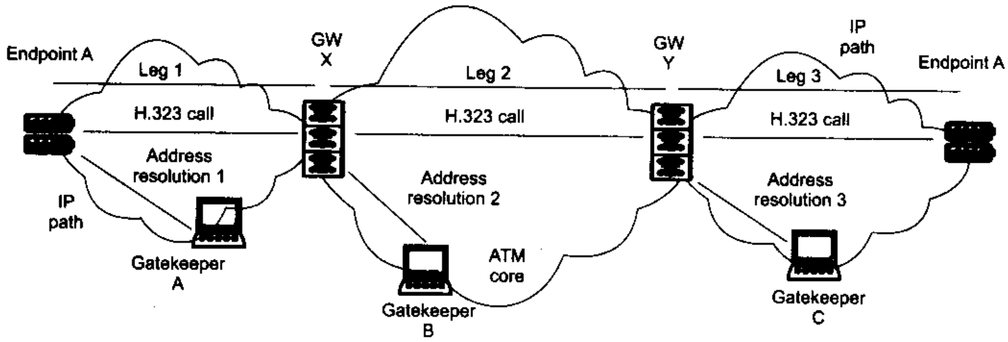

Figure 4.10 The ATM core in an Internet backbone scenario focusing on an H.323 call between endpoints A and B [4.21]. ©2000 IEEE.

channel to gatekeeper A, with which it is registered, to place calls to endpoint B. Usually, endpoint A will know an alias address for endpoint B. Gatekeeper A will translate this address to an IP address and a transport protocol port number. Three gatekeepers are illustrating a scenario with three zones where endpoints A and B belong to different zones. As such, address resolution 1 provided by gatekeeper A tells endpoint A to use GW X as a gateway for calls to endpoint B.

H.323 includes two other ITU-T standards: H.225.0 and H.245. These define control messages to be exchanged among H.323 endpoints.

The H.225.0 signaling establishes the H.323 call. This implies the opening of a call-control channel, subsequently used to open media channels through H.245 procedures. Similar to H.225.0, the H.245 channel is defined by a set of transport addresses used for the exchange of H.245 control messages. The initial H.225.0 messages carry these transport addresses that are replaced with local addresses for the gateways as the H.225.0 messages are relayed across adjacent legs of a call. These addresses' replacements force the H.245 messages to be addressed to the H.323-H.323 gateways, breaking the end-to-end H.245 channel into three legs as well.

4.5 Multimedia Operating Systems

OSs manage computer resources, for example, CPU, memory, I/O devices, and so forth. They also hide the physical characteristics of the underlying hardware and provide an efficient and convenient environment for end-users. A multimedia OS extends the functionalities of OSs to accommodate multimedia data manipulation to provide an environment for real-time support. The major concerns include real-time support while simultaneously running traditional applications efficiently. The major concerns include real-time processing and QoS-based resource-management. A multimedia OS may be developed as an extension of a traditional OS or constructed using the microkernel architecture [4.22]. It should provide CPU management, memory management, I/O management and file system management.

CPU Management

Real-time processing can be achieved through efficient real-time scheduling. In the context of continuous media, a deadline can be the acceptable playback time of each frame. Therefore, it is a soft deadline and appears periodically. The challenges of multimedia scheduling are due to two conflicting goals: non-real-time processes and real-time processes. Non-real-time processes should not suffer from the execution of real-time processes, because multimedia applications equally depend on discrete and continuous media data. Real-time processes should be allowed to pre-empt non-real-time processes or other real-time processes with lower priorities.

The most important real-time scheduling approaches include Earliest Deadline First (EDF) and rate monitoring scheduling [4.23]. With EDF, each task is preemptive and is assigned a priority according to the deadline. The highest priority is assigned to the job with the earliest deadline, and tasks are executed in a priority order. When a new task arrives, the scheduler recomputes the priorities of all pending tasks and then reorganizes such that the order of the task being executed is preempted and the new task gets served immediately. The interrupted process is resumed later from the interruption point. Otherwise, the new task will be put in an appropriate position.

With rate-monotonic scheduling, each task is pre-empted and is assigned a priority according to the request rate. The highest priority is assigned to the job with the highest rate. In contrast to EDF, such assignments are performed only at the connection establishment time and are maintained through the lifetime of the connection. For preemptive periodic tasks, rate-monotonic scheduling is optimal in the sense that no other-static algorithm can schedule a task that the rate-monotonic algorithm cannot also schedule [4.24].

Comparing these two algorithms, EDF is more dynamic. It has to be executed frequently and thus incurs higher scheduling overhead. The advantage is that it can achieve processor utilization up to 100%. On the other hand, a rate-monotonic algorithm is static because the priority assignment is only calculated once. Because the priorities are assigned according to the request rate, more context switches occur in rate-monotonic scheduling than EDF. The worst-case upper bound of the process use is about 69% even though, on the average, the use is suitable for continuous media applications because it has no scheduling overhead and is optimal for periodic jobs.

Memory Management

The memory manager allocates memory to processes. Continuous media data is typically very large in size and requires stringent timing requirements. One solution is to avoid swapping and to lock continuous media data in memory during the process [4.24]. This approach, however, may affect resource use. Other important practical implementation techniques include using scatter buffers and passing pointers. With scatter buffers or scatter loading, the address space of a process is loaded into possibly discontinuous regions of memory. This tends to be more space efficient than loading into a single continuous region, but may result in fragmentation. With passing pointers, objects are passed by reference rather than having to pass the objects themselves. This may result again in more efficient usage of memory space.

IO Management

The main function of the I/O subsystem is to transfer multimedia information between the main memory and the network adapter or multimedia peripherals (camera, loudspeaker, CD-ROM drive, microphone, disk, keyboard, monitor, and so forth). The important issues include device management, interrupt latency, and real-time transmission. Device management integrates all hardware components and provides a uniform interface for the control and management of these devices. Multimedia applications are I/O intensive. The continuous media frames will not frequently interrupt the kernel and lower the system throughput. There are three strategies to alleviate this problem: changing the internal structure of the kernel to make it highly preemptive, including a set of safe pre-emption points to the existing kernel or converting the current kernel to a user program and running on top of a microkernel [4.25]. Real-time I/O is necessary for most multimedia applications to ensure the continuity of each stream and the synchronization of multiple related streams. With advances in networking and communication technologies, it is possible to achieve network bandwidth well above a gigabit per second. The network I/O becomes the bottleneck that limits overall system performance. Therefore, the focus is to improve I/O system throughput.

File System Management

File management is responsible for making efficient use of the storage capacity and for providing file access and control of stored data to the users. Multimedia file systems demand additional real-time requirements to handle continuous media streams.

4.6 Distributed Multimedia Servers

Servers are an integral part of the multimedia environment such as digital libraries, in-house training systems, VOD or near-VOD services and so forth. In a typical environment, a multimedia server is connected through an interconnection to clients who request information on demand. The servers not only store and provide information to clients, but they perform various management operations as well (billing, accounting, encryption and so forth). One of the most important elements of a multimedia service environment is a videostream. Videostream elements require high I/O bandwidth at the server for their delivery and large amounts of memory for their storage [4.26].

A number of issues make the design of the video server difficult. First, a video server needs to provide video services simultaneously to multiple clients, to guarantee the QoS to multiple clients and to guarantee the QoS to each client. Second, a video server needs to manage system resources, including CPU, disk and memory, and it needs to schedule network activity to get the maximum use from the resources, while not overloading the system. Third, a video server needs to be able to support a variety of operations, such as playback, fast forward, slow forward, pause, resume, indexing and scrolling. Finally, a user watching a video may change from one service to another service (for example, from playback to fast-forward or from playback to slow forward). A video server should support these dynamic service changes while efficiently using system resources [4.27].

In general, servers can be classified into two large categories: centralized and distributed. In centralized servers, a high-end system stores servers and manages the video streams, and in a distributed server, a collection of workstations or PCs may constitute the server. A lot of effort has been concentrated on distributed server environments because it is not clear which type of server provides the most cost-effective solution [4.28, 4.29, 4.30].

In order to describe precisely the activities between a client and a server, we can divide them into three levels: session, transaction and services. A session is a connection between a client and a server. Its lifetime is defined as the duration from the login to the logout of a video client. Within a session, a user can execute two types of transactions, query or video. A query transaction asks for metadata of the video server. A video transaction is a sequence of video-related activities and is encapsulated by the open and close of a video file. The major difficulties between a query transaction and a video transaction are that the latter involves time-critical, video-related activities, whereas the former does not. After a video transaction is established, a user can initiate a video service, terminate a video service, or switch from one video service to another where a video service is a playback, fast forward, or slow forward operation on the chosen video [4.31].

The most important parameters in a multimedia server are its I/O bandwidth requirements and its storage requirements. The I/O bandwidth designates how many clients can be simultaneously served. On the other hand, the available amount of storage determines the number of videostreams that can be stored in the server. In a centralized server, one can easily observe that no client request for a videostream can be blocked as long as I/O bandwidth is available. This is not necessarily true for a distributed server [4.32].

In distributed multimedia servers where a client requests different videostreams that may have different probabilities, placement of videostreams is an important parameter because it may result in an unbalanced request to the system's stations, and thus to high blocking probabilities of requests.

Example 4.4 Assume that a client request arrives at the server for a stream that is stored on a station with I/O bandwidth that is allocated to other served streams. Although aggregate I/O bandwidth may be available in the system, the request has to be blocked. One solution to resolve such problems is to provide copies of videostreams in multiple stations so that alternative stations can serve a request. This demonstrates that placement and replication of videostreams is an important issue in distributed servers.

4.6.1 Multimedia Packing

Multimedia packing is a method that achieves load and storage balancing in the distributed multimedia server with two basic operations:

- Placement of videostreams in server stations and replication of a small subset of videostreams
- Weighted scheduling of client request for the replicated videostreams

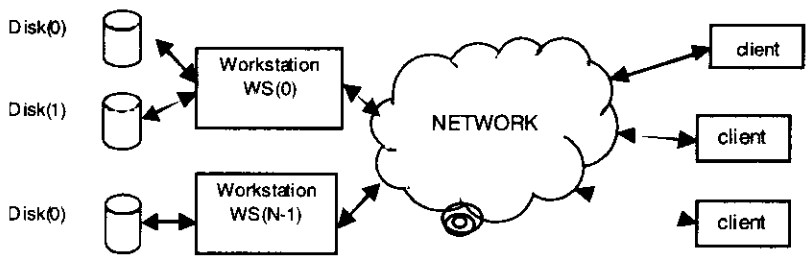

Figure 4.11 Distributed server configuration environment.

This results in the probability of a video request being the same for all stations in the distributed server. When client requests are uniformly routed to disks (or servers, assuming that a server corresponds to one disk), the blocking probability of a client's request is minimized [4.30]. In addition to traffic load balancing, multimedia packing achieves storage balancing, because it results in approximately an equal number of videostreams stored in all stations. The number of videostreams stored on any two workstations differs by a number less than or equal to two.

Consider a distributed server configuration in an environment as shown in Figure 4.11. A distributed video server is employed to deliver video services to clients. The server is composed of N stations WS(0),...,WS(N-1) storing a total of M different videostreams $VS_0, \ldots, VS_{M-1}$ and is connected to a total of C clients through a network. Here, we have $M >> N$. Clients generate requests for videostreams to the server. Videostreams are requested with different probabilities: videostream VS_i is required with probability P_i. Thus, we have $\sum_{i=0}^{M-1} P_i = 1$. Given a placement of videostreams to the workstations, the probability that a client's request is routed to a certain workstation WS(i) is equal to the cumulative probability SP_i of the videostreams stored on WS(i). Multimedia packing achieves a cumulative probability $SP_i = 1/N$ for $0 \le i \le (N-1)$. Multimedia packing also achieves storage balancing by resulting in a configuration with $|NS_i - NS_j| \le 2$ for all i,j. Here NS_i is the number of streams stored on workstation WS(i).

4.7 Distributed Multimedia Applications

Multimedia integration and real-time networking create a wide range of opportunities for multimedia applications. Distributed multimedia applications have several requirements with respect to the service achieved for them by the communication system. These requirements depend on the type of application and on its usage scenario. Furthermore, the requirements of applications regarding the communication services can be divided into traffic and functional requirements. The traffic requirements include transmission bandwidth, delay and reliability. They depend on the quality of the datastreams. The functional requirements are multicast transmission and the ability to define coordinated sets of unidirectional streams. On the other hand, the reliability requirements are sometimes lower than for traditional communication applications, for example, if a fault-tolerant data-encoding scheme is used. Furthermore, retransmissions, which are traditionally used for the provisioning of reliability, increase the end-to-end delay and are often worse than lost data for multimedia applications.

The traffic requirements can be satisfied by the use of resource management mechanisms. They establish a relationship between transmitted data and resources and ensure that the audio-visual data are transmitted in a timely manner. For this, during the transmission of data, the information about the resource needs must be available at all nodes participating in the distributed applications, that is, end systems and routers. Hence, a resource must be reserved, and a state must be created in these nodes, which means that a connection is established. This connection should then be used for the transmission of data.

For various multimedia applications, multiple receivers are interested in receiving the same data. For instance, in a talk distributed using the network, all listeners must receive the same data. Sending each person a single copy wastes resources because, for parts of the path from the sender to the receivers, the same nodes are traversed. Thus, multicast should be used, which provides for the transmission of a single copy of data to multiple receivers, as shown in Figure 4.2. In addition to reduced network load, multicast also lowers the processing load of the sender. Multicast must not be limited to a single sender. In conferencing scenarios, it is usual to have several senders who normally do not use the resources at the same time.

The delivery of audio-visual data to large receiver groups, such as the distribution of Internet Engineering Task Force (IETF) meetings across the multicast backbone, must also take into account that the resource capabilities and the participations can vary widely from high-speed network links and fast workstations to low-end personal computers connected using relatively narrow band links. Therefore, support for heterogeneous systems must be provided (heterogeneous is with respect to networks as well as to end-system capabilities).

According to the different requirements imposed upon the information, communication and computing subsystems and distributed multimedia applications may be classified into three types: ITV, telecooperation and hypermedia.

4.7.1 ITV

ITV requires a very high transmission rate and stringent QoS guarantees. It is therefore difficult to provide such broadband services across a low-speed network, such as the current Internet, due to its low bandwidth and best-effort-only services. ITV typically demands point-to-point switched connections, good customer services and excellent management for information sharing, billing and security. The bandwidth requirement is asymmetric in that the bandwidth of a downstream channel that carries video programs from the server to the user is much higher than that of the upstream channel from the user to the server [4.33].

The four main components of ITV systems are a home terminal commonly known as the Set-Top Box (STB) or Customer Premises Equipment (CPE), an access network, a network-based server and a powerful user-friendly interface. The STB (or CPE) typically takes the form of a box sitting on the top of the TV set. This STB connects to both the television and an external communication network using a subscriber drop or loop. When interactive services are offered, the complexity and cost of the STB increases. The new STB must provide both the analog functions and the digital functions of audio and video decompression, demodulation to recover the digital feeds, decryption, an upstream modem for communicating consumer control requests back to the program source and a user-friendly interface [4.33]. ITV servers are a collection of computing, storage and communications equipment that implements interactive video services. A service may require that more than one server be implemented, or a server may implement more than one service. All subsystems of the server communicate with one another using local high-bandwidth interconnect and a switch. This architecture can be used to scale capabilities of the server incrementally, and it provides isolation between the various subsystems. Two key components of the server technology are logical organization of the multimedia samples in a file system or database and techniques by which media components can be continuously recorded or played back from the server. The multimedia database server must ensure that the recording and presentation follow a real-time data rate. User interface designs for ITV are more involved than those for standard TV, owing to the richness of the types of possible interactions with the consumers. They will vary among different applications. An important ITV service is VoD.

VoD

VoD provides electronic video-rental services across the broad band network [4.34]. Customers are allowed to select programs from remote massive video archives, view them at the time they want without leaving the comfort of their homes and interact with the programs. A VoD system that satisfies requirements, like any video, or any VCR or like user interaction, is called a true VoD system and is said to provide true VoD services. Otherwise, it is called a near VoD system. One way to allow true VoD services is to have a dedicated videostream for each customer. This is not only expensive, but is wasteful of the 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 may be used. This allows multiple users accessing the same video to be served by the same videostream. Although batching complicates the provision of user interactions, it increases the system capability in peaks of the number of customers that the system can support [4.35, 4.36, 4.37].

Example 4.5 In a VoD system, multiple identical copies (streams) of the same video program are broadcast every five minutes. A user is served by one of the streams, and user interactions are simulated by jumping to a different stream. Not all user interactions can be simulated in this fashion. Even for those that can be simulated, the effect is not exactly what the user requires. For example, one cannot issue fast forward because there is no stream in the fast-forward mode. Also, one may pause for 5, 10, 15 and so forth minutes, but not for seven minutes from the original stream.

A protocol called Split-and-Merge (SAM) has been proposed to provide true VoD services while fully exploiting the benefit of batching, thereby reducing the per-user video delivery cost.

To implement the complete set of interactive services, a VoD system contains many components, including the video server, transport network, Subscriber Terminal Unit (STU) and ser-

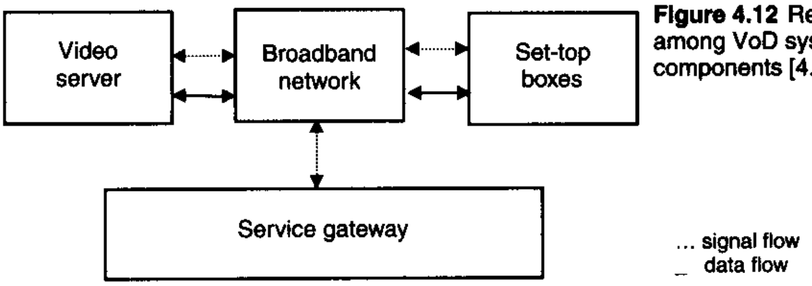

Figure 4.12 Relationships among VoD system components [4.2]. ©1997 IEEE.

vice gateway. The VoD system architecture and the relationship among its components is shown in Figure 4.12.

A video server consists of massive storage and media controllers (Figure 4.13). The video server stores a large number of digitized videos and serves a considerable number of simultaneous video requests to the same or to different videos on demand. The storage media consist of magnetic disks, optical disks, and magnetic tape and are usually organized hierarchically for cost-effectiveness. Under this configuration, popular videos are stored in the disks. Less popular ones are stored in tape devices with terabyte capacity and are retrieved as necessary to the disk drive for processing. The video server may be located at the local or regional switch of the network provider or at remote information archives. The basic functions supported by video servers include request handling, random access, and user interactions, in addition to admission control and QoS guarantees. The networked video jukebox and the video library system are two examples of video server prototypes [4.38, 4.39].

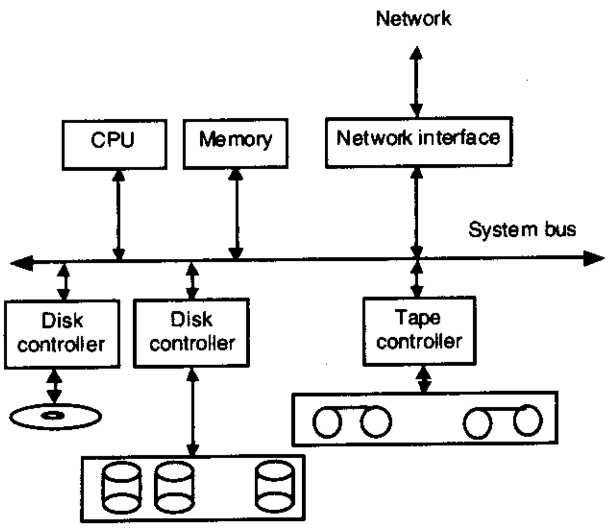

Figure 4.13 Video server architecture [4.2]. ©1997 IEEE.

The transport network delivers video programs from the video server to the customers. The network must have a very high data rate and must satisfy the real-time delivery constraints of video traffic. It consists of two major components: the backbone network with high-speed switches and the local access network. The backbone network links the remote video server at geographically dispersed locations and the regional, national or international information archives. The trend is toward a Synchronous Optical Network (SONET) backbone with ATM switching because of the low error rate, high data transfer rate, bandwidth on demand and seamless services.

The STU (or set-top box) along with the television monitor and the infrared controller (that is, remote control) serves as the bridge between the subscribers and the system. The major functions of STU include receiving the incoming videostreams; demodulating; demultiplexing and decoding the signals; performing the necessary signal conversion, such as D/A transformation for playback on the TV monitor; and sending outgoing control messages. STU must accommodate the heterogeneity of technologies and formats from various controls to services [4.40]. The usefulness of an STU is in its adaptation to the diversity of access networks, service providers, applications and user interfaces.

A service gateway 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 the following:

- Directory services to provide menu browsing and program scheduling
- Mapping from service identity to corresponding location and program provider
- Controlling, coordinating and signaling for multimedia session establishment, maintenance and disconnection
- System management, including operation management, fault management, configuration, resource management and performance management
- Subscriber profile maintenance and billing
- Secure communication to prevent unauthorized access, including authentication, encryption and scrambling

Video server placement is an important design issue in VoD systems. The alternatives include centralized video servers, hierarchical video servers and fully replicated distributed video servers. A centralized video server system is relatively simple to manage. All requests are sent to and served at one site. Hierarchical server placement exploits the user access pattern and the nonuniform popularity of videos. The distributed server system distributes the video copies to many switches located closer to the users, thus alleviating the congestion in the network and the bottleneck due to the central server, but at the expense of higher cost [4.41, 4.42]. However, managing distributed servers is more complex. One has to decide which video and how many copies to maintain at each distributed server. In addition, due to varying rates of requirements, the video offerings at each distributed server need to be changed periodically. Which alternative is highly preferable depends on the tradeoff between storage and communication costs, the application needs, the underlying infrastructure and other factors. In Li et al. [4.34], a performance model that may be used to evaluate the requirements of network bandwidth and server storage is proposed. Hence, we obtain the tradeoff between communication and storage costs for various placement alternatives.

4.7.2 Telecooperation

Telecooperation, also known as Computer-Supported Cooperative Work (CSCW) refers to a system that provides an electronic shared workspace to geographically dispersed users with communication, collaboration and coordination supports [4.43, 4.44, 4.45]. Group communication provides an electronic channel for the users to exchange messages either synchronously or asynchronously. It allows individuals to cooperate regardless of time and location constraints. Sharing information among groups is the key to effective collaboration.

Telecooperation requires multicast, multipoint and multiservice network support for group distribution. In contrast to the strict requirement on the video quality of ITV, telecooperation, such as videophone and desktop conferencing, allows lower picture quality and therefore has a lower bandwidth requirement. It is possible to provide such services with the development of real-time transport protocols across the Internet. Telecooperation requires powerful multimedia database systems rather than continuous media servers with the support of visual query and content-based indexing and retrieval.

CSCW may be classified into four different interactions, which are shown in Figure 4.14:

- Centralized synchronous
- Distributed synchronous
- Centralized asynchronous
- Distributed asynchronous

Synchronous and asynchronous refer to the time dimension, while centralized and distributed refer to the space dimension. Synchronous exchanges demand real-time communication, but distributed interactions require broadcast or multicast support for group distribution.

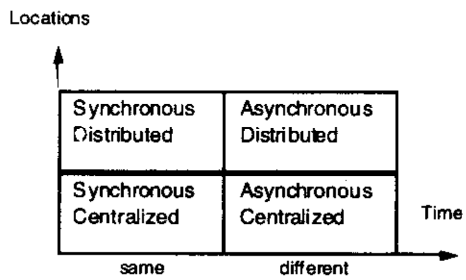

Figure 4.14 Classification of interactions by time and space for CSCW [4.2]. ©1997 IEEE.

The centralized synchronous mode requires face-to-face interactions. Examples are applications in the meeting room. The distributed synchronous mode provides real-time interaction in groups dispersed at different locations. Examples include network chat, real-time joint editing, multimedia conferencing and videophone. This type of application poses the greatest challenge in the design of group collaboration systems. The centralized asynchronous mode refers to those activities held at the same place, but at different times. News groups and electronic bulletin boards are such examples. The distributed asynchronous mode allows the exchange of messages within the group asynchronously. Electronic mail and voice mail are examples.

Telecooperation Infrastructure

A telecooperation infrastructure provides a robust framework to facilitate group work and to share information. It consists of a network model, a system model and a communication protocol model [4.46]. The telecooperation network model defines the functional elements and their relationships in the telecooperation system. It includes a multicast, multipoint and multiservice network and a collection of group collaboration agents. The group multicast, multipoint and multiservice network connects and distributes multimedia materials to remote participants. Figure 4.15 contains the block scheme of the network model. The collaboration agent includes the hardware and software that provide the necessary facilities and functionalities for cooperation and management of group work.

The telecooperation system model consists of five major modules as shown in Figure 4.16: cooperation control, application sharing, conferencing, interface and database. The cooperation control module administers a dynamic set of participants during cooperation sessions. The major functions include access control, group dynamic control and floor control. Access control validates membership in a group activity, registers a session, initiates or closes a session, and modifies the membership from one session to another. Group dynamic control allows participants to add in or drop out of a session dynamically. Floor control allows only one participant to own the floor at a time, that is, only one user can interact with programs at a time during the session. Floor control policies have two variations: centralized control and distributed control [4.47]. The

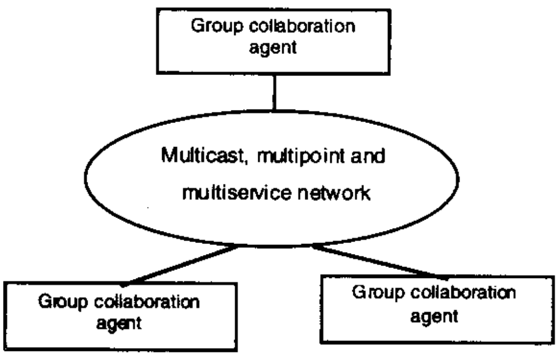

Figure 4.15 The telecooperation network model [4.2]. ©1997 IEEE.

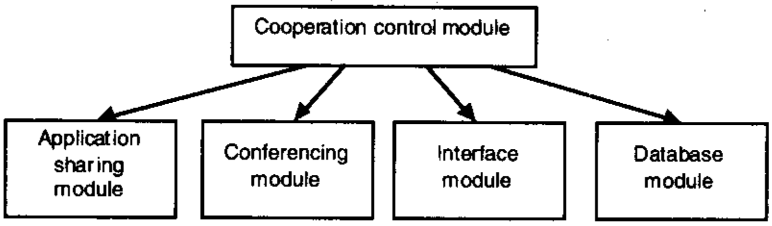

Figure 4.16 The telecooperation system model [4.2]. ©1997 IEEE.

application-sharing module handles the shared activities of the participants. It provides outputs to all participants with appropriate views while coordinating and controlling the inputs to satisfy the cooperation.

The application share module has two possible implementations: centralized or replicated. In the centralized approach (that is, client-server architecture), shared applications are run in a dedicated server. This control server processes all user requests and distributes the results to all local machines for display. In the replicated approach, every local machine keeps and runs a copy of the shared applications, and the input and processing events of the floor holder are distributed to all other sites [4.48]. This approach is more tolerant of machine failures and more scalable for large and heterogeneous user environments. It also facilitates the support of different views at the local displays. The major characteristics with the replicated approach are the maintenance of data consistency and cooperation control among the participants. The conferencing module supports asynchronous messaging and real-time audio-visual communications for distributed, multiparty collaboration. For video and audio transmissions, full-duplex communication channels and real-time transport services are required among the participants. The interface module is concerned with the display of the shared data. It supports private and shared windows for the participants to share only part of the information and allows the shared information to be presented to different users in different ways. The database module stores shared data and knowledge. Concurrency control is required to resolve the conflicts in the shared data and shared operations between participants. In contrast to database concurrency control (for example, locking shared data for exclusive access), group concurrency control needs to consider real-time and cooperation issues [4.49, 4.50].

The communication protocol model provides the protocols to exchange information with groups. Two kinds of protocols are required: user presentation and group work management. User-presentation protocols are concerned with client interactions, such as opening, closing or dynamically joining and leaving a telecooperation session. The group-work management protocols perform the communication between clients and servers, such as registering active sessions and inquiring about the current status of cooperative work.

Telecooperative Applications

Of the telecooperative applications, three are important:

- Multimedia email

- Collaborative authorship applications
- Multimedia conferencing

Multimedia email is an electronic messaging system that allows all moments to exchange multimedia messages asynchronously. Email is the most widely used service on the Internet. Unlike the older Simple Mail Transfer Protocol (SMTP) standard, which understands only ASCII characters, Multipurpose Internet Mail Extensions (MIME) specifies a set of encoding rules and header extensions that describe new types of contents (image, video or audio) and the structure of multimedia messages to embed multimedia information.

Collaborative authorship applications mean the activity of collaborative editing and composing a multimedia document by a group of people. This type of application can be either synchronous or asynchronous. Each member works on a part, and the completed product is the combination of all individual parts [4.51, 4.52].

Multimedia conferencing supports participants with distributed multiparty synchronous collaboration to simulate face-to-face interactions or real-time telephone conversations. The service may range from point-to-point videophone to multipoint conferencing. In videophone, the hardware requirements include a telephone equipped with a video camera, which transmits low-quality images at low frame rates through existing phone lines. In desktop conferencing, the hardware requirements include desktop PCs or workstations equipped with audio-visual facilities. The electronic meeting room, or Group-Decision Support System (GDSS), uses several networked workstations, large computer-controlled public displays and audio and video equipment to simulate a face-to-face meeting electronically [4.53, 4.54].

A video conferencing system uses both intraframes and interframes. The intraframe is only sent for the first picture or after a change of scene. Intraframe does not have motion estimation for the DCT [4.55], quantization, zigzag scan, and variable length and Huffman coding are used for each macroblock. The inverse quantization and the IDCT for the quantized frame form a referenced frame. The input frame uses motion estimation by comparing the input frame to the referenced neighbor frame (the reconstructed neighbor picture) to find the motion vectors. If the difference between the input block and the referenced block is below the threshold, no information need be sent. Otherwise, the difference is transformed by DCT, quantized, zigzag scanned and coded using variable length and Huffman coding. A referenced frame can be generated by reconstructing the frame using inverse quantization and the IDCT on the quantized difference and by adding the motion-compensated picture to this difference. A loop filter that removes the high-frequency noise can be used to improve the visual effects.

There are two approaches to implementing CSCW systems: collaboration transparency and collaboration awareness [4.48, 4.56]. Collaboration transparency performs multiuser cooperation on existing single-user applications using an application-sharing system. One key feature of collaboration-transparent applications is that all the users have to use the same application-sharing system and usually only one participant is able to access or control the shared windows [4.43]. Collaboration awareness performs multiuser access using the development of a new special-purpose application explicitly to handle collaboration. This approach embeds collaborative facilities within the applications and usually allows multiple users simultaneous access to the shared windows.

Telemedicine

Although a telemedicine concept is very simple (we acquire medical data from appropriate devices and transfer it to other centers), its realization is very difficult due to very hard technical requirements, particularly to transmit store and search for an extremely large number of large files, as medical images [4.57]. Within the last 10 years, various image-processing, transmission and archiving systems have been developed for medical applications. These have been focused in the areas of radiology and pathology. They are now finding their way into such areas as cardiology, neurology, orthopedics and surgery. The medical images, acquired from different imaging devices—Computer Tomography (CT), Magnetic Resonance Imaging (MRI), Nuclear Medicine (NM) imaging, Ultrasound (US) imaging, different radiology images, images from digital microscopes, and so forth—are in different formats having different spatial and level resolutions. The large volume of image data requires image compression. Different approaches are derived for ensuring the required level of QoS with cost-effectiveness [4.58]. Videostreams are usually compressed before being transferred across a network. Experimental studies with the lossy compression algorithm, such as JPEG or wavelet-based transforms, confirm that compression ratios of 10:1 or 20:1 produce no perceptible differences in the quality of the medical image. At the user level, a perceived QoS is defined as the percentage of diagnosis producing the same results as those carried out with original images.

4.7.3 Hypermedia Applications

Hypermedia applications are retrieval services and require point-to-point or multipoint-to-point and switched services. They also require user interfaces, powerful authoring and presentation tools. Some applications are particularly suited for hypermedia, such as encyclopedias, dictionaries and other reference books [4.59]. They are composed of a number of independent units that are seldom accessed sequentially, but rather by selection and cross-reference with other entries. On the Internet, even for technical papers and reports that are considered more suitable for linear reading in sequence, there is an increasing tendency to include Hypertext Markup Language (HTML). The terms "hypertext" and "hypermedia" are usually distinguished. The former refers to a system with text-only information, and the latter refers to multimedia data.

A hypermedia system may be tracked as an application of database systems because it provides flexible access to multimedia information and a novel method to structure and manage data. A hypermedia system allows the user more freedom in assimilating and exploring information, as the conventional database has well-defined structures and manipulation languages for data processing [4.60].

Basic Features of a Hypermedia System

The basic features of a hypermedia system are the following [4.61]:

- Information may be divided into several smaller units or nodes. Each node may contain single media data, such as text, source code, graphics, video, audio or animation, or combinations of two or more media types. The contents in different nodes may have different authors.

- The links that interconnect the units of information can be bidirectional to facilitate backward traversals. Node contents are displayed by activating links.
- The system consists of a network of nodes and links, which may be distributed across multiple sites and remote servers in a computer network.
- Linear reading is used to move up and down among the document pages at a node. Nonsequential browsing allows users to jump back and forth in the information space and to navigate the hypermedia network.
- With authoring tools, users can build their own information structure for various purposes, through creation, manipulation and linkage of shared information units.
- It is necessary to maintain a database system to manage the shared multimedia data. In addition to the database functions, such as data-integrity guarantee, query processing, concurrency control, failure recovery and security mechanisms, support of rich modeling, document hierarchy and hypertext link should be included.

Due to proprietary storage mechanisms and document formats, most current hypermedia systems are not amenable to interoperability, and it is difficult to integrate the materials created in different systems. To make hypermedia systems completely open and interchangeable, various hypermedia models and document architectures have been developed [4.62, 4.63, 4.64].

The Web

A nice property of the Web is that it shields the implementation details of both the formats and the access protocols and presents a uniform interface to users. The Web is evolving toward a universal information hyperspace with a distributed, collaborative infrastructure. In practice, the Web is a vast collection of digitized information documents stored in computers connected to the worldwide Internet. It is accessible through Web browsers, such as Netscape and Microsoft Explorer, and runs on a request-and-response paradigm known as Hypertext Transport Protocol (HTTP) to allow online service to information located in remote sites. HTML is the standard document representation format for Web pages.

The Web documents created by HTML are static in nature. The author of a Web page determines its contents. After it is stored on the Web server, each browser requesting the document obtains the same response. In contrast, in dynamic documents, the contents are generated by application programs dynamically during runtime, depending on the current input from the user. Two approaches may be used to create such dynamic documents. The first is exemplified by the Common Gateway Interface (CGI), which is a widely used technology. As soon as a browser requests a document, the Web server invokes a CGI program that creates a document according to the input information. The server then forwards this document to the client browser. The other approach allows the dynamics of documents even after they have been loaded into the client browsers. A popular example of this approach is Java. Whenever a browser requests a document, the server sends a copy of a computer program that the browser must run locally. As a result, the display of the document can change continuously and can also interact with user inputs, alleviating the burden of the servers. For example, one may use this approach to display animation of one's home page. There is no central authority on the Web. Anyone can create a Web document and reference other documents. The Uniform Resource Locator (URL) is the specification scheme to locate a document unambiguously on the Web.

Example 4.6 In a URL such as *commsci.usc.edu/faculty/li.html, http* indicates the access method and *commsci.usc.edu* is the machine name. Other methods to retrieve information from the Web include File Transfer Protocol (FTP) and Telnet.

4.8 Concluding Remarks

A DMS is an integrated communication, computing, and information system that enables the processing, management, delivery and presentation of synchronized multimedia information. Such a system includes discrete media data (text, data and images) and continuous media data (video or audio). It enhances human communications by exploiting both visual and aural senses and provides the ultimate flexibility in work and entertainment by allowing one to collaborate with participants, to view video movies on demand, to access online digital libraries from the desktop, and so forth. Solutions for DMS issues have to be developed and have to provide a complete multimedia communication infrastructure, which is needed to support distributed multimedia applications. More user control and interactivity are desired. Faster processors and hardware, higher network bandwidth and data compression ratios, as well as improvements in a variety of related technologies, are necessary. Standardization is also important to accommodate the heterogeneity of techniques and to provide probability of applications. For the applications to be commercially viable, cost is an important consideration.