Module 4

Switched Communication Networks

Lesson 6

Asynchronous Transfer Mode Switching (ATM)

Specific Instructional Objectives

On completion on this lesson, the student will be able to:

- State the need for ATM
- Explain the concept of cell switching
- Specify the architecture of ATM
- Explain the operation of Virtual connections and switching types used
- Explain switching fabric of ATM
- Explain the functions of the three ATM layers

4.6.1 Introduction

Asynchronous Transfer Mode (ATM) is an International Telecommunication Union-Telecommunications Standards Section (ITU-T) standard for cell relay wherein information for multiple service types, such as voice, video, or data, is conveyed in small, fixed-size cells. ATM networks are connection-oriented. Asynchronous transfer mode (ATM) is a technology that has its history in the development of broadband ISDN in the 1970s and 1980s. Technically, it can be viewed as an evolution of packet switching. Like packet switching protocols for data (e.g., X.25, frame relay, Transmission Control Protocol and Internet protocol (TCP IP]), ATM integrates the multiplexing and switching functions, is well suited for bursty traffic (in contrast to circuit switching), and allows communications between devices that operate at different speeds. Unlike packet switching, ATM is designed for high-performance multimedia networking. ATM technology has been implemented in a very broad range of networking devices. The most basic service building block is the ATM virtual circuit, which is an end-to-end connection that has defined end points and routes but does not have bandwidth dedicated to it. Bandwidth is allocated on demand by the network as users have traffic to transmit. ATM also defines various classes of service to meet a broad range of application needs. This lesson provides an overview of ATM protocols, services, and operation.

4.6.2 Benefits of ATM

The high-level benefits delivered through ATM services deployed on ATM technology using international ATM standards can be summarized as follows:

- **Dynamic bandwidth for bursty traffic** meeting application needs and delivering high utilization of networking resources; most applications are or can be viewed as inherently bursty, for example voice is bursty, as both parties are neither speaking at once nor all the time; video is bursty, as the amount of motion and required resolution varies over time.
- **Smaller header** with respect to the data to make the efficient use of bandwidth.
- **Can handle Mixed network traffic very efficiently:** Variety of packet sizes makes traffic unpredictable. All network equipments should incorporate elaborate software systems to manage the various sizes of packets. ATM handles these problems efficiently with the fixed size cell.
- **Cell network:** All data is loaded into identical cells that can be transmitted with complete predictability and uniformity.

- **Class-of-service support** for multimedia traffic allowing applications with varying throughput and latency requirements to be met on a single network.
- **Scalability in speed** and network size supporting link speeds of T1/E1 to OC–12 (622 Mbps).
- **Common LAN/WAN architecture** allowing ATM to be used consistently from one desktop to another; traditionally, LAN and WAN technologies have been very different, with implications for performance and interoperability. But ATM technology can be used either as a LAN technology or a WAN technology.
- **International standards compliance** in central-office and customer-premises environments allowing for multivendor operation.

4.6.3 ATM Devices and the Network Environment

ATM is a cell-switching and multiplexing technology that combines the benefits of circuit switching (guaranteed capacity and constant transmission delay) with those of packet switching (flexibility and efficiency for intermittent traffic). It provides scalable bandwidth from a few megabits per second (Mbps) to many gigabits per second (Gbps). Because of its asynchronous nature, ATM is more efficient than synchronous technologies, such as *time-division multiplexing (TDM).*

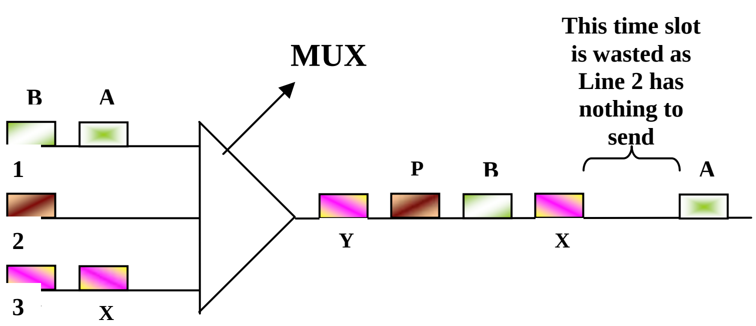

Figure 4.6.1 Normal TDM operation

With TDM, each user is assigned to a time slot, and no other station can send in that time slot as shown in Fig. 4.6.1. If a station has much data to send, it can send only when its time slot comes up, even if all other time slots are empty. However, if a station has nothing to transmit when its time slot comes up, the time slot is sent empty and is wasted.

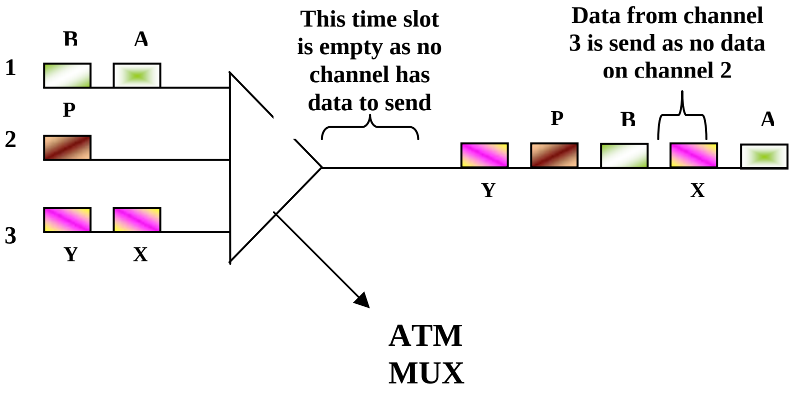

Figure 4.6.2 Asynchronous multiplexing of ATM

Because ATM is asynchronous, time slots are available on demand with information identifying the source of the transmission contained in the header of each ATM cell. Figure 4.6.2 shows how cells from 3 inputs have been multiplexed. At the first clock tick input 2 has no data to send, so multiplexer fills the slot with the cell from third input. When all cells from input channel are multiplexed then output slot are empty.

ATM Devices

An *ATM network* is made up of an *ATM switch* and *ATM endpoints*. An ATM switch is responsible for cell transit through an ATM network. The job of an ATM switch is well defined. It accepts the incoming cell from an ATM endpoint or another ATM switch. It then reads and updates the cell header information and quickly switches the cell to an output interface towards its destination. An ATM endpoint (or end system) contains an ATM network interface adapter. Examples of ATM endpoints are workstations, routers, digital service units (DSUs), LAN switches, and video coder-decoders (Codec's).

ATM Network Interfaces

An ATM network consists of a set of ATM switches interconnected by point-to-point ATM links or interfaces. ATM switches support two primary types of interfaces: UNI and NNI as shown in Fig. 4.6.3. The UNI (User-Network Interface) connects ATM end systems (such as hosts and routers) to an ATM switch. The NNI (Network-Network Interface) connects two ATM switches. Depending on whether the switch is owned and located at the customer's premises or is publicly owned and operated by the telephone company, UNI and NNI can be further subdivided into public and private UNIs and NNIs. A private UNI connects an ATM endpoint and a private ATM switch. Its public counterpart connects an ATM endpoint or private switch to a public switch. A private NNI connects two ATM switches within the same private organization. A public one connects two ATM switches within the same public organization.

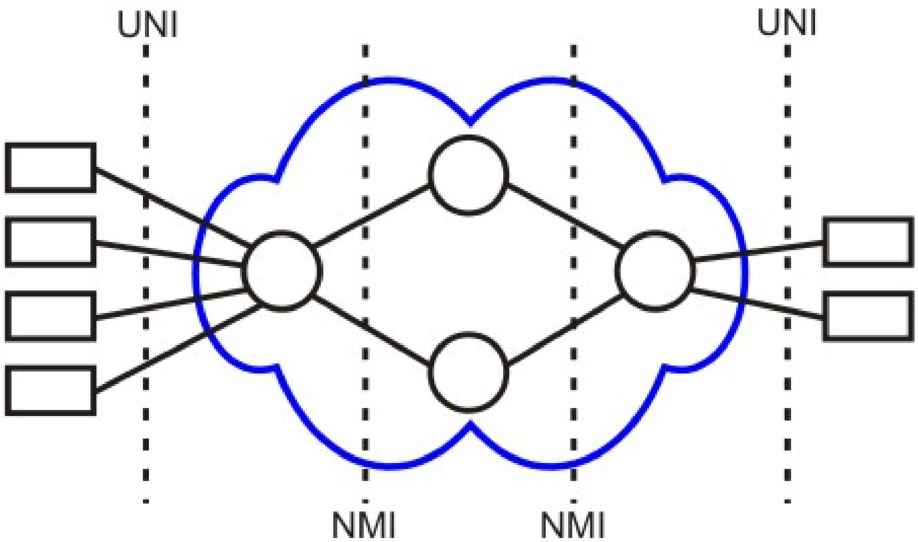

Figure 4.6.3 UNI and NNI interfaces of the ATM

4.6.4 ATM Cell Format

ATM transfers information in fixed-size units called *cells*. Each cell consists of 53 octets, or bytes as shown in Fig. 4.6.4. The first 5 bytes contain cell-header information, and the remaining 48 contain the payload (user information). Small, fixed-length cells are well suited to transfer voice and video traffic because such traffic is intolerant to delays that result from having to wait for a large data packet to download, among other things.

Header **5 bytes**	**Payload** **48 bytes**

Figure 4.6.4 ATM cell Format

An ATM cell header can be one of two formats: UNI or NNI. The UNI header is used for communication between ATM endpoints and ATM switches in private ATM networks. The NNI header is used for communication between ATM switches. Figure 4.6.5 depicts the ATM UNI cell header format, and the ATM NNI cell header format. Unlike the UNI, the NNI header does not include the Generic Flow Control (GFC) field. Additionally, the NNI header has a Virtual Path Identifier (VPI) field that occupies the first 12 bits, allowing for larger trunks between public ATM switches.

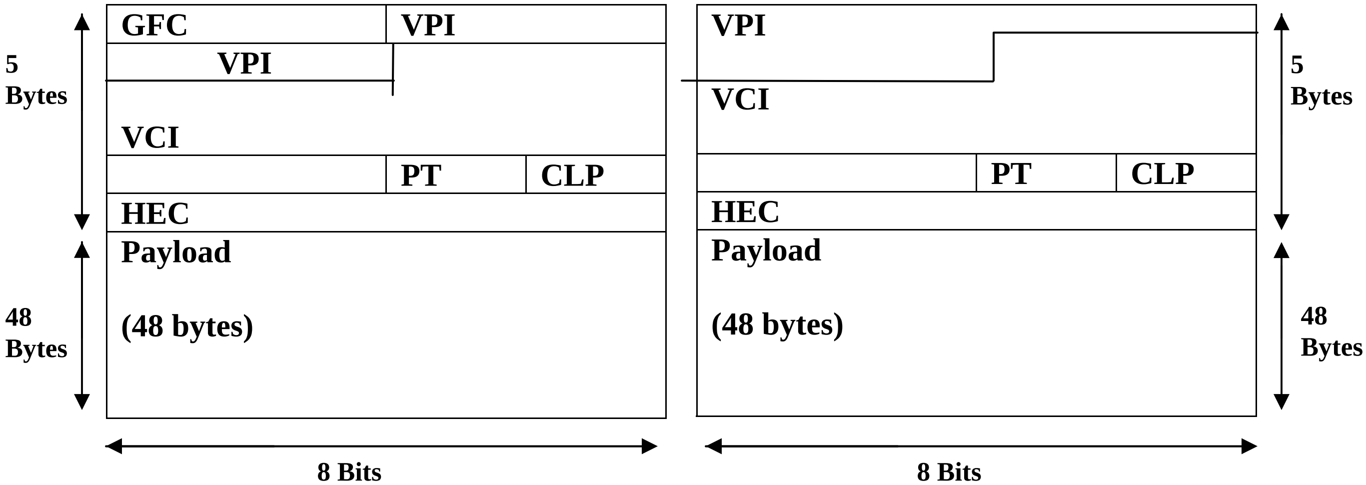

Figure 4.6.5 (a) UNI Cell Format

Figure 4.6.5 (b) NNI Cell Format

ATM Cell Header Fields

The following descriptions summarize the ATM cell header fields shown in Fig. 4.6.5.

- **Generic Flow Control (GFC)—**Provides local functions, such as identifying multiple stations that share a single ATM interface. This field is typically not used and is set to its default value of 0 (binary 0000).
- **Virtual Path Identifier (VPI)—**In conjunction with the VCI, identifies the next destination of a cell as it passes through a series of ATM switches on the way to its destination.
- **Virtual Channel Identifier (VCI)—**In conjunction with the VPI, identifies the next destination of a cell as it passes through a series of ATM switches on the way to its destination.
- **Payload Type (PT)—**Indicates in the first bit whether the cell contains user data or control data. If the cell contains user data, the bit is set to 0. If it contains control data, it is set to 1. The second bit indicates congestion (0 = no congestion, 1 = congestion), and the third bit indicates whether the cell is the last in a series of cells that represent a single AAL5 frame (1 = last cell for the frame).
- **Cell Loss Priority (CLP)—**Indicates whether the cell should be discarded if it encounters extreme congestion as it moves through the network. If the CLP bit equals 1, the cell should be discarded in preference to cells with the CLP bit equal to 0.
- **Header Error Control (HEC)—**Calculates checksum only on the first 4 bytes of the header. HEC can correct a single bit error in these bytes, thereby preserving the cell rather than discarding it.

4.6.5 ATM Virtual Connections

ATM standard defines two types of ATM connections: virtual path connections (VPCs), which contain virtual channel connections (VCCs) as shown in Fig. 4.6.6. A virtual channel connection (or virtual circuit) is the basic unit, which carries a single stream of cells, in order, from user to user. A collection of virtual circuits can be bundled together into a virtual path connection. A virtual path connection can be created from end-to-end across an ATM network. In this case, the ATM network does not route cells belonging to a particular virtual circuit. All cells belonging to a particular virtual path are routed the same way through the ATM network, thus resulting in faster recovery in case of major failures. In this case, all the switches within the ATM network are only VP switches, i.e. they switch the cells only on the basis of VPIs. Only the switches, which are connected to the subscribers are VP/VC switches, i.e. they use both VPIs and VCIs to switch the cell. This configuration is usually followed so that the intermediate switches can do switching much faster.

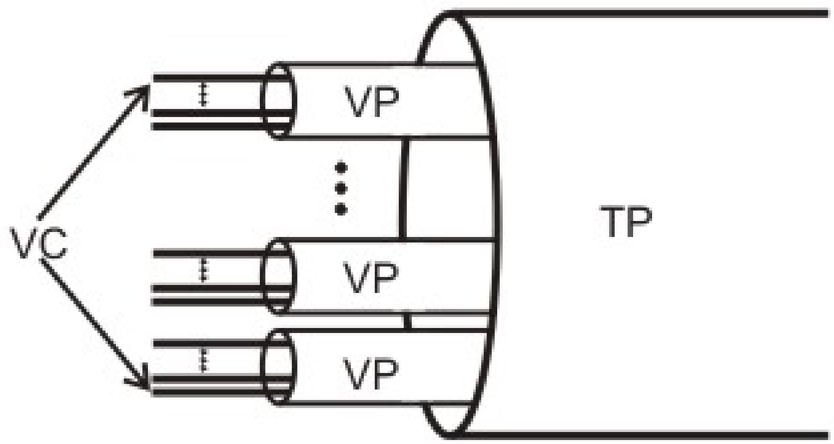

Figure 4.6.6 Virtual channel connections of ATM

An ATM network also uses virtual paths internally for the purpose of bundling virtual circuits together between switches. Two ATM switches may have many different virtual channel connections between them, belonging to different users. These can be bundled by two ATM switches into a virtual path connection. This can serve the purpose of a virtual trunk between the two switches. This virtual trunk can then be handled as a single entity by perhaps, multiple intermediate virtual paths cross connects between the two virtual circuit switches.

ATM Switching Operations

The basic operation of an ATM switch is straightforward: The cell is received across a link with a known VPI/VCI value. The switch looks up the connection value in a local translation table to determine the outgoing port (or ports) of the connection and the new VPI/VCI value of the connection on that link. The switch then retransmits the cell on that outgoing link with the appropriate connection identifier.

Incoming		Outgoing		
VPI	VCI	VPI	VCI	Interface
10	122	11	41	1
121	213	10	158	1
12	11	211	111	2
11	151	321	210	2

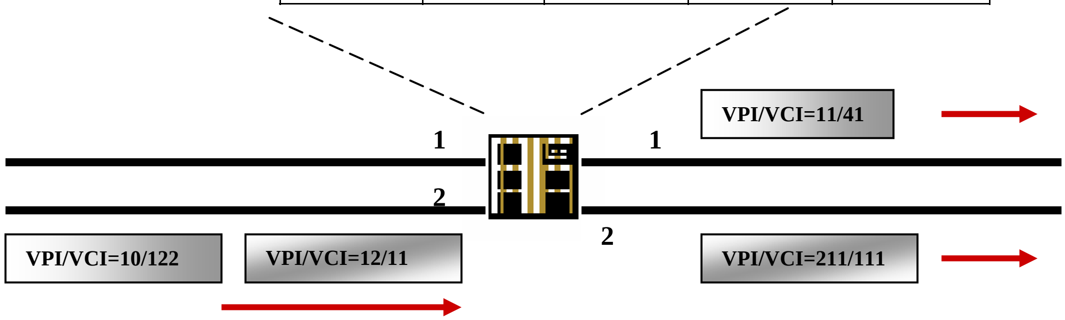

Figure 4.6.7 A VP/VC ATM switch table

Because all VCIs and VPIs have only local significance across a particular link, these values are remapped, as necessary, at each switch. Figure 4.6.7 and Fig. 4.6.8 shows a VP-VC switch and an only VP switch, respectively. Usually the intermediate switches are only VPI switches while switches connected to the users are VPI/VCI switches.

Incoming	Outgoing	
VPI	VPI	Interface
22	65	1
121	99	2
312	201	1
11	21	2

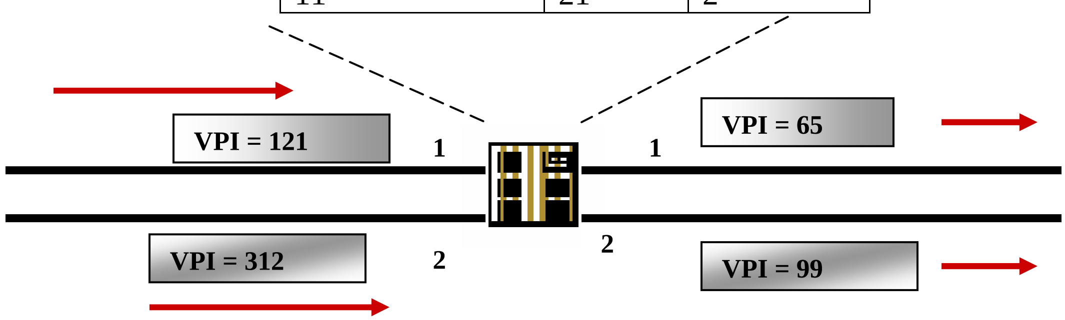

Figure 4.6.8 VP ATM switch table

To make the switching more efficient, ATM uses two types of switches namely, VP switch and VP-VC switch. A VP switch route cells only on the basis of VPI, here VPIs change but VCIs remain same during switching. On the other hand, VP-VC switch uses the complete identifier, i.e. both VPI and VCI to route the cell. We can think of a VP-VC switch as a combination of Only VP and Only VC switch.

4.6.6 ATM Reference Model

The ATM architecture uses a logical model to describe the functionality that it supports. ATM functionality corresponds to the physical layer and part of the data link layer of the OSI reference model.

The ATM reference model, as shown in Fig. 4.6.9, consists of the following planes, which span all layers:

- **Control**—This plane is responsible for generating and managing signaling requests.
- **User**—This plane is responsible for managing the transfer of data.
- **Management**—This plane contains two components:
    - Layer management manages layer-specific functions, such as the detection of failures and protocol problems.
    - Plane management manages and coordinates functions related to the complete system.

The ATM reference model consists of the following ATM layers:

- **Physical layer**—Analogous to the physical layer of the OSI reference model, the ATM physical layer manages the medium-dependent transmission.

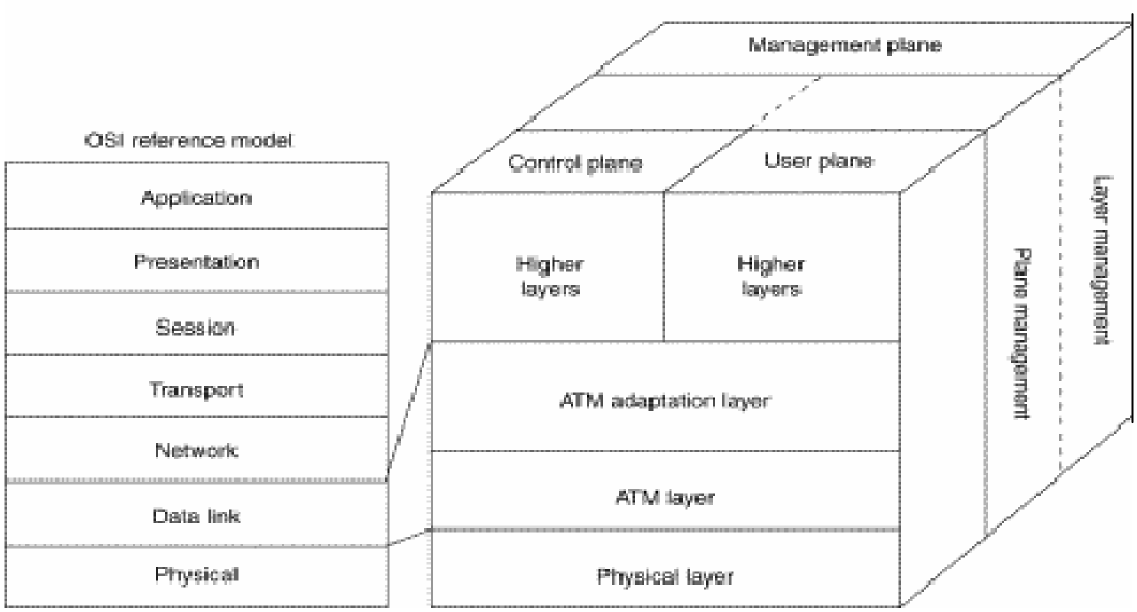

Figure 4.6.9 ATM reference model

- **ATM layer—**Combined with the ATM adaptation layer, the ATM layer is roughly analogous to the data link layer of the OSI reference model. The ATM layer is responsible for the simultaneous sharing of virtual circuits over a physical link (cell multiplexing) and passing cells through the ATM network (cell relay). To do this, it uses the VPI and VCI information in the header of each ATM cell.
- **ATM adaptation layer (AAL)—**Combined with the ATM layer, the AAL is roughly analogous to the data link layer of the OSI model. The AAL is responsible for isolating higher-layer protocols from the details of the ATM processes. The adaptation layer prepares user data for conversion into cells and segments the data into 48-byte cell payloads.

Finally, the higher layers residing above the AAL accept user data, arrange it into packets, and hand it to the AAL.

The ATM Physical Layer

The main functions of the ATM physical layer are as follows:

- Cells are converted into a bit stream,
- The transmission and receipt of bits on the physical medium are controlled,
- ATM cell boundaries are tracked,
- Cells are packaged into the appropriate types of frames for the physical medium.

The ATM physical layer is divided into two parts: the physical medium-dependent (PMD) sub layer and the transmission convergence (TC) sub layer.

The PMD sub layer provides two key functions.

- It synchronizes transmission and reception by sending and receiving a continuous flow of bits with associated timing information.
- It specifies the physical media for the physical medium used, including connector types and cable.

The TC sub layer has four functions:

- Cell delineation, it maintains ATM cell boundaries, allowing devices to locate cells within a stream of bits.
- Generates and checks the header error control code to ensure valid data.
- Cell-rate decoupling, maintains synchronization and inserts or suppresses idle (unassigned) ATM cells to adapt the rate of valid ATM cells to the payload capacity of the transmission system.
- Transmission frame adaptation packages ATM cells into frames acceptable to the particular physical layer implementation.

ATM Layer

The ATM layer provides routing, traffic management, switching and multiplexing services. It processes outgoing traffic by accepting 48-byte segment from the AAL sub-layers and transforming them into 53-byte cell by addition of a 5-byte header. The cell header format is already discussed in section 4.6.4. And the switching part and virtual connections were discussed in 4.6.5.

Adaptation Layers

ATM adaptation layers allow existing packet networks to connect to ATM facilities. AAL Protocol accepts transmission from upper layer services (e.g.: packet data) and map them into fixed-sized ATM cells. These transmissions can be of any type, variable or fixed data rate. At the receiver, this process is reversed and segments are reassembled into their original formats and passed to the receiving services. Instead of one protocol for all types of data, the ATM standard divides the AAL layer into categories, each supporting the requirements of different types of applications. There are four types of data streams that are identified: Constant-bit rate, variable bit-rate, connection oriented packet data transfer, connectionless packet data transfer. In addition to dividing AAL by category (AAL1, AAL2 and so on), ITU-T also divides it on the basis of functionality. Each AAL layer is actually divided into two layers: the **convergence** sub-layer and **Segmentation and reassembly** (SAR) sub-layer. Table 4.6.1 below gives a brief description of these data streams and various ATM adaptation layers which are used for each of them.

Table 4.6.1 Mapping of various data types and ATM adaptation layers

Service Class	**Quality of Service Parameter**	**ATM Adaptation layers**
Constant Bit rate (CBR)	This class is used for emulating circuit switching. The cell rate is constant with time. CBR applications are quite sensitive to cell-delay variation. Examples of applications that can use CBR are telephone traffic (i.e., nx64 kbps), videoconferencing, and television.	**AAL1:** AAL1, a connection-oriented service, is suitable for handling constant bit rate sources (CBR), such as voice and videoconferencing. AAL1 requires timing synchronization between the source and the destination. For this reason, AAL1 depends on a medium, such as SONET, that supports clocking. The AAL1 process prepares a cell for transmission in three steps. First, synchronous samples (for example, 1 byte of data at a sampling rate of 200 microseconds) are inserted into the Payload field. Second, Sequence Number (SN) and Sequence Number Protection (SNP) fields are added to provide information that the receiving AAL1 uses to verify that it has received cells in the correct order. Third, the remainder of the Payload field is filled with enough single bytes to equal 48 bytes.

<table>
<tr><td>Variable Bit Rate - non-real time (VBR–NRT)</td><td>This class allows users to send traffic at a rate that varies with time depending on the availability of user information. Statistical multiplexing is provided to make optimum use of network resources. Multimedia e-mail is an example of VBR–NRT.</td><td rowspan="2">AAL 2: The AAL2 process uses 44 bytes of the cell payload for user data and reserves 4 bytes of the payload to support the AAL2 processes.

VBR traffic is characterized as either real-time (VBR-RT) or as non-real-time (VBR-NRT). AAL2 supports both types of VBR traffic.</td></tr>
<tr><td>Variable bit rate–real time (VBR–RT)</td><td>This class is similar to VBR–NRT but is designed for applications that are sensitive to cell-delay variation. Examples for real-time VBR are voice with speech activity detection (SAD) and interactive compressed video.</td></tr>
<tr><td>Connection oriented packet transfer or available bit rate (ABR)</td><td>This class of ATM services provides rate-based flow control and is aimed at data traffic such as file transfer and e-mail. Although the standard does not require the cell transfer delay and cell-loss ratio to be guaranteed or minimized, it is desirable for switches to minimize delay and loss as much as possible. Depending upon the state of congestion in the network, the source is required to control its rate. The users are allowed to declare a minimum cell rate, which is guaranteed to the connection by the network.</td><td rowspan="2">AAL3/4: AAL3/4 supports both connection-oriented and connectionless data. AAL3/4 prepares a cell for transmission in four steps. First, the convergence sub layer (CS) creates a protocol data unit (PDU) by prepending a beginning/end tag header to the frame and appending a length field as a trailer. Second, the segmentation and reassembly (SAR) sub layer fragments the PDU and prepends a header to it. Then the SAR sub layer appends a CRC-10 trailer to each PDU fragment for error control. Finally, the completed SAR PDU becomes the Payload field of an ATM cell to which the ATM layer prepends the standard ATM header.

AAL 5: AAL5 is the primary AAL for data and supports both connection-oriented and connectionless data. It is used to transfer most non-SMDS data, such as classical IP over ATM and LAN Emulation (LANE). AAL5 also is known as the simple and efficient adaptation layer (SEAL)</td></tr>
<tr><td>Connectionless data transfer or unspecified bit rate (UBR)</td><td>This class is the catch-all, other class and is widely used today for TCP/IP.</td></tr>
</table>

4.6.7 ATM Applications

ATM is used in both LANs and WANs; let's have a look at few of the possible applications.

ATM WANs: ATM is basically a WAN technology that delivers cell over long distances. Here ATM is mainly used to connect LANs or other WANs together. A router between ATM network and the other network serves as an end point. This router has two stacks of protocols: one belonging to ATM and other belonging to other protocol.

ATM LANs: High data rate (155 and 622 Mbps) of ATM technology attracted designers to think of implementing ATM technology in LANs too. At the surface level, to implement an ATM LAN ATM switch will replace the traditional Ethernet switch, in a switched LAN. But few things have to be kept in mind and software modules would be needed to map the following differences between the two technologies:

- **Connectionless versus connection-oriented**: ATM is a virtual connection oriented technology, while traditional Ethernet uses connectionless protocols.
- **Physical address versus virtual circuit identifier**: In the Traditional LAN packets are routed based on the source and destination addresses, while in ATM cells are routed based on the virtual circuit identifiers (VPI-VCI pair).

LAN Emulation: LAN Emulation (LANE) is a standard defined by the ATM Forum that gives to stations attached via ATM the same capabilities that they normally obtain from legacy LANs, such as Ethernet and Token Ring. As the name suggests, the function of the LANE protocol is to emulate a LAN on top of an ATM network. Specifically, the LANE protocol defines mechanisms for emulating either an IEEE 802.3 Ethernet or an 802.5 Token Ring LAN.

Multimedia virtual private networks and managed services: Service providers are building on their ATM networks to offer a broad range of services. Examples include managed ATM, LAN, voice and video services (these being provided on a per-application basis, typically including customer-located equipment and offered on an end-to-end basis), and full-service virtual private-networking capabilities (these including integrated multimedia access and network management).

Frame-relay backbones: Frame-relay service providers are deploying ATM backbones to meet the rapid growth of their frame-relay services to use as a networking infrastructure for a range of data services and to enable frame relay to ATM service internetworking services.

Internet backbones: Internet service providers are likewise deploying ATM backbones to meet the rapid growth of their frame-relay services, to use as a networking infrastructure for a range of data services, and to enable Internet class-of-service offerings and virtual private intranet services.

Residential broadband networks: ATM is the networking infrastructure of choice for carriers establishing residential broadband services, driven by the need for highly scalable solutions.

Carrier infrastructures for the telephone and private-line networks: Some carriers have identified opportunities to make more-effective use of their SONET/SDH fiber infrastructures by building an ATM infrastructure to carry their telephony and private-line traffic.

Fill In The Blanks:

1. ATM is abbreviated as ________________ Mode.
2. ATM is an ITU-T standard for _______ relay.
3. ATM is a technology that has its history in the development of _____________ in the 1970s and 1980s
4. ATM is _________________ that's why the cells, which follow the same path may not reach destination in order.
5. _________ layer in ATM reformats the data received from other network.
6. _________ AAL type supports the data stream that has constant bit rate.
7. AAL3/4 supports both __________________ and __________ data.
8. ________ supports connectionless data transfer or unspecified bit rate (UBR).
9. VCI is abbreviated as __________________ identifier.
10. VPI is abbreviated as __________________ identifier.

Answers:

1. Asynchronous Transfer
2. Cell
3. Broadband ISDN
4. virtual circuit switching
5. ATM physical layer
6. AAL1
7. Connection oriented, connectionless
8. AAL5
9. Virtual Circuit
10. Virtual Path

Short Questions:

1. Explain how virtual connection is defined in ATM?

Ans: ATM is based on packet switching by setting up a virtual circuit. Several nodes are connected with the hosts and various switches. It sets up a logical path, not a permanent path. A table is setup in each station in the path and information is entered in these tables. Once a virtual path is created, there is an entry about the root through data communication takes place. Again another table (information also provided) will provide information about the virtual circuit. Routing takes place based on these tables.

2. Explain few benefits of ATM.

Ans: The high-level benefits delivered through ATM services can be summarized as follows:

- **Dynamic bandwidth for bursty traffic** meeting application needs and delivering high utilization of networking resources; most applications are or can be viewed as inherently bursty, for example voice is bursty, as both parties are neither speaking at once nor all the time; video is bursty, as the amount of motion and required resolution varies over time.
- **Smaller header** with respect to the data to make the efficient use of bandwidth.
- **Can handle Mixed network traffic very efficiently:** The variety of packet sizes make traffic unpredictable. All network equipments should incorporate elaborate software systems to manage the various sizes of packets. ATM handles these problems efficiently with the fixed size cell.
- **Cell network:** All data is loaded into identical cells that can be transmitted with complete predictability and uniformity.
- **Class-of-service support** for multimedia traffic allowing applications with varying throughput and latency requirements to be met on a single network.
- **Scalability in speed** and network size supporting link speeds of T1/E1 to OC–12 (622 Mbps).

3. What is the difference between UNI and NNI?

Ans: ATM switches support two primary types of interfaces: UNI and NNI. The UNI (User-Network Interface) connects ATM end systems (such as hosts and routers) to an ATM switch. The NNI (Network-Network Interface) connects two ATM switches. Depending on whether the switch is owned and located at the customer's premises or is publicly owned and operated by the telephone company, UNI and NNI can be further subdivided into public and private UNIs and NNIs.

4. What do you mean by LAN emulation?

Ans: LAN Emulation: ***LAN Emulation (LANE)*** is a standard defined by the ATM Forum that gives to stations attached via ATM the same capabilities that they normally obtain from legacy LANs, such as Ethernet and Token Ring. As the name suggests, the function of the LANE protocol is to emulate a LAN on top of an ATM network. Specifically, the LANE protocol defines mechanisms for emulating either an IEEE 802.3 Ethernet or an 802.5 Token Ring LAN.

5. Explain Payload Type field in the ATM frame Format.

Ans: Payload Type (PT)—Indicates in the first bit whether the cell contains user data or control data. If the cell contains user data, the bit is set to 0. If it contains control data, it is set to 1. The second bit indicates congestion (0 = no congestion, 1 = congestion), and the third bit indicates whether the cell is the last in a series of cells that represent a single AAL5 frame (1 = last cell for the frame).

6. Explain HEC field in the ATM frame Format.
Ans: Header Error Control (HEC)—Calculates checksum only on the first 4 bytes of the header. HEC can correct a single bit error in these bytes, thereby preserving the cell rather than discarding it.

7. Explain different planes in ATM references model.
Ans: The ATM reference model is composed of the following planes, which span all layers:

- **Control—**This plane is responsible for generating and managing signaling requests. **User—**This plane is responsible for managing the transfer of data.
- **Management—**This plane contains two components:
  - Layer management manages layer-specific functions, such as the detection of failures and protocol problems.
  - Plane management manages and coordinates functions related to the complete system.