CHAPTER 1

Introduction

1.1 Evolution of ATM

With the advent of high-speed telecommunications technology, many new services began to emerge in the 1980's [1]—[3]. These generally consisted of voice, data, image, and video, and varied in bandwidth from 56/64 kb/s all the way to broadband rates of 155 Mb/s or higher. While much of the user traffic (such as voice and video) in these services was traditional in nature and required constant bandwidth for the entire duration of a call, there were many services where the traffic was bursty. In other words, the user remained inactive for quite some time, and then suddenly emitted bursts of data at a high speed for a short duration. Thus, even though the bandwidth, when averaged over a long time, did not vary much, the user may have required a much higher bandwidth for a short duration. In some cases, the applications were connection-oriented, where a connection had to be established before the user information exchange could take place. In this case, for a multimedia service, a single user might have required multiple channels simultaneously. In other instances, applications were connectionless, such as *email* and other LAN-type data communications.

As the high-speed technology began to develop, the CCITT Study Group XVIII undertook the task of defining a standard that would provide all these different services from a public or private network over a single, integrated interface. While information transfer across any interface may involve many functional layers, e.g., service modules, the transfer mode, the physical transmission, etc., Study Group XVIII concentrated on the transfer mode that included only switching and multiplexing — functions similar to those of a link-layer protocol.

In the early stages, the standards group considered the possibility of using a synchronous transfer mode (STM) [1]. Here, the transmitter sends out its data in periodic frames, each consisting of a number of time slots, and each slot being allocated to a particular application for the entire duration of a call. A generic frame format is depicted in Figure 1-1, where each frame is shown to be preceded by a distinct synchronizing pattern.[1] In STM, frames with

1. In some cases, as in T1, the synchronizing pattern may be distributed on a bit-by-bit basis over the full length of a frame.

Figure 1-1 A repetitive frame structure in STM.

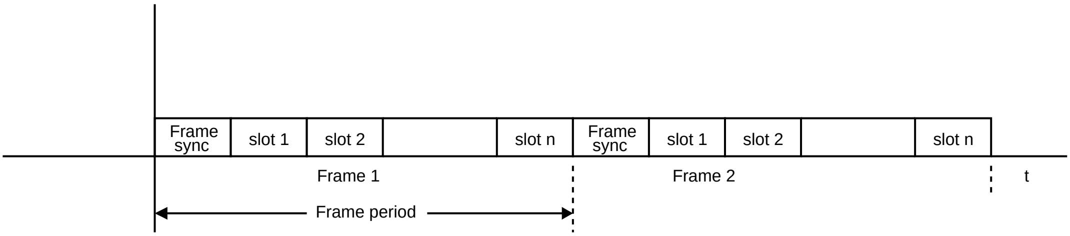

all their time slots appear periodically on a link, even if there is no information to send over a slot.

Narrow-band-ISDN (N-ISDN) is an example of STM where the frame structure is 2B + D for the Basic Rate ISDN (BRI) interface, 23B + D for the Primary Rate ISDN (PRI) interface for North American countries, and 30B + D for the PRI interface for European countries. Here, each B and D channel is 64 kb/s. In N-ISDN, while an application is usually allocated a single B channel, the PRI standard allows for higher bandwidths to be allocated to a user by defining new bandwidth structures such as $H_0 = 6B$, $H_{10} = 23B$, and $H_{11} = 30B$. Thus, at a PRI interface, in a general case, bandwidths are assigned according to the following expression:

$$n_1B + n_2H_0 + n_3H_{10}$$

where the numbers n_1, n_2, and n_3 are such that the available bandwidth of the PRI interface is never exceeded. Thus, for example, a user can be assigned up to four H_0 channels, a single H_{10} channel, or four B channels together with a single H_0 channel. In fact, Study Group XVIII initially considered a frame structure of this type, but one that would allow a much wider bandwidth and a more general and flexible allocation. For example, two new channel types were defined: H_2 with a bandwidth of 32.768 Mb/s, and H_4 with a bandwidth of 132.032 − 138.240 Mb/s. Bandwidth allocation schemes were considered whereby the variable bandwidth requirement of an application could be met by assigning an appropriate number of these channels or fragments thereof in some predetermined combinations. It was observed that with ATM, even if the requested bandwidths remained constant throughout the life of a call, there would be situations, depending on the user demands, where fractions of the bandwidth available on an interface would be unusable, thus resulting in inefficient utilization of the bandwidth. More importantly, with STM, it would be very difficult to meet the bandwidth requirement of bursty traffic by dynamically assigning a variable number of time slots.

In view of these limitations, a new transfer mode — known as the Asynchronous Transfer Mode (ATM) — was suggested. In this transfer mode, there is no concept of a periodic frame of the type illustrated in Figure 1-1. Instead, data is sent out in fixed-size blocks called cells. Whenever a user has information to send, it transmits as many cells as needed according to its bandwidth requirement at any instant, thus enabling an ATM switch to avoid the complexity that is inherent in dynamic allocation of time slots in an STM protocol. Each cell contains in its header portion a virtual channel identifier (VCI) that can be used as a label to identify a logical channel in much the same way as the position of a time slot indicates a specific channel in STM. Many of these logical channels may be multiplexed at an interface point, and as the information is transferred from one end-point to another, a given VCI may be mapped to different VCIs.

Even though the advantages of ATM were obvious, a number of issues were raised about its suitability. First, it was argued that delay-sensitive, constant-bit-rate services such as voice and high-bandwidth video could not be switched over ATM, and as such, might require STM with its circuit-switching capability to coexist with ATM in the same interface. Second, it was thought that even if ATM were suitable for all broadband services, hybrid switches consisting of both STM and ATM fabrics might still be needed in the early stages of ATM deployment to provide a transition from existing STM-based networks. Notwithstanding these problems,[2] there was general agreement among the participants in the standards organizations to proceed with the ATM approach. In 1989, CCITT Recommendation I.121 accepted ATM as the transfer mode solution for Broadband-ISDN (B-ISDN) [4]. Subsequently, a number of other CCITT recommendations were published on different aspects of the ATM protocol [5]—[8].

In 1991, an international organization called the ATM Forum was established to promote the understanding and development of the ATM technology. Since then, the Forum has developed a number of technical specifications on various aspects of the ATM protocol, and has produced white papers and computer-based training courses. For example, specifications have been developed on circuit and LAN emulation over ATM, the integrated layer management interface, physical, medium-dependent interfaces for ATM, etc.[3]

1.2 What is ATM?

The Asynchronous Transfer Mode is a packet switching and multiplexing technique. Even though the word "asynchronous" appears in its description, it *is not in any way an asynchronous transmission procedure.* Because of the way it has been designed, it is particularly suitable for high-bandwidth and low-delay applications. As indicated before, with ATM, information is sent out in fixed-size packets or cells, each containing in its header a VCI that provides a means for creating multiple logical channels and multiplexing them as needed. Because the cells have a fixed size, they may contain unused bits.

---

2. As we shall see later in this book, these problems were successfully tackled by CCITT Study Group XVIII, and subsequently by the ATM Forum. For example, the latter has standardized procedures for circuit emulation over ATM, which, in essence, provide constant bandwidths for delay-sensitive voice and video for the entire duration of a call. Similarly, procedures are being developed to support different applications using a single ATM switching fabric.

3. For a complete list of the approved ATM Forum Specifications, see the Technical Specifications section of the ATM Forum Web site.

Figure 1-2 The ATM protocol layers.

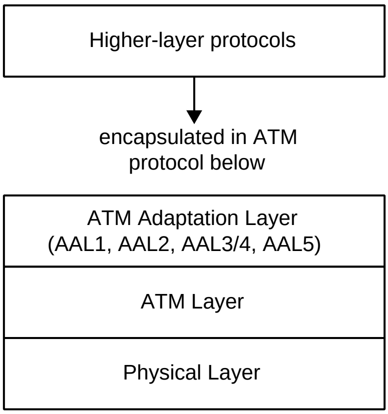

ATM is actually a very simple protocol: it merely transfers data from one point to another, and does not, by itself, provide any error recovery. However, ATM has been designed to interoperate with other, existing protocols. In fact, it can accommodate almost any upper-layer protocols that support end-to-end error recovery.

The complete ATM protocol stack is shown in Figure 1-2. The higher-layer protocols shown in the top-most box are application-specific. For example, they could be the standard file transfer protocol at the application layer with transmission control protocol (TCP) at the transport layer and Internet protocol (IP) at the network layer. Or, they could be simply the network-layer protocol for call and connection controls with a suitable application layer on top. As the name implies, the ATM adaptation layer (AAL) "adapts" the upper-layer packets to the ATM layer below. While the details might vary from one service to another (e.g., connectionless data services, connection-oriented data services, constant-bit-rate data services, variable-bit-rate data services, etc.), this adaptation is achieved by adding a header, a trailer, and some fill octets to the upper-layer packets and segmenting them into fixed-size ATM cells. Below this layer is the so-called ATM layer. This layer can be thought of as a link-layer protocol. However, in some respects, it is different from other link-layer protocols. For example, with the high-level data link control (HDLC) or Q.921 link access procedures on the D channel (LAPD) protocol, the length of a frame varies — it may vary from 2 octets to 256 octets. With the IEEE 802.3 protocol, the length of a packet, excluding the preamble and sync bits, may be anywhere from 64 octets to 1518 octets. In ATM, cell length is fixed at 53 octets. Also, unlike the link-layer protocols, ATM does not provide for acknowledgment procedures at the receiver. Errors can be detected in a cell, but not corrected. It is assumed that the transmission medium is highly reliable.

Blocks of user data of variable lengths from upper layers are passed to the ATM Adaptation Layer (AAL), which adds headers, trailers, padding octets, and/or cyclic redundancy

Figure 1-3 Label switching in an ATM network. Router R1 gets a packet with Label L1 in its header. It uses L1 to route the packet to R2, replacing L1 with Label L2.

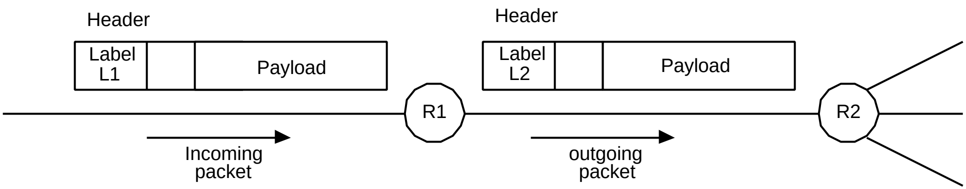

check (CRC) bits according to some rules that depend upon the service type. Each resulting data block is segmented into smaller blocks, which are then encapsulated into 53-octet cells in the ATM layer. It is these ATM cells that are transmitted to the destination. Many different media-dependent, physical-layer interfaces can be used for ATM.

1.3 Advantages of ATM

There are many advantages of an ATM network. Some of them are listed below.

- *Label switching*— The ATM protocol, like the frame relay protocol, is ideally suited for label switching. It works in the following way: In a traditional packet-switched network, when a packet arrives at a router, it examines its layer 3 header and routes the packet to the next hop along an appropriate route based upon the destination address. Since the network layer address generally contains much more information than would be required in making the routing decision, the layer 3 routing process is relatively complex. In label switching, the layer 3 address is mapped to a shorter identifier, which is called a label. It is important to emphasize here that a label is not an explicit address of an end-point. When the packet is routed to the next hop, the label is sent along with it as part of the header so that the router at the next hop can use it to derive subsequent routing information (see Figure 1-3). In ATM, labels are formed with the 24-bit virtual path identifier (VPI) and virtual channel identifier (VCI) fields.

  Label switching offers a number of advantages. First, since packets can now be routed using a label as an index into the switch memory to determine the next hop, and since labels are shorter than IP addresses, it is easier to build a label-switching router. Second, if IP packets are to be routed between any two end-points in an ATM network in the traditional way, either virtual circuits must be connected in a full mesh configuration among ATM switches, or cut-through switched virtual channels (SVCs) must be established using an appropriate protocol.[4] Label switching obviates the need for such mesh connections and reduces the number of peer routers that need to communicate

4. For example, the Next Hop Resolution Protocol (NHRP).

Figure 1-4 Illustration of service delays in a packet-switched network. (a) Fixed cell packets, as in ATM. (b) Variable-size packets, as in Ethernets.

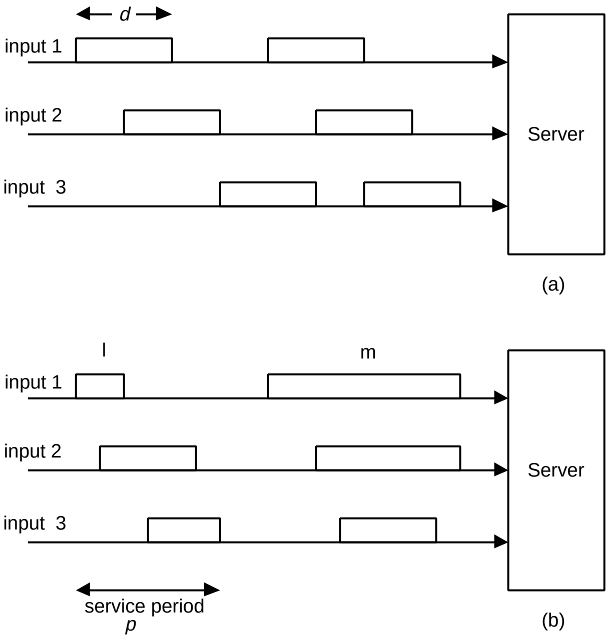

with each other. Consequently, a label-switching network is cheaper and faster. Third, label switching in an ATM environment is similar in many respects to label switching in other protocols [9], [10]. It may, therefore, be possible to use common methods for packet forwarding and even network management.

- *Low latency*—An important feature of the ATM protocol is its low latency and seamless capacity to span LANs and wide area networks (WANs). ATM's low latency results from the fact that all packets in the ATM layer have a fixed length. To see this, consider Figure 1-4(a), which shows a server with three inputs. Assume that all packets have a fixed size, *d.* The packets on any input link may arrive randomly with respect to the other input lines. If the server scans the inputs every *d* seconds, which is the length of a packet, then the average service delay for any packet on any input is *d/2* seconds. Perfect scheduling is possible here because the packet size is known *a priori.* For example, if the data rate on each input line is 25 Mb/s, then for 53-octet ATM cells, the service period is 16.96 microseconds, and the average delay is 8.48 microseconds.

  Next, consider the case where the size of a packet varies from some minimum value, say *L,* to a maximum value, *M.* This is shown in Figure 1-4(b) and is applicable to Ethernet. In this case, the server must scan the inputs frequently enough to match the length of the shortest packets; otherwise, these packets will be subjected to long delays. For example, if the packet size varies from 2 octets to 100 octets, the server should scan the inputs every 640 ns, which may be an excessive burden on the CPU. If the service period is increased to 16.96 microseconds, the average delay is about 13 times the size of the shortest packet. If all packets are fixed-length, scheduling of network resources is much easier.

- *High-speed and high-bandwidth*— Because of its low latency, ATM is particularly suitable for applications that require high-speed transport and high bandwidths. For example, one can use ATM in a network backbone that interconnects traditional LANs such as Ethernets and token ring LANs. Currently, there are many high-speed LANs such as gigabit Ethernets and metropolitan area networks (MANs) covering tens of kilometers in diameter using such protocols as fiber-distributed data interface (FDDI) and dual queue distributed bus (DQDB). ATM could very well form the basis of the new generation of high-speed LANs and MANs. Furthermore, ATM is equally suitable for applications that do not require high speed or high capacity. For example, currently, vendors are offering ATM switches at 25.6, 44.736, 51.84, 155, and 622 Mb/s bandwidths with various cabling types — two- or four-wire category 3 unshielded twisted pair (UTP), multi-mode and single-mode fibers, DS3, and T1/E1 copper circuits.
- *Integrated network*— Normally, a packet protocol is only suitable for bursty, variable-bit-rate services, and would not be able to transport information that is sensitive to delays. For example, the public-switched telephone network (PSTN) can only transport circuit-switched information. An X.25 or frame relay network can handle only packet-switched data. The ATM protocol has been designed such that it can carry not only bursty, variable-bit-rate services, but also delay-sensitive information such as voice and video that would normally be carried by circuit-switched networks. In fact, over the last few years, a rich set of procedures and protocols that enable ATM to support many different services have been developed by the ATM Forum — connection-oriented as well as connectionless, constant-bit-rate as well as variable, and they provide different qualities of service according to the application and customer needs. Thus, with ATM, it is possible to provide all different services with a single, integrated network.
- *Integrated access from customer premises*— ATM provides a means for achieving integrated access to broadband services from public or private networks. Services such as compact disc (CD)-quality music, pay-per-view movie channels, high-definition TV, high-speed Internet data downloading, etc. can all be combined with traditional, circuit-switched voice and low-speed data services and then presented over a single ATM pipe to the customer premises, where a set-top box would demultiplex these services. The user could even request special services from a network provider using an upstream control channel.
- *Interworks with existing protocols and legacy LANs*— There are many instances where new applications would almost certainly require the bandwidth and speed of an ATM network. One such example involves collaboration among different research organiza-

tions with high-resolution, high-bandwidth imaging data. The ATM network, if installed, would still be able to interwork with traditional data networking protocols and legacy LANs such as Ethernet, token ring, and FDDI. Thus, the existing network infrastructure needs to be augmented only when or where necessary, leading to a graceful but less expensive evolution to the new technology.

- *Bandwidth-on-demand*—Bandwidth-on-demand is another innate benefit of ATM. In private networks, higher bandwidths can be requested by users. However, generally, they must be provisioned through network managers, and cannot be assigned dynamically at connection setup time. With some private networks that are equipped with inverse multiplexing capability at both ends, it may be possible to request and obtain increased bandwidth dynamically. Even then, the range is rather limited. With ATM, users may request a desired bandwidth when originating a call, and the network would attempt to dynamically allocate the requested bandwidth only if the customer had subscribed to this feature at subscription time. Furthermore, for ATM networks, there are traffic and congestion control mechanisms in place which, in the event of congestion, allow the network to maintain the quality of service for each customer with minimum degradation. Thus, initially, one could install a network with only a minimum amount of reserve capacity, and add to the network only when the demand for bandwidth has grown to a point where it is no longer possible to provide each customer with the subscribed quality of service.

1.4 Some Examples of ATM Applications

Many different applications of ATM will be discussed in relevant chapters of this book. For the time being, however, we will present a few examples of common applications.

1.4.1 High-Bandwidth ATM Backbone

Because of its low latency, ATM is inherently suitable for a high-bandwidth backbone network. This is shown in Figure 1-5. The legacy LANs are connected to the ATM network through bridges and routers. Integrated data, voice, and video can now be delivered to desktops. This is possible because ATM is capable of supporting diverse, upper-layer protocols so that the existing network infrastructure can be connected to an ATM network without any modification. ATM servers and high-speed ATM workstations can also be directly connected.

1.4.2 ATM Switch in a Central Office

In the post-divestiture era, a local exchange carrier must provide interfaces between its local switches and any inter-exchange carrier that wants to provide inter-exchange service between two Local Access Transport Areas (LATAs) or between two points of the same LATA separated by a second LATA. Currently, a special type of switch, called an access tandem, is used for this purpose. It takes the inter-exchange traffic from the local switches in a central office, multiplexes them, and routes them to each destination inter-exchange carrier. Because these

Figure 1-5 ATM is used in a backbone network, inter-connecting different types of traditional LANs. Conversion to/from the ATM takes place at the bridge and router as shown.

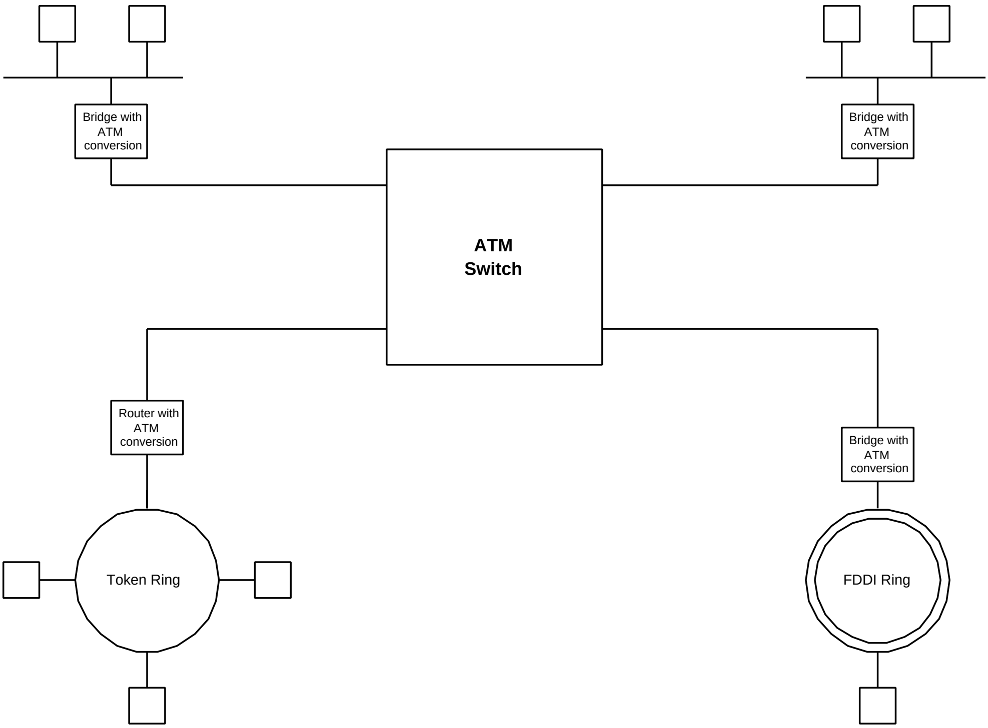

tandem switches have a fixed capacity, as the traffic grows with time, they must either have large spare capacity to start with so that they meet the growing demand with time, or they must be replaced by larger systems when the traffic exceeds their capacity. ATM switches, because of their efficient bandwidth management and inherently low delays, are ideal for this application. This is shown in Figure 1-6. The ATM network can similarly handle inter-exchange signaling information, data for operations and maintenance, exchange of data to or from advanced intelligent networks, etc. Thus, in essence, ATM provides a multi-service platform for public networks. The system has the following features:

- AAL Type 1 or 2 — constant-bit-rate, connection-oriented, with timing relation between source and destination required.
- Dynamic routing based on real-time routing criteria.
- Low latency.
- High bandwidth.

Figure 1-6 An ATM switch in a central office provides interfaces to inter-exchange carriers, acting as a high-speed tandem switch of an inter-exchange carrier at the point of connection with a local exchange carrier.

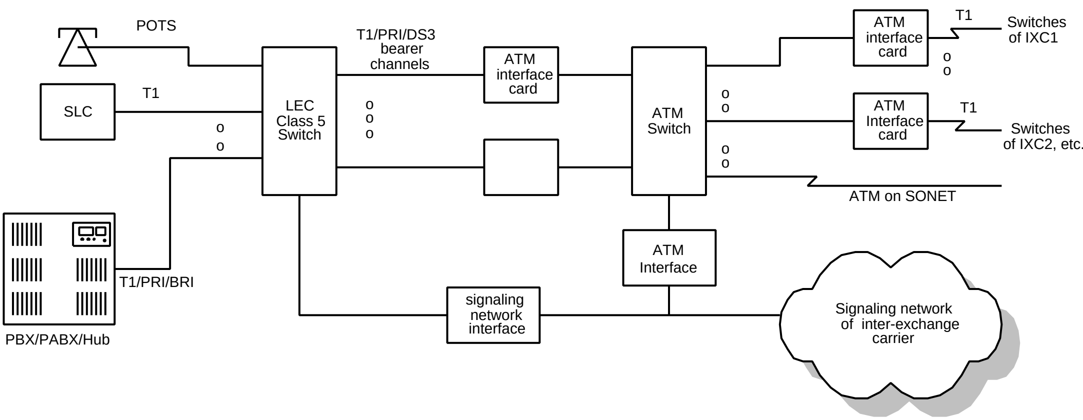

Figure 1-7 A simplified functional block diagram of a mobile communications network.

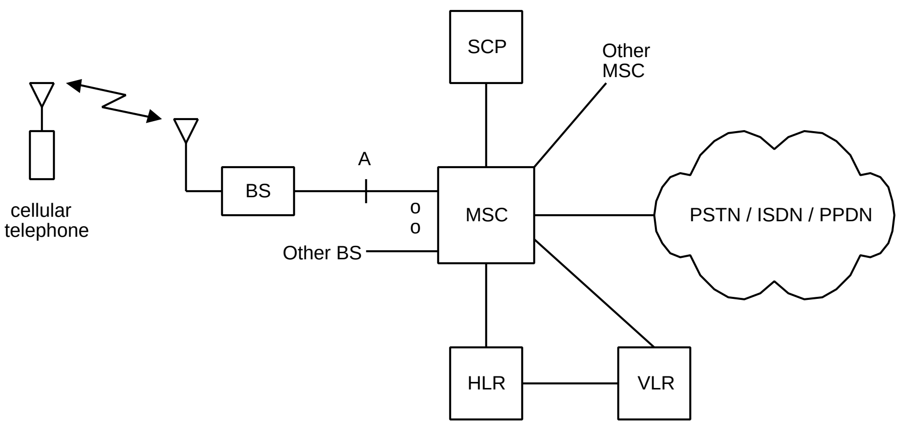

1.4.3 ATM in Mobile Communications Systems

A mobile communications system is another interesting application of the ATM technology. ATM is currently being used in TIA/EIA standard IS-634-A, which defines an interface between a base station (BS) and a mobile switching center (MSC) [13]. A simplified version of the network reference model for this interface is shown in Figure 1-7. Base stations, which provide radio communications with mobile stations, form a wireless network. Calls to or from a mobile station are switched by an MSC via the base stations. An MSC is usually connected to a PSTN, an ISDN, a Public Packet Data Network (PPDN), or any combination thereof, and may also be connected to other MSCs as well. The home location register (HLR) contains a database of the phone numbers and services for all subscribers to the home system. The visitor location register (VLR) contains a database of the phone numbers and services for all visitors to the system. The entity SCP (service control point) is a database system that performs the so-called intelligent network (IN) functions such as translation of the destination number for 800-type calls, credit card validation, voice mail and voice recognition systems, etc.

IS-634-A defines the interface at reference point A between an MSC and a BS. While it specifies the interface requirements for all types of user traffic and signaling information exchanged over this reference point, ATM is used to transport only the following information:

1. The coded user traffic (e.g., user data, or 64-kb/s PCM voice packetized to lower bit rates) and signaling information between an MSC and a BS. The purpose of the signaling information is to allocate the radio channels that transport the user traffic. Separate logical channels carry the user traffic and signaling information. These interface functions are designated as the A3 interface.
2. The signaling information between the source BS that initially serves a call and any other BS that supports the call (i.e., the target BS). This interface function is designated as the A7 interface.

Figure 1-8 shows the protocol stack for these interfaces. Notice that at the ATM adaptation layer, AAL5 is used for the signaling information and AAL2 for the user traffic.

The use of ATM for the next generation of mobile communications systems has been proposed by a number of authors [11], [12]. A possible architecture for such a network that would be able to support different mobile communications systems, e.g., AMPS, CDMA,

Figure 1-8 The protocol stack for the A3 and A7 interfaces.

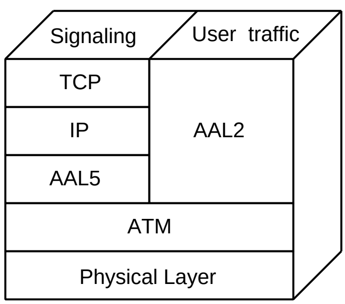

Figure 1-9 Architecture of the next-generation, mobile communications network.

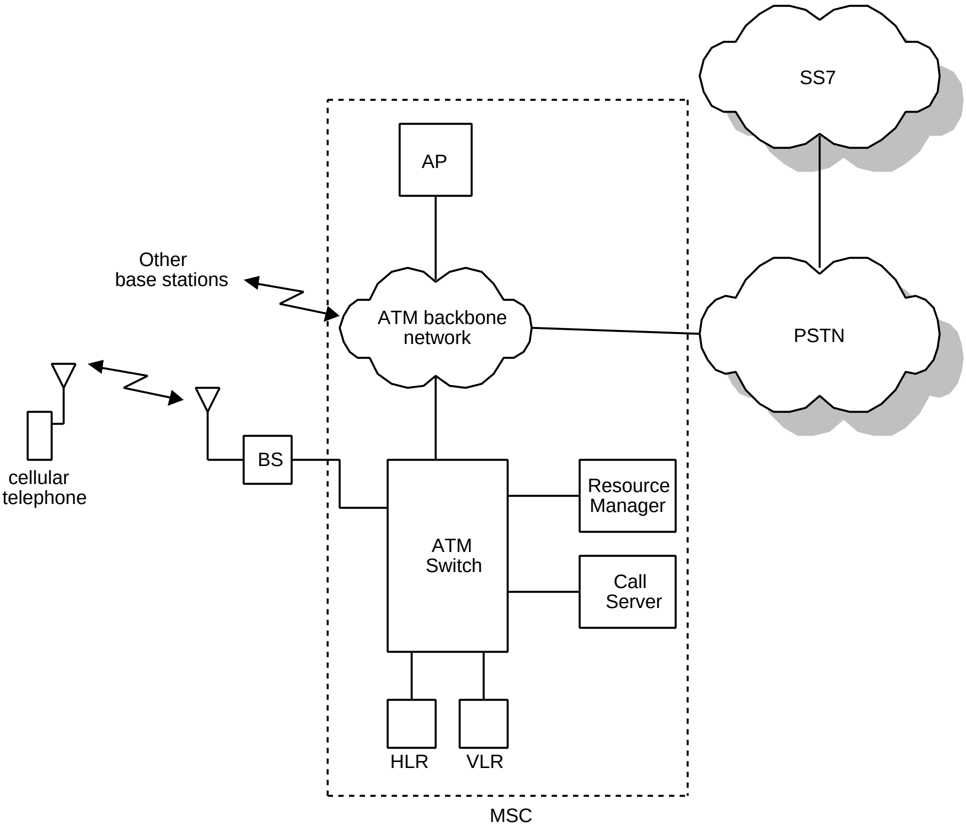

TDMA, and GSM,[5] and at the same time provide interoperability among them, is shown in Figure 1-9. In this architecture, ATM is the underlying technology of the interconnection network. It supports multimedia services using both the native-mode ATM and existing TCP/IP protocols according to service needs, and it provides bandwidth-on-demand with guaranteed quality of service. Each BS contains multiple radios and operates over either a subset of channels available to a system, as in AMPS, or all channels, as in the CDMA system. The block designated as the Resource Manager manages the radios and channels via the ATM switch. The Call Server implements call-associated functions and interacts again through the ATM switch with the Resource Manager whenever a mobile has to be connected to a radio, disconnected from a radio, or reassigned to a new channel.[6]

---

5. AMPS — Advanced Mobile Phone Service; CDMA — Code Division Multiple Access; TDMA — Time Division Multiple Access; GSM — Global System for Mobile Communications.

6. As a mobile moves from one sector of a cell to another sector, or from one cell to another, it is necessary to switch the mobile from one channel to another. This is called a hard handoff. In a CDMA

Figure 1-10 Video conferencing in a private network using an ATM switch. ATM multimedia conferencing is based on ITU standards H.310 and H.321. Copyright 1995 ATM Forum.

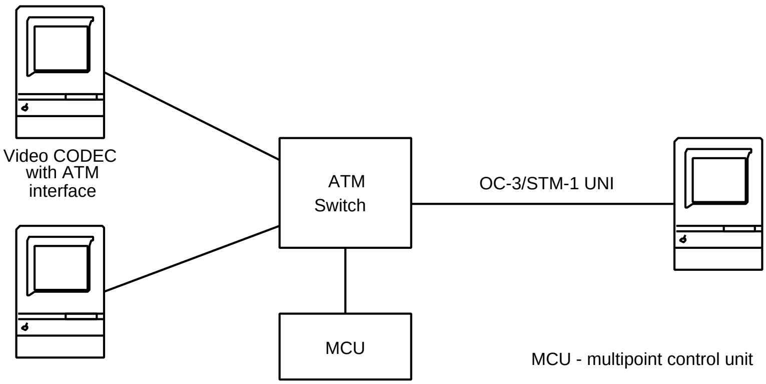

The ATM backbone network permits multiple base stations in a mobile service area to be connected together, provides interoperability among various access networks (e.g., AMPS, CDMA, TDMA, and GSM), and connects the mobile serving area to the existing PSTN. New services and features are added to the network via the entity marked AP (adjunct processor) on a per-need basis. Since the APs are connected to the ATM backbone network, the same services are uniformly available throughout the system at increased efficiency and reduced cost. The network marked SS7 (Signaling System 7) forms the basis of the advanced intelligent networks (AINs), and provides such services as 800-type calls, credit card validation, home location register for mobile communications systems, voice mail and voice recognition systems, etc.

1.4.4 Video Conferencing over ATM

Because of its low latency and high-bandwidth capability, ATM is ideally suited for video conferencing. In fact, ATM switches that provide full-motion, high-quality video conferencing over OC-3/STM-1 user network interfaces are available in the market. They support the use of a multipoint control unit (MCU) to control both conferences and broadcasts of audio and video signals. A generic network configuration for video conferencing is shown in Figure 1-10. Similarly, some commercial switches provide multimedia services, integrating high-resolution, bi-directional video, stereo-quality audio, and high-speed (48 kb/s) data.

---

system, there is another type of handoff that does not involve switching the channels; as a mobile moves from one location to another, a new BS may begin to serve the mobile using the same frequency band as the old BS. This is known as a soft handoff. During the life of a call, a mobile may go through a number of handoffs.

1.4.5 Mass Distribution of Multimedia Information in Real Time

Because of its high bandwidth and low latency, ATM is ideal for distributing high-quality multimedia programs in real time. Figure 1-11 shows how a cable TV provider or news service organization can deliver multimedia services in real time to its customers. The analog signal is converted into digital form with an MPEG-2 (Motion Pictures Expert Group) encoder. Higher layer protocols such as ITU Recommendations H.310 or H.321 are used to convert the encoder output into a form suitable for transmission over ATM. This ATM signal is transported over an OC-3 link to the internet, or one or more cable TV networks where it is first converted into an analog form so that it can be broadcast over a network. The encoded signal is stored on a server so that it can be broadcast at a later time, or locally reviewed and edited. This system has the following advantages:

- Because ATM can support different qualities of service according to customer needs, the system can serve a much larger number of users.
- Since the system is capable of large bandwidths, it can provide multiple programs of high quality videos simultaneously.

Figure 1-11 The use of ATM to provide multimedia services over the Internet or cable TV networks.

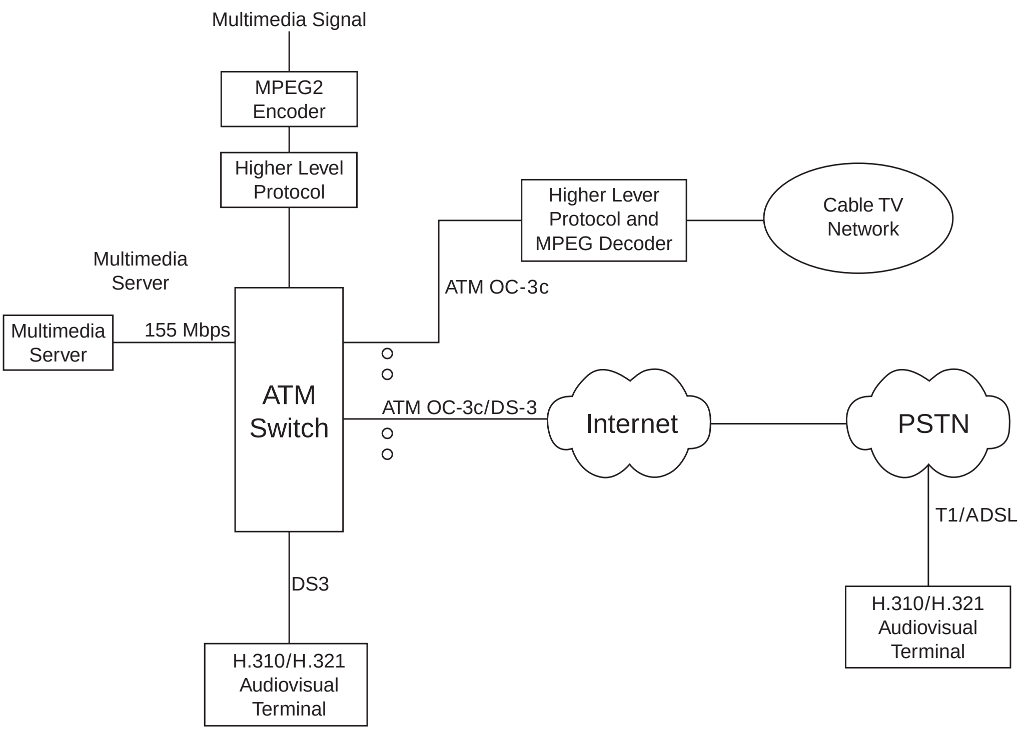

- Since delays caused by ATM switches are inherently low, events can be broadcast in real time using digital techniques for high quality videos.

A number of other applications for ATM are described in Reference [14].

1.5 Summary

In this chapter, we traced the evolution of ATM and presented a brief description of the protocol. ATM offers a number of advantages. For example, it is capable of label switching, provides low latency, and supports not only traditional, variable-bit-rate data services but also delay-sensitive voice and video. As such, it can be used to provide integrated access from customer premises to a broadband network. A number of supporting protocols have been developed for ATM that allow it to interwork with traditional data networking protocols and legacy LANs. Thus, the existing network infrastructure requires augmentation only when necessary, leading to a graceful and less expensive evolution of the ATM technology. Another important aspect of ATM is the provision of bandwidth-on-demand and quality of service. Congestion control mechanisms have been developed that guarantee the quality of service requested by each user. We discussed these advantages in detail and closed the chapter with a number of examples that illustrate the benefits of ATM.

1.6 References

[1] S. E. Minzer, "Broadband ISDN and Asynchronous Transfer Mode (ATM)," *IEEE Commun. Mag.*, Sept. 1989, pp. 17–24.

[2] J. S. Turner, "New Directions in Communications," *IEEE Commun. Mag.*, vol. 24, no. 10, Oct. 1986, pp. 8–15.

[3] G. Hayward, L. Linnell, D. Mahoney, and L. Smoot, "A Broadband ISDN local access system using emerging-technology components," *Proc. of Int'l Switching Symp. '87,* vol. 3, 1987, pp. 597–601.

[4] CCITT Recommendation I.121, "Broadband aspects of ISDN," Blue Book, Geneva, Switzerland, 1989.

[5] ITU-T Recommendation I.150, "B-ISDN Asynchronous Transfer Mode Functional Characteristics," Nov. 1995.

[6] ITU-T Recommendation I.361, "B-ISDN ATM Layer Specification," Nov. 1995.

[7] ITU-T Recommendation I.363, "B-ISDN ATM Adaptation Layer (AAL) Specifications," Mar. 1993.

[8] ITU-T Recommendation I.413, "B-ISDN User-Network Interface," Mar. 1993.

[9] R. Callon, et al., "A Framework for Multiprotocol Label Switching," Internet Engineering Task Force, Network Working Group, *Internet Draft*, Nov. 1997.

[10] E. Rosen, et al., "Multiprotocol Label Switching Architecture," Internet Engineering Task Force, Network Working Group, *Internet Draft*, July 1998.

[11] G. E. Fry, et al., "Next generation wireless networks," *Bell Labs Tech. J.*, Vol. 1, No. 2, Autumn 1996, pp. 88–96.

[12] E. Ayanoglu, et al., "Mobile information infrastructure," *Bell Labs Tech. J.*, Vol. 1, No. 2, Autumn 1996, pp. 143–163.
[13] TIA/EIA/IS-634-A MSC-BS Interface (A-Interface), July 1998.
[14] ATM Forum: ATM Application Notes, NetWorld+Interop 97, Las Vegas, May 6–8, 1997.