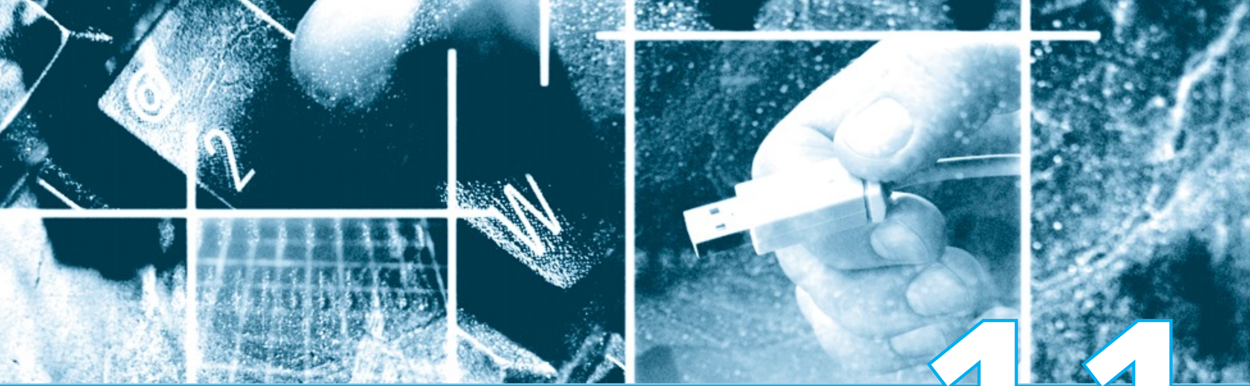

CHAPTER 11

ASYNCHRONOUS TRANSFER MODE

One man had a vision of railways that would link all the mainline railroad termini. His name was Charles Pearson and, though born the son of an upholsterer, he became Solicitor to the city of London. There had previously been a plan for gaslit subway streets through which horse-drawn traffic could pass. This was rejected on the grounds that such sinister tunnels would become lurking places for thieves. Twenty years before his system was built, Pearson envisaged a line running through "a spacious archway," well-lit and well-ventilated. His was a scheme for trains in a drain.

—*King Solomon's Carpet.* Barbara Vine (Ruth Rendell)

KEY POINTS

- ATM is a streamlined packet transfer interface. ATM makes use of fixed-size packets, called cells. The use of a fixed size and fixed format results in an efficient scheme for transmission over high-speed networks.
- Some form of transmission structure must be used to transport ATM cells. One option is the use of a continuous stream of cells, with no multiplex frame structure imposed at the interface. Synchronization is on a cell-by-cell basis. The second option is to place the cells in a synchronous time-division multiplex envelope. In this case, the bit stream at the interface has an external frame based on the Synchronous Digital Hierarchy (SDH).
- ATM provides both real-time and non-real-time services. An ATM-based network can support a wide range of traffic, include synchronous TDM streams such as T-1, using the constant bit rate (CBR) service; compressed voice and video, using the real-time variable bit rate (rt-VBR) service; traffic with specific quality-of-service requirements, using the non-real-time VBR (nrt-VBR) service; and IP-based traffic using the available bit rate (ABR), unspecified bit rate (UBR), and guaranteed frame rate (GFR) services.

Asynchronous transfer mode (ATM), also known as cell relay, takes advantage of the reliability and fidelity of modern digital facilities to provide faster packet switching than X.25.

11.1 PROTOCOL ARCHITECTURE

Asynchronous transfer mode is in some ways similar to packet switching using X.25 and to frame relay. Like packet switching and frame relay, ATM involves the transfer of data in discrete chunks. Also, like packet switching and frame relay, ATM allows multiple logical connections to be multiplexed over a single physical interface. In the

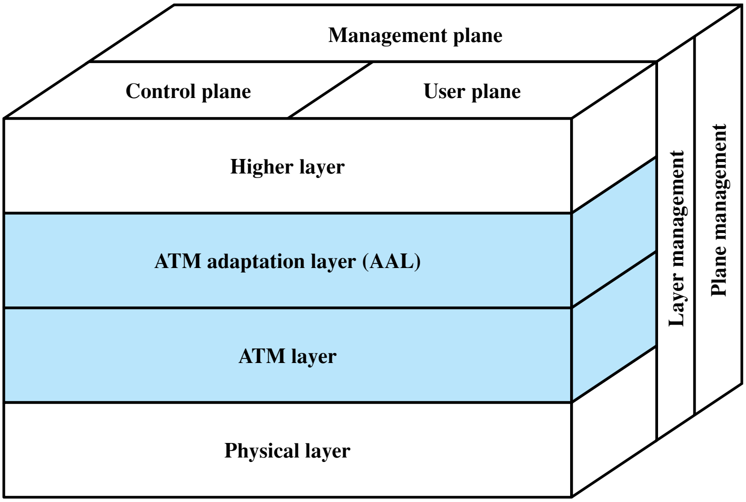

Figure 11.1 ATM Protocol Architecture

case of ATM, the information flow on each logical connection is organized into fixed-size packets, called **cells**.

ATM is a streamlined protocol with minimal error and flow control capabilities. This reduces the overhead of processing ATM cells and reduces the number of overhead bits required with each cell, thus enabling ATM to operate at high data rates. Further, the use of fixed-size cells simplifies the processing required at each ATM node, again supporting the use of ATM at high data rates.

The standards issued for ATM by ITU-T are based on the protocol architecture shown in Figure 11.1, which illustrates the basic architecture for an interface between user and network. The physical layer involves the specification of a transmission medium and a signal encoding scheme. The data rates specified at the physical layer range from 25.6 Mbps to 622.08 Mbps. Other data rates, both higher and lower, are possible.

Two layers of the protocol architecture relate to ATM functions. There is an ATM layer common to all services that provides packet transfer capabilities, and an ATM adaptation layer (AAL) that is service dependent. The ATM layer defines the transmission of data in fixed-size cells and defines the use of logical connections. The use of ATM creates the need for an adaptation layer to support information transfer protocols not based on ATM. The AAL maps higher-layer information into ATM cells to be transported over an ATM network, then collects information from ATM cells for delivery to higher layers.

The protocol reference model involves three separate planes:

- **User plane:** Provides for user information transfer, along with associated controls (e.g., flow control, error control)
- **Control plane:** Performs call control and connection control functions
- **Management plane:** Includes plane management, which performs management functions related to a system as a whole and provides coordination between all the planes, and layer management, which performs management functions relating to resources and parameters residing in its protocol entities

11.2 ATM LOGICAL CONNECTIONS

Logical connections in ATM are referred to as **virtual channel connections** (VCCs). A VCC is analogous to a virtual circuit in X.25; it is the basic unit of switching in an ATM network. A VCC is set up between two end users through the network and a variable-rate, full-duplex flow of fixed-size cells is exchanged over the connection. VCCs are also used for user-network exchange (control signaling) and network-network exchange (network management and routing).

For ATM, a second sublayer of processing has been introduced that deals with the concept of virtual path (Figure 11.2). A **virtual path connection** (VPC) is a bundle of VCCs that have the same endpoints. Thus, all of the cells flowing over all of the VCCs in a single VPC are switched together.

The virtual path concept was developed in response to a trend in high-speed networking in which the control cost of the network is becoming an increasingly higher proportion of the overall network cost. The virtual path technique helps contain the control cost by grouping connections sharing common paths through the network into a single unit. Network management actions can then be applied to a small number of groups of connections instead of a large number of individual connections.

Several advantages can be listed for the use of virtual paths:

- **Simplified network architecture:** Network transport functions can be separated into those related to an individual logical connection (virtual channel) and those related to a group of logical connections (virtual path).
- **Increased network performance and reliability:** The network deals with fewer, aggregated entities.
- **Reduced processing and short connection setup time:** Much of the work is done when the virtual path is set up. By reserving capacity on a virtual path connection in anticipation of later call arrivals, new virtual channel connections can be established by executing simple control functions at the endpoints of the virtual path connection; no call processing is required at transit nodes. Thus, the addition of new virtual channels to an existing virtual path involves minimal processing.
- **Enhanced network services:** The virtual path is used internal to the network but is also visible to the end user. Thus, the user may define closed user groups or closed networks of virtual channel bundles.

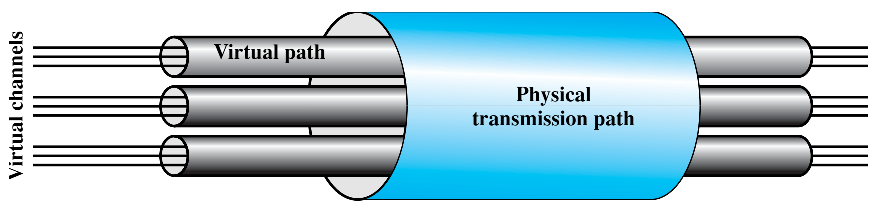

Figure 11.2 ATM Connection Relationships

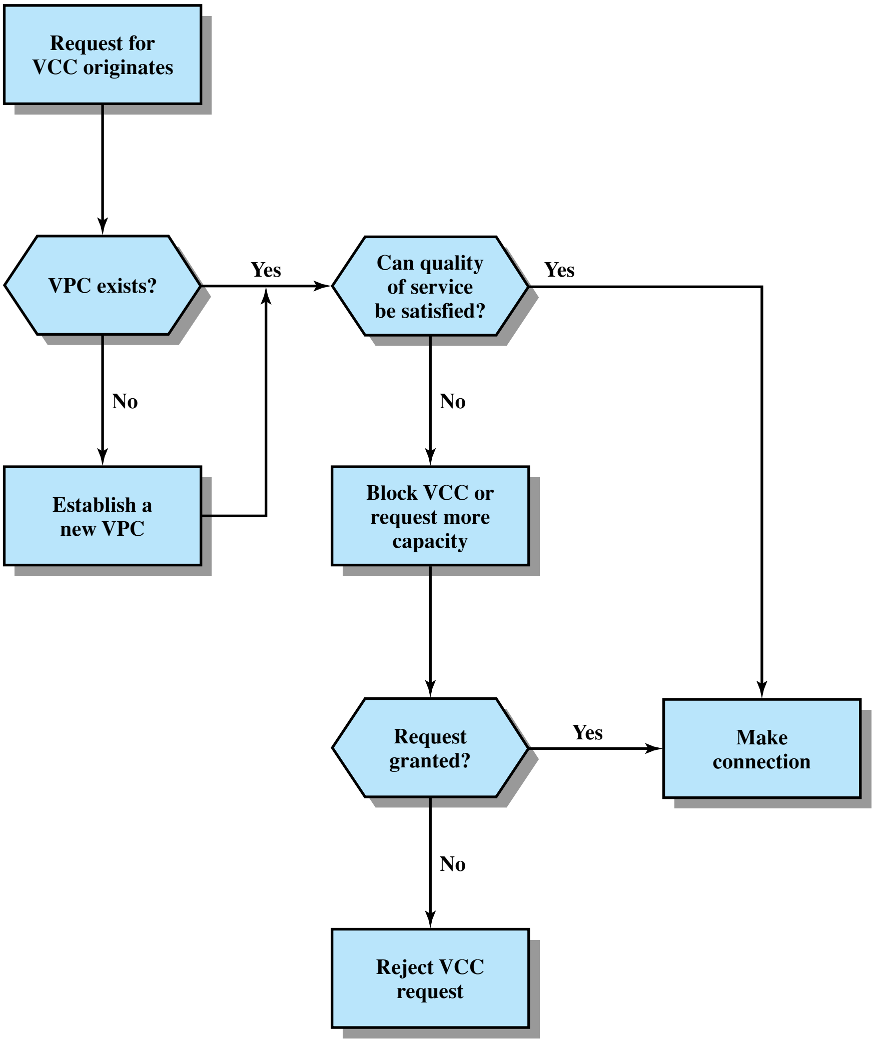

Figure 11.3 Call Establishent Using Virtual Paths

Figure 11.3 suggests in a general way the call establishment process using virtual channels and virtual paths. The process of setting up a virtual path connection is decoupled from the process of setting up an individual virtual channel connection:

- The virtual path control mechanisms include calculating routes, allocating capacity, and storing connection state information.
- To set up a virtual channel, there must first be a virtual path connection to the required destination node with sufficient available capacity to support the virtual channel, with the appropriate quality of service. A virtual channel is set up by storing the required state information (virtual channel/virtual path mapping).

Table 11.1 Virtual Path/Virtual Channel Terminology

Virtual Channel (VC)	A generic term used to describe unidirectional transport of ATM cells associated by a common unique identifier value.
Virtual Channel Link	A means of unidirectional transport of ATM cells between a point where a VCI value is assigned and the point where that value is translated or terminated.
Virtual Channel Identifier (VCI)	A unique numerical tag that identifies a particular VC link for a given VPC.
Virtual Channel Connection (VCC)	A concatenation of VC links that extends between two points where ATM service users access the ATM layer. VCCs are provided for the purpose of user-user, user-network, or network-network information transfer. Cell sequence integrity is preserved for cells belonging to the same VCC.
Virtual Path	A generic term used to describe unidirectional transport of ATM cells belonging to virtual channels that are associated by a common unique identifier value.
Virtual Path Link	A group of VC links, identified by a common value of VPI, between a point where a VPI value is assigned and the point where that value is translated or terminated.
Virtual Path Identifier (VPI)	Identifies a particular VP link.
Virtual Path Connection (VPC)	A concatenation of VP links that extends between the point where the VCI values are assigned and the point where those values are translated or removed, i.e., extending the length of a bundle of VC links that share the same VPI. VPCs are provided for the purpose of user-user, user-network, or network-network information transfer.

The terminology of virtual paths and virtual channels used in the standard is a bit confusing and is summarized in Table 11.1. Whereas most of the network-layer protocols that we deal with in this book relate only to the user-network interface, the concepts of virtual path and virtual channel are defined in the ITU-T Recommendations with reference to both the user-network interface and the internal network operation.

Virtual Channel Connection Uses

The endpoints of a VCC may be end users, network entities, or an end user and a network entity. In all cases, cell sequence integrity is preserved within a VCC: that is, cells are delivered in the same order in which they are sent. Let us consider examples of the three uses of a VCC:

- **Between end users:** Can be used to carry end-to-end user data; can also be used to carry control signaling between end users, as explained later. A VPC between end users provides them with an overall capacity; the VCC organization of the VPC is up to the two end users, provided the set of VCCs does not exceed the VPC capacity.

- **Between an end user and a network entity:** Used for user-to-network control signaling, as discussed subsequently. A user-to-network VPC can be used to aggregate traffic from an end user to a network exchange or network server.
- **Between two network entities:** Used for network traffic management and routing functions. A network-to-network VPC can be used to define a common route for the exchange of network management information.

Virtual Path/Virtual Channel Characteristics

ITU-T Recommendation I.150 lists the following as characteristics of virtual channel connections:

- **Quality of service (QoS):** A user of a VCC is provided with a QoS specified by parameters such as cell loss ratio (ratio of cells lost to cells transmitted) and cell delay variation.
- **Switched and semipermanent virtual channel connections:** A switched VCC is an on-demand connection, which requires a call control signaling for setup and tearing down. A semipermanent VCC is one that is of long duration and is set up by configuration or network management action.
- **Cell sequence integrity:** The sequence of transmitted cells within a VCC is preserved.
- **Traffic parameter negotiation and usage monitoring:** Traffic parameters can be negotiated between a user and the network for each VCC. The network monitors the input of cells to the VCC, to ensure that the negotiated parameters are not violated.

The types of traffic parameters that can be negotiated include average rate, peak rate, burstiness, and peak duration. The network may need a number of strategies to deal with congestion and to manage existing and requested VCCs. At the crudest level, the network may simply deny new requests for VCCs to prevent congestion. Additionally, cells may be discarded if negotiated parameters are violated or if congestion becomes severe. In an extreme situation, existing connections might be terminated.

I.150 also lists characteristics of VPCs. The first four characteristics listed are identical to those for VCCs. That is, QoS; switched and semipermanent VPCs; cell sequence integrity; and traffic parameter negotiation and usage monitoring are all also characteristics of a VPC. There are a number of reasons for this duplication. First, this provides some flexibility in how the network service manages the requirements placed upon it. Second, the network must be concerned with the overall requirements for a VPC, and within a VPC may negotiate the establishment of virtual channels with given characteristics. Finally, once a VPC is set up, it is possible for the end users to negotiate the creation of new VCCs. The VPC characteristics impose a discipline on the choices that the end users may make.

In addition, a fifth characteristic is listed for VPCs:

- **Virtual channel identifier restriction within a VPC:** One or more virtual channel identifiers, or numbers, may not be available to the user of the VPC but may be reserved for network use. Examples include VCCs used for network management.

Control Signaling

In ATM, a mechanism is needed for the establishment and release of VPCs and VCCs. The exchange of information involved in this process is referred to as control signaling and takes place on separate connections from those that are being managed.

For VCCs, I.150 specifies four methods for providing an establishment/release facility. One or a combination of these methods will be used in any particular network:

1. **Semipermanent VCCs** may be used for user-to-user exchange. In this case, no control signaling is required.
2. If there is no preestablished call control signaling channel, then one must be set up. For that purpose, a control signaling exchange must take place between the user and the network on some channel. Hence we need a permanent channel, probably of low data rate, that can be used to set up VCCs that can be used for call control. Such a channel is called a **meta-signaling channel**, as the channel is used to set up signaling channels.
3. The meta-signaling channel can be used to set up a VCC between the user and the network for call control signaling. This **user-to-network signaling virtual channel** can then be used to set up VCCs to carry user data.
4. The meta-signaling channel can also be used to set up a **user-to-user signaling virtual channel**. Such a channel must be set up within a preestablished VPC. It can then be used to allow the two end users, without network intervention, to establish and release user-to-user VCCs to carry user data.

For VPCs, three methods are defined in I.150:

1. A VPC can be established on a **semipermanent** basis by prior agreement. In this case, no control signaling is required.
2. VPC establishment/release may be **customer controlled**. In this case, the customer uses a signaling VCC to request the VPC from the network.
3. VPC establishment/release may be **network controlled**. In this case, the network establishes a VPC for its own convenience. The path may be network-to-network, user-to-network, or user-to-user.

11.3 ATM CELLS

The asynchronous transfer mode makes use of fixed-size cells, consisting of a 5-octet header and a 48-octet information field. There are several advantages to the use of small, fixed-size cells. First, the use of small cells may reduce queuing delay for a high-priority cell, because it waits less if it arrives slightly behind a lower-priority cell that has gained access to a resource (e.g., the transmitter). Second, it appears that fixed-size cells can be switched more efficiently, which is important for the very high data rates of ATM [PARE88]. With fixed-size cells, it is easier to implement the switching mechanism in hardware.

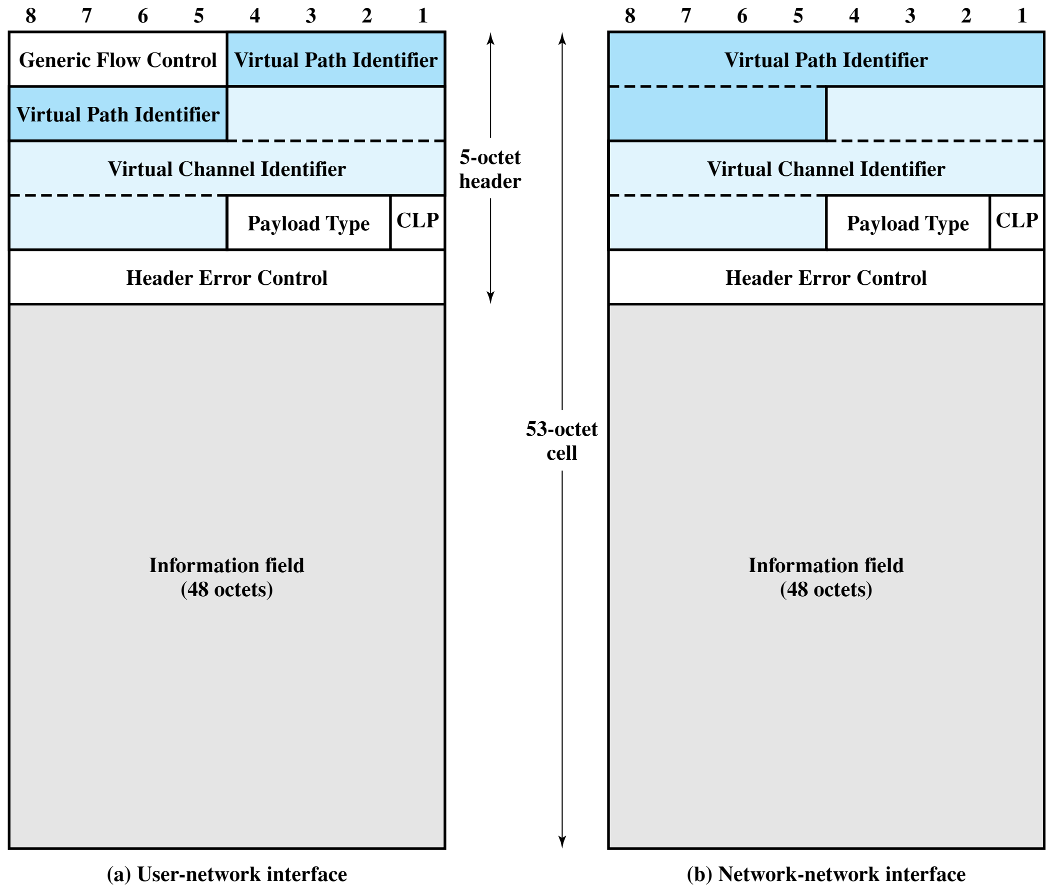

Figure 11.4 ATM Cell Format

Header Format

Figure 11.4a shows the cell header format at the user-network interface. Figure 11.4b shows the cell header format internal to the network.

The **Generic Flow Control** (GFC) field does not appear in the cell header internal to the network, but only at the user-network interface. Hence, it can be used for control of cell flow only at the local user-network interface. The field could be used to assist the customer in controlling the flow of traffic for different qualities of service. In any case, the GFC mechanism is used to alleviate short-term overload conditions in the network.

I.150 lists as a requirement for the GFC mechanism that all terminals be able to get access to their assured capacities. This includes all constant-bit-rate (CBR) terminals as well as the variable-bit-rate (VBR) terminals that have an element of guaranteed capacity (CBR and VBR are explained in Section 11.5). The current GFC mechanism is described in a subsequent subsection.

The **Virtual Path Identifier** (VPI) constitutes a routing field for the network. It is 8 bits at the user-network interface and 12 bits at the network-network interface. The

Table 11.2 Payload Type (PT) Field Coding

PT Coding	Interpretation		
0 0 0	User data cell,	congestion not experienced,	SDU-type = 0
0 0 1	User data cell,	congestion not experienced,	SDU-type = 1
0 1 0	User data cell,	congestion experienced,	SDU-type = 0
0 1 1	User data cell,	congestion experienced,	SDU-type = 1
1 0 0	OAM segment associated cell		
1 0 1	OAM end-to-end associated cell		
1 1 0	Resource management cell		
1 1 1	Reserved for future function		

SDU = Service Data Unit
OAM = Operations, Administration, and Maintenance

latter allows support for an expanded number of VPCs internal to the network, to include those supporting subscribers and those required for network management. The **Virtual Channel Identifier** (VCI) is used for routing to and from the end user.

The **Payload Type** (PT) field indicates the type of information in the information field. Table 11.2 shows the interpretation of the PT bits. A value of 0 in the first bit indicates user information (that is, information from the next higher layer). In this case, the second bit indicates whether congestion has been experienced; the third bit, known as the Service Data Unit (SDU)[1] type bit, is a one-bit field that can be used to discriminate two types of ATM SDUs associated with a connection. The term *SDU* refers to the 48-octet payload of the cell. A value of 1 in the first bit of the Payload Type field indicates that this cell carries network management or maintenance information. This indication allows the insertion of network-management cells onto a user's VCC without impacting the user's data. Thus, the PT field can provide inband control information.

The **Cell Loss Priority** (CLP) bit is used to provide guidance to the network in the event of congestion. A value of 0 indicates a cell of relatively higher priority, which should not be discarded unless no other alternative is available. A value of 1 indicates that this cell is subject to discard within the network. The user might employ this field so that extra cells (beyond the negotiated rate) may be inserted into the network, with a CLP of 1, and delivered to the destination if the network is not congested. The network may set this field to 1 for any data cell that is in violation of an agreement concerning traffic parameters between the user and the network. In this case, the switch that does the setting realizes that the cell exceeds the agreed traffic parameters but that the switch is capable of handling the cell. At a later point in the network, if congestion is encountered, this cell has been marked for discard in preference to cells that fall within agreed traffic limits.

[1]This is the term used in ATM Forum documents. In ITU-T documents, this bit is referred to as the ATM-user-to-ATM-user (AAU) indication bit. The meaning is the same.

The **Header Error Control** (HEC) field is used for both error control and synchronization, as explained subsequently.

Generic Flow Control

I.150 specifies the use of the GFC field to control traffic flow at the user-network interface (UNI) in order to alleviate short-term overload conditions. The actual flow control mechanism is defined in I.361. GFC flow control is part of a proposed controlled cell transfer (CCT) capability intended to meet the requirements of non-ATM LANs connected to a wide area ATM network [LUIN97]. In particular, CCT is intended to provide good service for high-volume bursty traffic with variable-length messages. In the remainder of this subsection, we examine the GFC mechanism, as so far standardized.

When the equipment at the UNI is configured to support the GFC mechanism, two sets of procedures are used: uncontrolled transmission and controlled transmission. In essence, every connection is identified as either subject to flow control or not. Of those subject to flow control, there may be one group of controlled connections (Group A) that is the default, or controlled traffic may be classified into two groups of controlled connections (Group A and Group B); these are known, respectively, as the one-queue and two-queue models. Flow control is exercised in the direction from the subscriber to the network by the network side.

First, we consider the operation of the GFC mechanism when there is only one group of controlled connections. The controlled equipment, called terminal equipment (TE), initializes two variables: TRANSMIT is a flag initialized to SET (1), and GO_CNTR, which is a credit counter, is initialized to 0. A third variable, GO_VALUE, is either initialized to 1 or set to some larger value at configuration time. The rules for transmission by the controlled device are as follows:

1. If TRANSMIT = 1, cells on uncontrolled connections may be sent at any time. If TRANSMIT = 0, no cells may be sent on either controlled or uncontrolled connections.
2. If a HALT signal is received from the controlling equipment, TRANSMIT is set to 0 and remains at zero until a NO_HALT signal is received, at which time TRANSMIT is set to 1.
3. If TRANSMIT = 1 and there is no cell to transmit on any uncontrolled connections, then
    - —If GO_CNTR > 0, then the TE may send a cell on a controlled connection. The TE marks that cell as a cell on a controlled connection and decrements GO_CNTR.
    - —If GO_CNTR = 0, then the TE may not send a cell on a controlled connection.
4. The TE sets GO_CNTR to GO_VALUE upon receiving a SET signal; a null signal has no effect on GO_CNTR.

Table 11.3 Generic Flow Control (GFC) Field Coding

	Uncontrolled	Controlling → controlled		Controlled → controlling	
		1-Queue Model	**2-Queue Model**	**1-Queue Model**	**2-Queue Model**
First bit	0	HALT(0)/ NO_HALT(1)	HALT(0)/ NO_HALT(1)	0	0
Second bit	0	SET(1)/NULL(0)	SET(1)/NULL(0) for Group A	cell belongs to controlled(1)/ uncontrolled(0)	cell belongs to Group A(1)/ or not (0)
Third bit	0	0	SET(1)/NULL(0) for Group B	0	cell belongs to Group B(1)/ or not (0)
Fourth bit	0	0	0	equipment is uncontrolled(0)/ controlled(1)	equipment is uncontrolled(0)/ controlled(1)

The HALT signal is used logically to limit the effective ATM data rate and should be cyclic. For example, to reduce the data rate over a link by half, the HALT command is issued by the controlling equipment so as to be in effect 50% of the time. This is done in a predictable, regular pattern over the lifetime of the physical connection.

For the two-queue model, there are two counters, each with a current counter value and an initialization value: GO_CNTR_A, GO_VALUE_A, GO_CNTR_B, and GO_VALUE_B. This enables the network to control two separate groups of connections.

Table 11.3 summarizes the rules for setting GFC bits.

Header Error Control

Each ATM cell includes an 8-bit HEC field that is calculated based on the remaining 32 bits of the header. The polynomial used to generate the code is $X^8 + X^2 + X + 1$. In most existing protocols that include an error control field, such as HDLC, the data that serve as input to the error code calculation are in general much longer than the size of the resulting error code. This allows for error detection. In the case of ATM, the input to the calculation is only 32 bits, compared to 8 bits for the code. The fact that the input is relatively short allows the code to be used not only for error detection but also, in some cases, for actual error correction. This is because there is sufficient redundancy in the code to recover from certain error patterns.

Figure 11.5 depicts the operation of the HEC algorithm at the receiver. At initialization, the receiver's error correction algorithm is in the default mode for single-bit error correction. As each cell is received, the HEC calculation and comparison is performed. As long as no errors are detected, the receiver remains in error correction mode. When an error is detected, the receiver will correct the error if it is a single-bit error or will detect that a multibit error has occurred. In either case, the receiver now moves to detection mode. In this mode, no attempt is made to

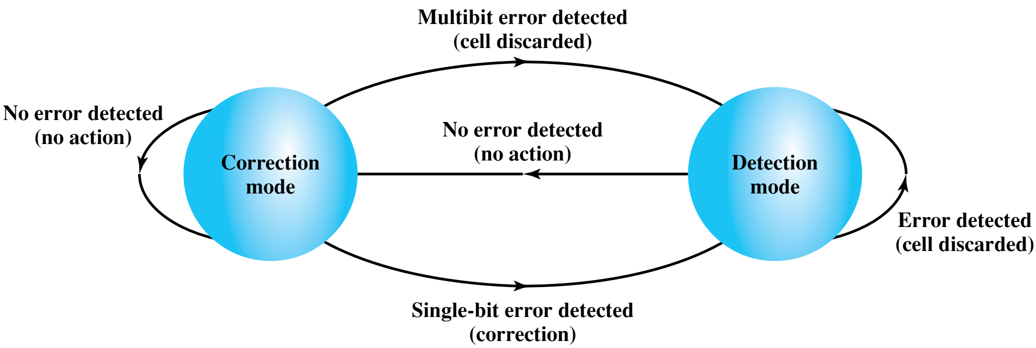

Figure 11.5 HEC Operation at Receiver

correct errors. The reason for this change is a recognition that a noise burst or other event might cause a sequence of errors, a condition for which the HEC is insufficient for error correction. The receiver remains in detection mode as long as errored cells are received. When a header is examined and found not to be in error, the receiver switches back to correction mode. The flowchart of Figure 11.6 shows the consequence of errors in the cell header.

The error protection function provides both recovery from single-bit header errors and a low probability of the delivery of cells with errored headers under bursty error conditions. The error characteristics of fiber-based transmission systems appear to be a mix of single-bit errors and relatively large burst errors. For some transmission systems, the error correction capability, which is more time-consuming, might not be invoked.

Figure 11.7, based on one in ITU-T I.432, indicates how random bit errors impact the probability of occurrence of discarded cells and valid cells with errored headers when HEC is employed.

11.4 TRANSMISSION OF ATM CELLS

I.432 specifies that ATM cells may be transmitted at one of several data rates: 622.08 Mbps, 155.52 Mbps, 51.84 Mbps, or 25.6 Mbps. We need to specify the transmission structure that will be used to carry this payload. Two approaches are defined in I.432: a cell-based physical layer and an SDH-based physical layer.[2] We examine each of these approaches in turn.

Cell-Based Physical Layer

For the cell-based physical layer, no framing is imposed. The interface structure consists of a continuous stream of 53-octet cells. Because there is no external frame imposed in the cell-based approach, some form of synchronization is needed. Synchronization is

[2]The SDH-based approach is not defined for 25.6 Mbps.

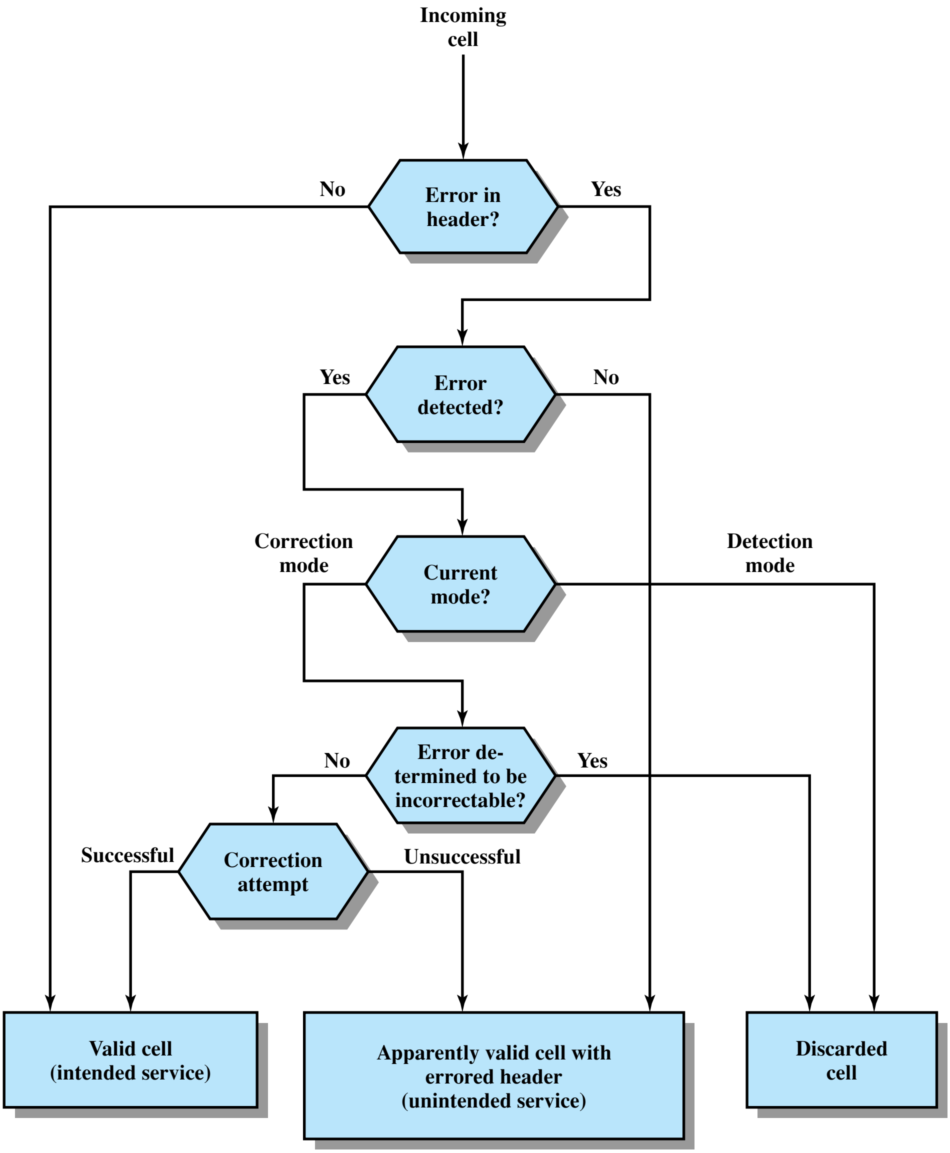

Figure 11.6 Effect of Error in Cell Header

achieved on the basis of the HEC field in the cell header. The procedure is as follows (Figure 11.8):

1. In the HUNT state, a cell delineation algorithm is performed bit by bit to determine if the HEC coding law is observed (i.e., match between received HEC and calculated HEC). Once a match is achieved, it is assumed that one header has been found, and the method enters the PRESYNC state.

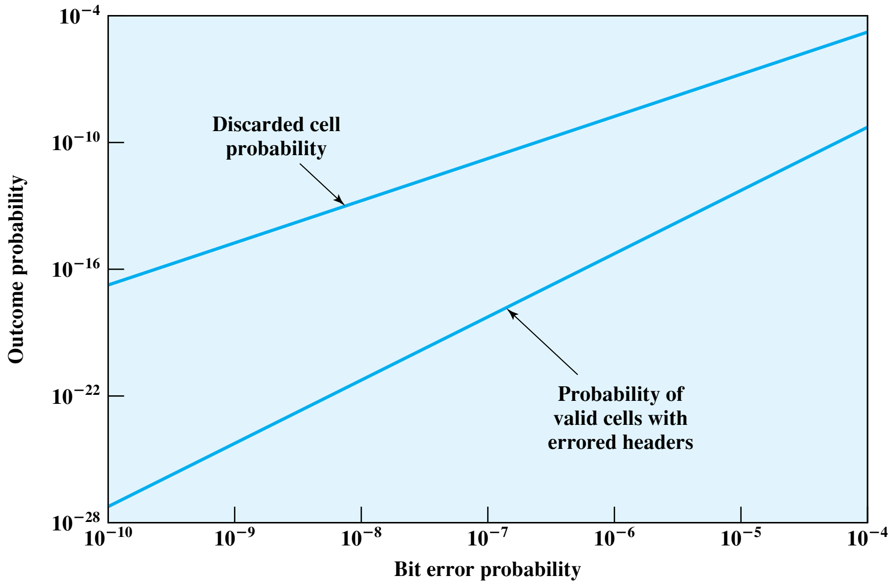

Figure 11.7 Impact of Random Bit Errors on HEC Performance

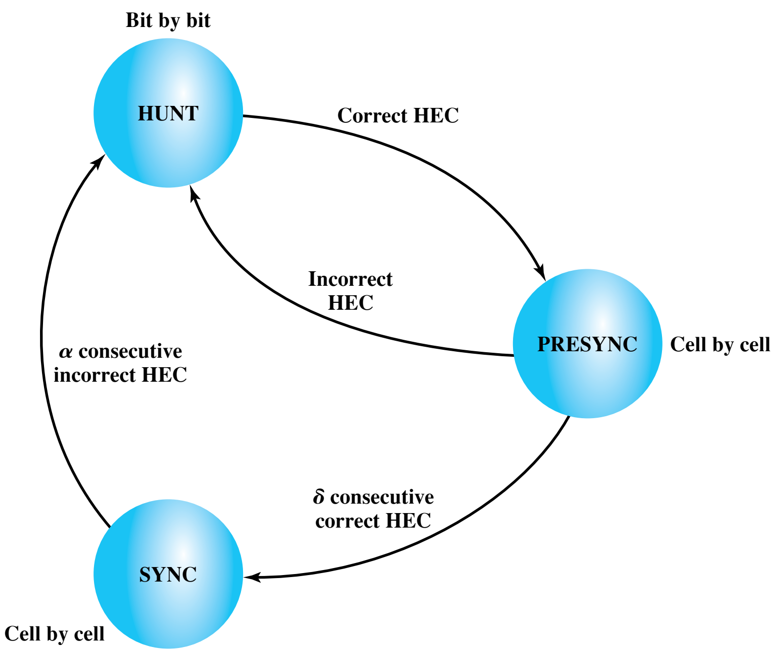

Figure 11.8 Cell Delineation State Diagram

2. In the PRESYNC state, a cell structure is now assumed. The cell delineation algorithm is performed cell by cell until the encoding law has been confirmed consecutively δ times.
3. In the SYNC state, the HEC is used for error detection and correction (see Figure 11.5). Cell delineation is assumed to be lost if the HEC coding law is recognized consecutively as incorrect α times.

The values of α and δ are design parameters. Greater values of δ result in longer delays in establishing synchronization but in greater robustness against false delineation. Greater values of α result in longer delays in recognizing a misalignment but in greater robustness against false misalignment. Figures 11.9 and 11.10, based on I.432, show the impact of random bit errors on cell delineation performance for various values of α and δ. The first figure shows the average amount of time that the receiver will maintain synchronization in the face of errors, with α as a parameter. The second figure shows the average amount of time to acquire synchronization as a function of error rate, with δ as a parameter.

The advantage of using a cell-based transmission scheme is the simplified interface that results when both transmission and transfer mode functions are based on a common structure.

SDH-Based Physical Layer

The SDH-based physical layer imposes a structure on the ATM cell stream. In this section, we look at the I.432 specification for 155.52 Mbps; similar structures are used at other data rates. For the SDH-based physical layer, framing is imposed

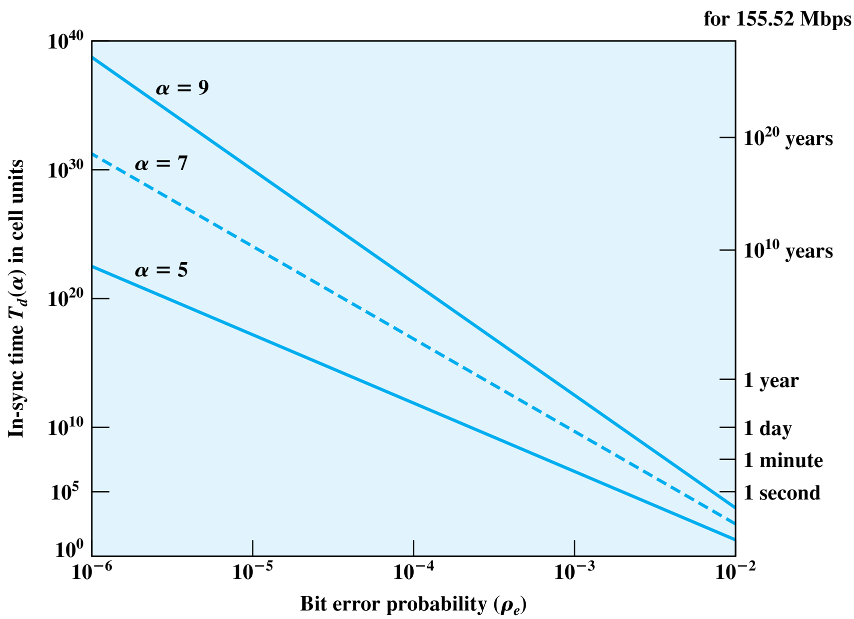

Figure 11.9 Impact of Random Bit Errors on Cell-Delineation Performance

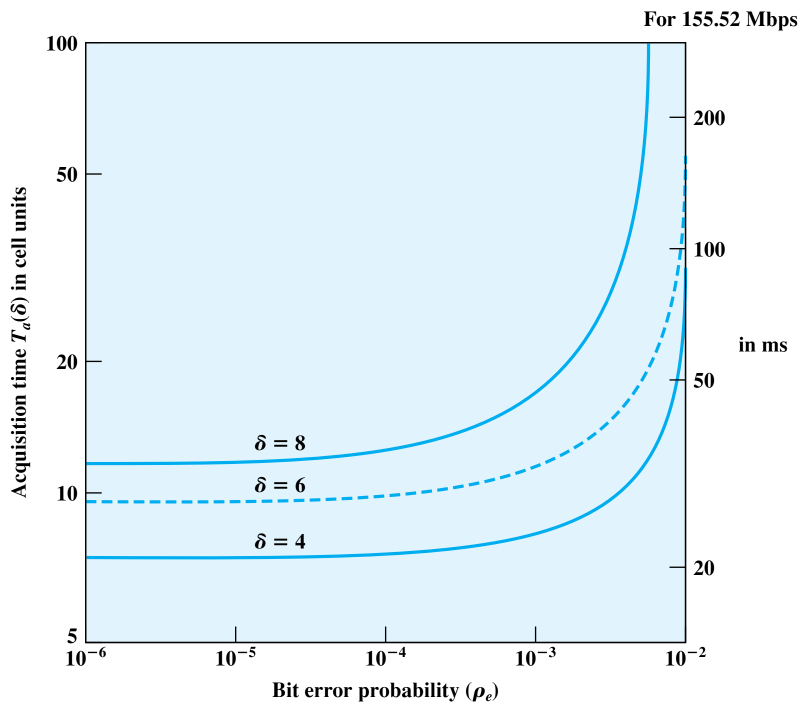

Figure 11.10 Acquisition Time versus Bit-Error Probability

using the STM-1 (STS-3) frame. Figure 11.11 shows the payload portion of an STM-1 frame (see Figure 8.11). This payload may be offset from the beginning of the frame, as indicated by the pointer in the section overhead of the frame. As can be seen, the payload consists of a 9-octet path overhead portion and the remainder, which contains ATM cells. Because the payload capacity (2340 octets) is not an integer multiple of the cell length (53 octets), a cell may cross a payload boundary.

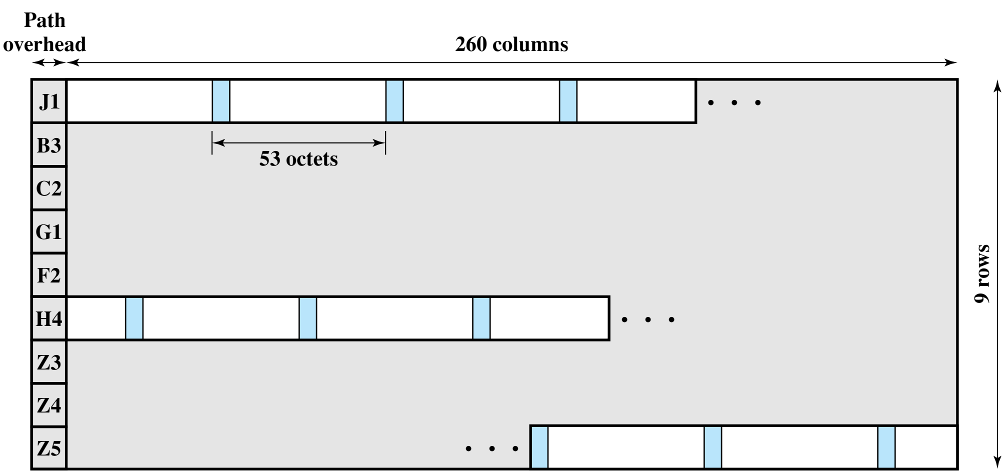

Figure 11.11 STM-1 Payload for SDH-Based ATM Cell Transmission

The H4 octet in the path overhead is set at the sending side to indicate the next occurrence of a cell boundary. That is, the value in the H4 field indicates the number of octets to the first cell boundary following the H4 octet. The permissible range of values is 0 to 52.

The advantages of the SDH-based approach include the following:

- It can be used to carry either ATM-based or STM-based (synchronous transfer mode) payloads, making it possible to initially deploy a high-capacity fiber-based transmission infrastructure for a variety of circuit-switched and dedicated applications and then readily migrate to the support of ATM.
- Some specific connections can be circuit switched using an SDH channel. For example, a connection carrying constant-bit-rate video traffic can be mapped into its own exclusive payload envelope of the STM-1 signal, which can be circuit switched. This may be more efficient than ATM switching.
- Using SDH synchronous multiplexing techniques, several ATM streams can be combined to build interfaces with higher bit rates than those supported by the ATM layer at a particular site. For example, four separate ATM streams, each with a bit rate of 155 Mbps (STM-1), can be combined to build a 622-Mbps (STM-4) interface. This arrangement may be more cost effective than one using a single 622-Mbps ATM stream.

11.5 ATM SERVICE CATEGORIES

An ATM network is designed to be able to transfer many different types of traffic simultaneously, including real-time flows such as voice, video, and bursty TCP flows. Although each such traffic flow is handled as a stream of 53-octet cells traveling through a virtual channel, the way in which each data flow is handled within the network depends on the characteristics of the traffic flow and the requirements of the application. For example, real-time video traffic must be delivered within minimum variation in delay.

We examine the way in which an ATM network handles different types of traffic flows in Chapter 13. In this section, we summarize ATM service categories, which are used by an end system to identify the type of service required. The following service categories have been defined by the ATM Forum:

- **Real-Time Service**
  —Constant bit rate (CBR)
  —Real-time variable bit rate (rt-VBR)
- **Non-Real-Time Service**
  —Non-real-time variable bit rate (nrt-VBR)
  —Available bit rate (ABR)
  —Unspecified bit rate (UBR)
  —Guaranteed frame rate (GFR)

Real-Time Services

The most important distinction among applications concerns the amount of delay and the variability of delay, referred to as jitter, that the application can tolerate. Real-time applications typically involve a flow of information to a user that is intended to reproduce that flow at a source. For example, a user expects a flow of audio or video information to be presented in a continuous, smooth fashion. A lack of continuity or excessive loss results in significant loss of quality. Applications that involve interaction between people have tight constraints on delay. Typically, any delay above a few hundred milliseconds becomes noticeable and annoying. Accordingly, the demands in the ATM network for switching and delivery of real-time data are high.

Constant Bit Rate (CBR) The CBR service is perhaps the simplest service to define. It is used by applications that require a fixed data rate that is continuously available during the connection lifetime and a relatively tight upper bound on transfer delay. CBR is commonly used for uncompressed audio and video information. Example of CBR applications include

- Videoconferencing
- Interactive audio (e.g., telephony)
- Audio/video distribution (e.g., television, distance learning, pay-per-view)
- Audio/video retrieval (e.g., video-on-demand, audio library)

Real-Time Variable Bit Rate (rt-VBR) The rt-VBR category is intended for time-sensitive applications; that is, those requiring tightly constrained delay and delay variation. The principal difference between applications appropriate for rt-VBR and those appropriate for CBR is that rt-VBR applications transmit at a rate that varies with time. Equivalently, an rt-VBR source can be characterized as somewhat bursty. For example, the standard approach to video compression results in a sequence of image frames of varying sizes. Because real-time video requires a uniform frame transmission rate, the actual data rate varies.

The rt-VBR service allows the network more flexibility than CBR. The network is able to statistically multiplex a number of connections over the same dedicated capacity and still provide the required service to each connection.

Non-Real-Time Services

Non-real-time services are intended for applications that have bursty traffic characteristics and do not have tight constraints on delay and delay variation. Accordingly, the network has greater flexibility in handling such traffic flows and can make greater use of statistical multiplexing to increase network efficiency.

Non-Real-Time Variable Bit Rate (nrt-VBR) For some non-real-time applications, it is possible to characterize the expected traffic flow so that the network can provide substantially improved QoS in the areas of loss and delay. Such applications can use the nrt-VBR service. With this service, the end system specifies a peak cell rate, a sustainable or average cell rate, and a measure of how bursty or clumped the cells may be. With this information, the network can allocate resources to provide relatively low delay and minimal cell loss.

The nrt-VBR service can be used for data transfers that have critical response-time requirements. Examples include airline reservations, banking transactions, and process monitoring.

Unspecified Bit Rate (UBR) At any given time, a certain amount of the capacity of an ATM network is consumed in carrying CBR and the two types of VBR traffic. Additional capacity is available for one or both of the following reasons: (1) Not all of the total resources have been committed to CBR and VBR traffic, and (2) the bursty nature of VBR traffic means that at some times less than the committed capacity is being used. All of this unused capacity could be made available for the UBR service. This service is suitable for applications that can tolerate variable delays and some cell losses, which is typically true of TCP-based traffic. With UBR, cells are forwarded on a first-in-first-out (FIFO) basis using the capacity not consumed by other services; both delays and variable losses are possible. No initial commitment is made to a UBR source and no feedback concerning congestion is provided; this is referred to as a **best-effort service**. Examples of UBR applications include

- Text/data/image transfer, messaging, distribution, retrieval
- Remote terminal (e.g., telecommuting)

Available Bit Rate (ABR) Bursty applications that use a reliable end-to-end protocol such as TCP can detect congestion in a network by means of increased round-trip delays and packet discarding. This is discussed in Chapter 20. However, TCP has no mechanism for causing the resources within the network to be shared fairly among many TCP connections. Further, TCP does not minimize congestion as efficiently as is possible using explicit information from congested nodes within the network.

To improve the service provided to bursty sources that would otherwise use UBR, the ABR service has been defined. An application using ABR specifies a peak cell rate (PCR) that it will use and a minimum cell rate (MCR) that it requires. The network allocates resources so that all ABR applications receive at least their MCR capacity. Any unused capacity is then shared in a fair and controlled fashion among all ABR sources. The ABR mechanism uses explicit feedback to sources to assure that capacity is fairly allocated. Any capacity not used by ABR sources remains available for UBR traffic.

An example of an application using ABR is LAN interconnection. In this case, the end systems attached to the ATM network are routers.

Figure 11.12 suggests how a network allocates resources during a steady-state period of time (no additions or deletions of virtual channels).

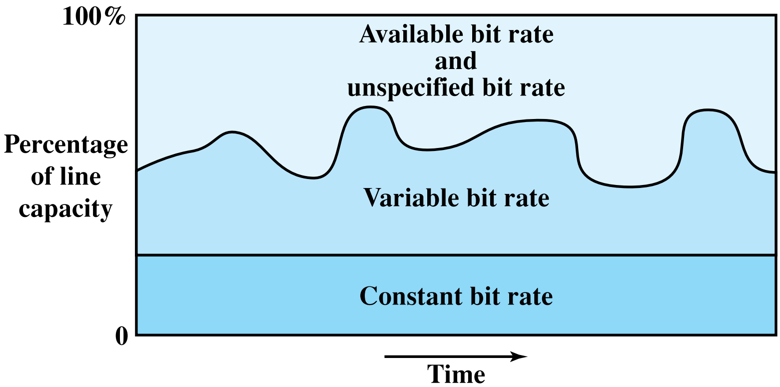

Figure 11.12 ATM Bit Rate Services

Guaranteed Frame Rate (GFR) The most recent addition to the set of ATM service categories is GFR, which is designed specifically to support IP backbone subnetworks. GFR provides better service than UBR for frame-based traffic, including IP and Ethernet. A major goal of GFR is to optimize the handling of frame-based traffic that passes from a LAN through a router onto an ATM backbone network. Such ATM networks are increasingly being used in large enterprise, carrier, and Internet service provider networks to consolidate and extend IP services over the wide area. While ABR is also an ATM service meant to provide a greater measure of guaranteed packet performance over ATM backbones, ABR is relatively difficult to implement between routers over an ATM network. With the increased emphasis on using ATM to support IP-based traffic, especially traffic that originates on Ethernet LANs, GFR may offer the most attractive alternative for providing ATM service.

One of the techniques used by GFR to provide improved performance compared to UBR is to require that network elements be aware of frame or packet boundaries. Thus, when congestion requires the discard of cells, network elements must discard all of the cells that comprise a single frame. GFR also allows a user to reserve capacity for each GFR VC. The user is guaranteed that this minimum capacity will be supported. Additional frames may be transmitted if the network is not congested.

11.6 RECOMMENDED READING AND WEB SITES

[MCDY99] and [BLAC99a] provide good coverage of ATM. The virtual path/virtual channel approach of ATM is examined in [SATO90], [SATO91], and [BURG91].

[GARR96] provides a rationale for the ATM service categories and discuss the traffic management implications of each. [ARMI93] and [SUZU94] discuss AAL and compare types 3/4 and 5.

ARMI93 Armitage, G., and Adams, K. "Packet Reassembly During Cell Loss." *IEEE Network*, September 1995.

BLAC99a Black, U. *ATM Volume I: Foundation for Broadband Networks.* Upper Saddle River, NJ: Prentice Hall, 1992.

BURG91 Burg, J., and Dorman, D. "Broadband ISDN Resource Management: The Role of Virtual Paths." *IEEE Communications Magazine*, September 1991.

GARR96 Garrett, M. "A Service Architecture for ATM: From Applications to Scheduling." *IEEE Network*, May/June 1996.

MCDY99 McDysan, D., and Spohn, D. *ATM: Theory and Application.* New York: McGraw-Hill, 1999.

SATO90 Sato, K.; Ohta, S.; and Tokizawa, I. "Broad-Band ATM Network Architecture Based on Virtual Paths." *IEEE Transactions on Communications*, August 1990.

SATO91 Sato, K.; Ueda, H.; and Yoshikai, M. "The Role of Virtual Path Crossconnection." *IEEE LTS*, August 1991.

SUZU94 Suzuki, T. "ATM Adaptation Layer Protocol." *IEEE Communications Magazine*, April 1995.

Recommended Web sites:

- **ATM Hot Links:** Excellent collection of white papers and links maintained by the University of Minnesota.
- **MFA Forum:** An industry forum that promotes ATM and other packet-based technologies. Contains white papers, vendor information, and links.
- **Cell Relay Retreat:** Contains archives of the cell-relay mailing list, links to numerous ATM-related documents, and links to many ATM-related Web sites.

11.7 KEY TERMS, REVIEW QUESTIONS, AND PROBLEMS

Key Terms

asynchronous transfer mode (ATM) ATM adaptation layer (AAL) available bit rate (ABR) cell loss priority (CLP) constant bit rate (CBR)	generic flow control (GFR) guaranteed frame rate (GFR) header error control (HEC) non-real-time variable bit rate (nrt-VBR) payload type	real-time variable bit rate (rt-VBR) service data unit (SDU) unspecified bit rate (UBR) variable bit rate (VBR) virtual channel virtual path

Review Questions

11.1. How does ATM differ from frame relay?
11.2. What are the relative advantages and disadvantages of ATM compared to frame relay?
11.3. What is the difference between a virtual channel and a virtual path?
11.4. What are the advantages of the use of virtual paths?
11.5. What are the characteristics of a virtual channel connection?
11.6. What are the characteristics of a virtual path connection?
11.7. List and briefly explain the fields in an ATM cell.
11.8. Briefly explain two methods for transmitting ATM cells.
11.9. List and briefly define the ATM service categories.

Problems

11.1 List all 16 possible values of the GFC field and the interpretation of each value (some values are illegal).

11.2 One key design decision for ATM was whether to use fixed or variable length cells. Let us consider this decision from the point of view of efficiency. We can define transmission efficiency as

$$N = \frac{\text{Number of information octets}}{\text{Number of information octets} + \text{Number of overhead octets}}$$

a. Consider the use of fixed-length packets. In this case the overhead consists of the header octets. Define

L = Data field size of the cell in octets

H = Header size of the cell in octets

X = Number of information octets to be transmitted as a single message

Derive an expression for *N*. *Hint*: The expression will need to use the operator $\lceil \cdot \rceil$, where $\lceil Y \rceil$ = the smallest integer greater than or equal to Y.

b. If cells have variable length, then overhead is determined by the header, plus the flags to delimit the cells or an additional length field in the header. Let Hv = additional overhead octets required to enable the use of variable-length cells. Derive an expression for *N* in terms of *X, H,* and Hv.

c. Let $L = 48$, $H = 5$, and $Hv = 2$. Plot N versus message size for fixed- and variable-length cells. Comment on the results.

11.3 Another key design decision for ATM is the size of the data field for fixed-size cells. Let us consider this decision from the point of view of efficiency and delay.

a. Assume that an extended transmission takes place, so that all cells are completely filled. Derive an expression for the efficiency *N* as a function of *H and L*.

b. Packetization delay is the delay introduced into a transmission stream by the need to buffer bits until an entire packet is filled before transmission. Derive an expression for this delay as a function of *L* and the data rate *R* of the source.

c. Common data rates for voice coding are 32 kbps and 64 kbps. Plot packetization delay as a function of *L* for these two data rates; use a left-hand *y*-axis with a maximum value of 2 ms. On the same graph, plot transmission efficiency as a function of *L;* use a right-hand *y*-axis with a maximum value of 100%. Comment on the results.

11.4 Consider compressed video transmission in an ATM network. Suppose standard ATM cells must be transmitted through five switches. The data rate is 43 Mbps.

a. What is the transmission time for one cell through one switch?

b. Each switch may be transmitting a cell from other traffic all of which we assume to have lower (non-preemptive for the cell) priority. If the switch is busy transmitting a cell, our cell has to wait until the other cell completes transmission. If the switch is free our cell is transmitted immediately. What is the maximum time from when a typical video cell arrives at the first switch (and possibly waits) until it is finished being transmitted by the fifth and last one? Assume that you can ignore propagation time, switching time, and everything else but the transmission time and the time spent waiting for another cell to clear a switch.

c. Now suppose we know that each switch is utilized 60% of the time with the other low priority traffic. By this we mean that with probability 0.6 when we look at a switch it is busy. Suppose that if there is a cell being transmitted by a switch, the average delay spent waiting for a cell to finish transmission is one-half a cell transmission time. What is the average time from the input of the first switch to clearing the fifth?

d. However, the measure of most interest is not delay but jitter, which is the variability in the delay. Use parts (b) and (c) to calculate the maximum and average variability, respectively, in the delay.

In all cases assume that the various random events are independent of one another; for example, we ignore the burstiness typical of such traffic.

11.5 In order to support IP service over an ATM network, IP datagrams must first be segmented into a number of ATM cells before sending them over the ATM network. As ATM does not provide cell loss recovery, the loss of any of these cells will result in the loss of the entire IP packet. Given

PC = cell loss rate in the ATM network

n = number of cells required to transmit a single IP datagram

PP = IP-packet loss rate

a. Derive an expression for PP, and comment on the resulting expression.

b. What ATM service would you use to get the best possible performance?