NEAR EAST UNIVERSITY

FACULTY OF ENGINEERING

DEPARTMENT OF COMPUTER ENGINEERING

BROADBAND ATM

GRADUATION PROJECT
COM 400

STUDENT: BERTAN GÖZLÜLÜ

STUDENT NO: 950444

SUPERVISOR: PROF.DR.FAKHREDDIN MAMEDOW

LEFKOSA , JUNE-2000

CONTENTS

CHAPTER 1

CHAPTER 4

CHAPTER 5

Acknowledgement

I would like to thank Prof. Dr. Fakhreddin Mamedov who is also the supervisor of my graduation project and he only and really helps me to get all the description about ATM and BROADBAND and also I want to thank to my theachers Assist.Prof.Dr Adnan Khashman and my supervisor Mr.Tayseer Alshanableh for helping me in my five year in this university. And also thanks to my family for giving me this chance. Also I thank to all my teachers in the university and my friends.

ABSTRACT

An effort to cope with these drawbacks of the existing telecommunication infrastructure was made with the development and implementation of the Narrowband ISDN (N-ISDN).

A consortium of a few hundred companies regrouped under the ATM Forum is designing the ATM. The ATM Forum publishes recommendations that define the ATM protocols.

The ATM Forum has defined network management protocols. These protocols classify the attributes of the different network elements. For instance, the attributes of a virtual circuit connection are its status (whether it is up or down), the current quality of service. The ATM layer provides for the transparent transfer of fixed size ATM layer Service Data Units (ATM-SDUs) between communicating upper layer entities.

Presently, the only technology that can fulfill all the requirements in terms of bandwidth, flexibility, and interactivity the digital video services have is the ATM protocol.

Aim of this thesis is the analysis and interpretation of the ATM and BROADBAND ATM Networks and working profiles.

INTRODUCTION

A consortium of a few hundred companies regrouped under the ATM Forum is designing the ATM. The ATM Forum publishes recommendations that define the ATM protocols.

In this project introduces the concepts of layered protocols. The focus is on the open systems interconnection (OSI) model, because the architecture of ATM is based on OSI architecture. For this book, this model is examined in a general manner, with descriptions on the rationale for its use, and how it is used in networks. Other books treat the OSI model in more detail. Of course, in the spirit of the subject matter of this book, the OSI model is examined in relation to ATM.

The latter part of the chapter also provides an overview of several of the OSI and Internet protocols that operate at the network and transport layers in these two architectures. Knowledge of these systems is important if the reader is to understand how ATM supports data communications systems.

To be able to transport the MPEG compressed digital video signals through the network, from the encoder to the set top box for instance, a very powerful transmission mechanism is needed. Presently, the only technology that can fulfill all the requirements in terms of bandwidth, flexibility, and interactivity the digital video services have is the ATM protocol, the central element of the B-ISDN. This section will provide the background and the essential technical details of B-ISDN needed to enable a more detailed understanding of the key networking technologies, referred to in the later section dealing with "Video in Broadband Networks.

The ATM network must be able to adapt to unforeseen traffic patterns, for example, unusual bursts of traffic from the various end-user devices or applications. Also, the network must be able to shed traffic in certain conditions to prevent or read to congestion. In so doing, it ml's be able to enforce an allowable peak cell rate for each VPI/VCI connection, This means that when a user's traffic load is presented to the net work beyond a maximum peak rate, then it nay be discarded.

In addition to these responsibilities, the network must be able t monitor (police) traffic it must be able to monitor all VPI/VCI connections, and verify their correctness (that they are properly mapped into the network and operate effectively). The network must be able to detect problems, and emit alarms when certain troubling events are encountered.

An effective ATM network must be designed with an understanding that both the user and the network assume responsibility for certain QOS operations between them The user is responsible for agreeing to service contract with the ATM network that stipulates rules on the us' of the network, such as the amount of traffic that can be submitted in measured time period In turn, the network assumes the responsibility of supporting the user QOS requirements.

The ATM layer provides for the transparent transfer of fixed size ATM layer Service Data Units (ATM-SDUs) between communicating upper layer entities (e.g., AAT-entities). This transfer occurs on a pre-established ATM connection according to a traffic contract. A traffic contract is comprised of a QoS class, a vector of traffic parameters, a conformance definition and other items as specified in section 3.6. Each ATM end-point is expected to generate traffic, which conforms to these parameters. Enforcement of the traffic contract is optional at the Private UNI. The Public Network is expected to monitor the offered load and enforce the traffic contract.

PREFACES

In this project I want to explain the ATM and Broadband ATM technologies and networks for helping me and the other students to get knowledge about this subject.

Chapter 1: Gives us what is ATM. Also network operations and routing in ATM is written. If we want to say in one word Asynchronous Transfer Mode was written in here.

Chapter 2: Here the Layered protocols and architecture of ATM is explained.

Chapter 3: The entering Broadband Network Tech. And Broadband ATM

Chapter 4: In this chapter traffic management is explained. Also allocation of bandwidth and UPC usage parameters was written.

Chapter 5: this chapter involves al the ATM and funtions of ATM layer management systems also ATM cell structures.

CHAPTER-1

ASYNCHRONOUS TRANSFER MODE

1.1Architecture

A consortium of a few hundred companies regrouped under the ATM Forum is designing the ATM. The ATM Forum publishes recommendations that define the ATM protocols.

Cells. The SNP can correct single bit errors and detect multiple bit errors in the sequence number.

AAL-5 carries IF packets or other payloads with little overhead The Cs packages the information into CS-SDU' s with a length equals to a multiple of 48 bytes The CS-SDU a header whose use is under study. The 2-byte alignment field (AL) is filler the padding (FAD) and length (LEN) are required to accommodate variable-length payloads The CRC is calculated over the complete CS-SDU. The SAR then puts the CS-SDU into back-to-back 48-byte cells with no additional control information. The higher layer (e.g detects lost cells. TOP or UDP

1.2 Network Operations and Maintenance

Consider a simple network that consists of a single ATM switch attached to a number of computers. The network manager may set up permanent virtual paths between the computers through the switch the manager allocates a specific VPI for each pair of computers. For instance, say that the virtual path between computer *A* and computer *B* uses VPI = 17. To communicate with *B*, computer *A* uses VPI = 17 and some VCI. The virtual path is setup permanently with resources reserved for it in the switch. For instance, VPI 17 may be allocated a transmission rate of 10 MBPS and 2 Mbytes of buffer space inside the switch to absorb traffic fluctuations.

In large network, it is wasteful to allocate resources permanently to possible virtual paths that might conceivably be set up. In such a network, it is preferable to set up Inn large network; it is wasteful to allocate resources permanently to all possible virtual paths on demand. A computer communicates with the network to request a connection, using a VPI/VCI reserved for call setup requests. Internally, the network switches use reserved VPI/VCI s to exchange the control information they need to set up and monitor connections. We explain some of the procedures that the network uses to setup connections in Section 1.2. Here we explain the basic ideas behind network maintenance.

The objective of network maintenance is to monitor the connections and to take corrective actions as needed. The maintenance protocols supervise the connections by using OAM (operations and maintenance) cells. Some OAM cells indicate an alarm, others signal that the destination fails to receive the user cells; others are "loopback" cells that an intermediate switch must loop back.

We examine a concrete scenario to illustrate the OAM protocols. In figure 1.1, a virtual path is set up between the computers A and B with VPI = 17. A few virtual circuit connections belong to this virtual path: these connections do not use the reserved VCIs 3 and 4. An OAM connection with VPI = 17 and VCI = 4 monitors the virtual path VCI = 17 from end to end. Another OAM connection (VPI = 17, VCI = 3) supervises the connection between the end user *A* and the first network switch *V*

At some time, *B* notices that the user cells from *A* do not arrive, as they should. The DAM protocol implemented in *B* sends a "far end receive failure" (EERF) cell to *A* along the (17,4) OAM connection, When it gets that cell, the OAM protocol in *A* starts a procedure to locate the faulty network element, The OAM sends a loopback cell along the OAM connection (17, 4), asking switch *W* to loop it back. We assume that the faulty element is between *V* and *W* and prevents the loopback cell from coming back to *A*. The OAM protocol inside *a* then sends a loopback cell along (17,4) asking *V* to loop it back. The cell comes back to *A*. The OAM protocol then knows that the failure is somewhere between *V* and *W* and the OAM protocol sends an alarm describing the problem to the network operator.

We sketched the maintenance procedure for a virtual path. The ATM Forum has defined similar procedures. With the corresponding format of OAM cells, for maintaining a virtual circuit, a transmission path. A communication link, and transmission equipment between electronic regenerators.

In addition to defining network maintenance protocols, the ATM Forum has defined network management protocols. These protocols classify the attributes of the different network elements. For instance, the attributes of a virtual circuit connection are its status (whether it is up or down), the current quality of service (cell loss rate, delay statistic), and descriptors of he traffic that it carries The attribute's of an ***ATM link*** include the maximum numbers of virtual circuits and virtual paths it can carry and the current numbers it caries. The network management protocols also specify how to read these attributes and how to modify those that can be controlled (such as turning off a faulty transmitter for repair).

Using the management protocols a network management software can construct a management information base (MIB) which maintains an up-to-data picture of the status of the network. The network operator uses this MIB to take corrective actions and to plan network modifications and upgrades.

Figure 1-1
OAM cells that monitor a VPC.

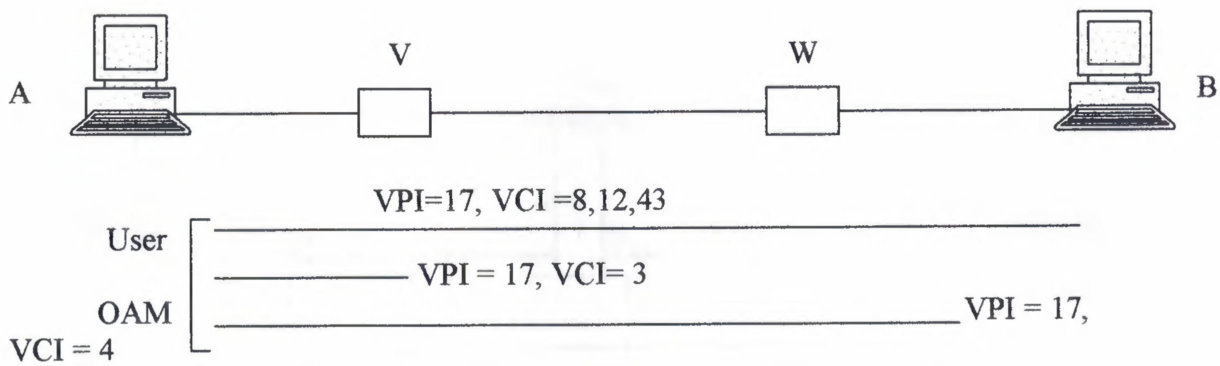

1.3 Routing in ATM

We explain how the switches allocate VIPs and the routing tables they maintain. We then outline some of the main ideas in the routing algorithms that the ATM Forum is considering for ATM Finally. We discuss two commonly used designs of ATM switches.

1.3.1 Routing Tables

Each ATM switch maintains a routing table that indicates on which link it should send the cells. A cell belongs to a virtual path that is identified by the VPI of the cell. Thus, the outgoing link is determined by the VPI of the cell.

However, there is a slight twist to this procedure. The VIPs of different virtual paths need only be unique on each link to differentiate them Consider the situation in Figure 1.2. In that network, three switches interconnect the five computers *A*, B.........*E*. three virtual paths-represented by thick lines-are set up successively: from *A* to *C* from B to *E* and from *C* to *D*. in that order. These virtual paths each carry a number or virtual circuit connections

Figure 1-2
The virtual paths and their VIPs along the links of an ATM network

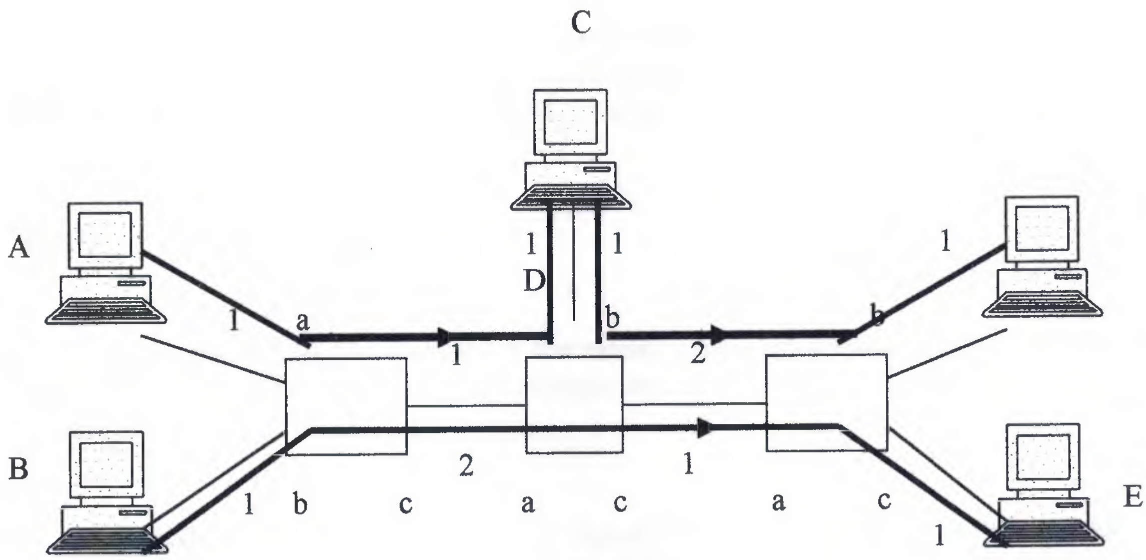

a, 1-c, 1
b, 1-c, 2

a, 1-b, 1
a, 2-c, 1

a, 1-c, 1
a, 2-b, 1

In the figure, the numbers on the thick lines are the VPI of each virtual path along each link of the network. When the virtual path from *A* to *C* is first set up, the network allocates the VPI number 1 to that path along each link, When the virtual path from *B* to F is set up next. The VPI number 1 is not used between *B* and the first switch, so that the network allocates that number to the path. Although the virtual path from *A* to *C* already uses the number 1 on other links choosing 1 for this new path does not create any risk of confusion. On the next link, the path from *B* to *E* cannot use the number 1, which is already used on that link. The network allocates the smallest free number on that link to the path, i.e., the number 2. The procedure continues for the other links and the other virtual path from *C* to *D*.

Note that the VPI of the path from B to *E* is I before the first switch, then 2 on the next link. Then 1 again for the next two links. Each switch maintains a routing table that it updates whenever a new connection is set up or a connection terminates. The table has one entry per connection, the entry has the following format: incoming link, incoming VPI-outgoing link, outgoing VPI (see figure).

1.4 Network Node Interface

The ATM Forum is considering a protocol called PNNI, for private network-node interface, for routing in private ATM networks. This protocol. Or something like it will probably be recommended for public ATM networks as well. The two main objectives of PNNI are that the routing should be based on the quality of service that connections request and that the algorithms and protocols should be scalable. PNNI has two parts: addressing and routing. We address these topics separately.

Routing

Consider the left part in Figure 1.3 the circles represent subnetworks and the lines are ATM links or some more complex connections such as virtual paths. Node *A* wants to set up a connection to D.

In a nutshell. The steps of the routing algorithm are as follows. Each link and each connection through a subnetwork is described by a vector of attributes x (i). These attributes are distributed to all the nodes by floating. Node A uses the attributes {x (1), x (9)} to calculate a preferred path to D. We assume that this preferred path is (A, B, C, D).

Figure 1-3
Left: a network of subnetworks. Right: crankback during call setup.

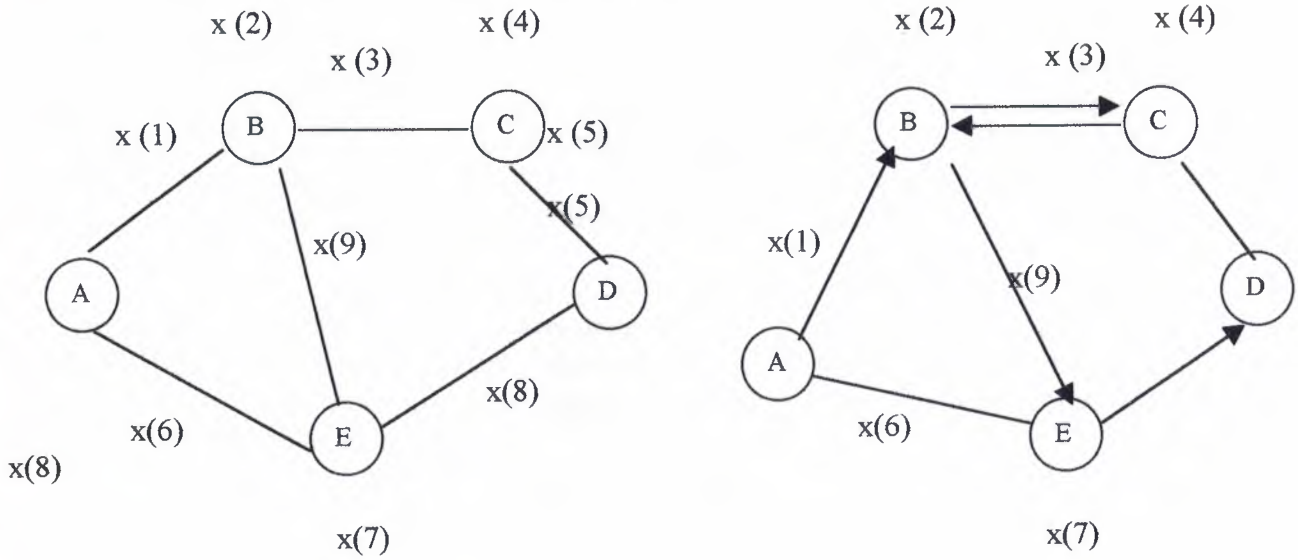

Node A then sends a connection request toward D along that preferred path using source routing. As it gets that request, B checks whether it can accept it. That is, *B* verifies it can carry the new call with traffic descriptors and QoS attributes specified in the request message. Assume that *B* accepts the request: it forwards it to *C*. Assume that *C* does not accept the request node ***C*** notifies B of its rejection.

Node *B* then calculates a preferred path to *D*. based on network status information, as A did earlier but with more current information. Say that B prefers the path (B. E, D). Node B then sends the request toward D along that path and the procedure continues in the figure, we assume that this second step succeeds and that the final path is (A, *B, E, D).* This mechanism is called *crankback*

How does A calculate a preferred path? The current ATM Forum recommendation specifies that the attributes .x (i} should describe for connection i the delay (maximum and maximum variation), the cell loss ratio, the "spare capacity," and some measure of desirability. Node A estimates the effect of adding the connection that it is trying to route on the attributes of the links Node *A* then uses these modified attributes to execute an algorithm similar to Dijkstra shortest path algorithm that determines the preferred path. You will note that since the paths have multiple attributes. the algorithm must weigh the different attributes to select a preferred path.

Addressing

Addressing is inspired from the hierarchical addressing of 1nternet, except that the number of levels in the hierarchy can be much larger (up to 105!) than just subnets and autonomous systems. Consider Figure 1.4. The network administrators assign ATM addresses to the nodes. The address of a node indicates its membership in the hierarchy. For instance, in Figure 1.4, the node with address 1.1.2 belongs to the subgroup 1 of group 1. We call this subgroup "group 1.1." Similarly, node; 2.2.3 belongs to the group 2.2. (We assumed a two-level hierarchy to keep the figure simple.)

The network is self-organizing. Each node talks to its neighbors to identify the members of the group of the some level, starting from the lowest level. Then the group members elect a leader which represents the group for the next level and the procedure continues until the top level is reached

In the example of the figure, the nodes 1.1.1, 1.1.2, and 1.1.3 discover that they all belong to group 1.1. Group 1.1 then elects a leader, say node 1.1.1. that we call "node 1.1" By flooding their link metrics inside group 1.1 the members of that group learn path metrics to each other and node 1.1, can calculate path metrics to other groups of the same hierarchical levels. Thus node 1.1 finds out it can reach node 1.2 and it can estimate the metric of the shortest path to that node. The nodes 2.1 and 2.2 perform the same steps The procedure then continues at the next level up. That is, nodes 2.1 and 2.2 learn that they both belong to group I and they elect a group leader, say node 1.1. We call this leader "node 1." (Thus, node 1 is in fact node 1.1.1 in our example.) Node 1 discovers it can reach node 2 and it can calculate the shortest path.

Figure 1-4

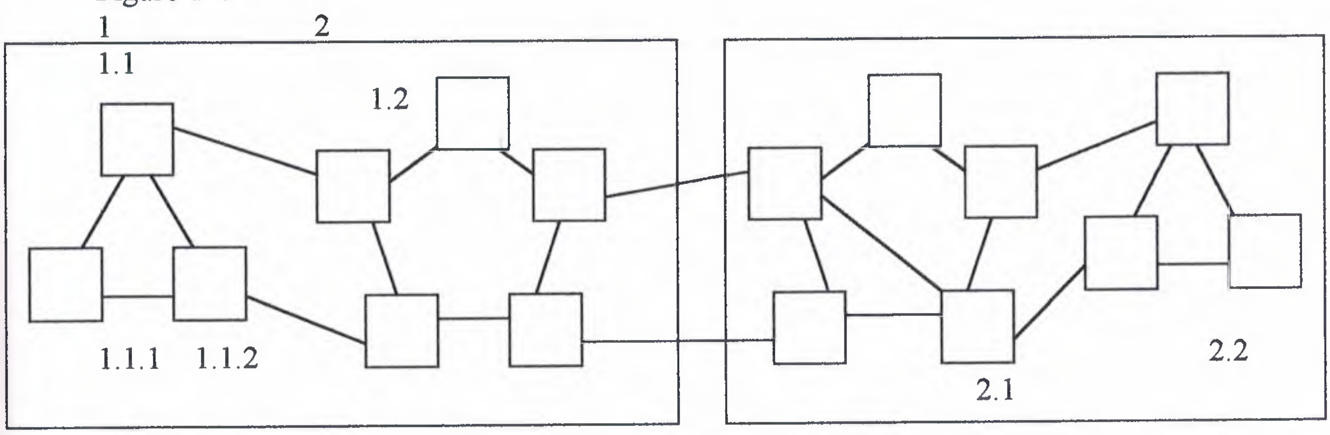

In this process, node 1.1 has a table with its group members and attributes of the paths to them. Node 1.1 gives that information to node 1. Node 1 gives the aggregate information it gets from nodes 1.1 and 1.2 to node 2. The information is then fed down to the group members Eventually node 1.1.2 gets the attributes of paths to 2.2.3. Node 1.1.2 can then use the procedure we explained above to select a path to 2.2.3 and send a request the request for a connection from 1.1.2 to 2.2.3 may then be accepted. Possibly after crankback.

1.5 Switch Designs

A number of designs have been studied for ATM switches, In this section we explain two designs that are used for fast switches with a limited number of ports (a few dozens). Other designs are better suited for switches with a very large number of ports (thousands) Please refer to the literature if you want a discussion of those designs. The two most commonly used ATM switch designs are shown in Figure 1.5: input buffer and output buffer

1.5.1 Input Buffer Switch

In an input buffer switch, cells enter a buffer as they arrive at the switch. Cells wait in line in the buffers until tile switching fabric can send them out on their output link, The advantage of this design is that the buffer memory runs only at the line rare. For instance, if the links operate at 2.4 Gbps. then the input buffers need to be able to store bits at that rate, not fasten

In an actual implementation, some fast electronic device converts the serial bit stream that arrives on a link into N parallel bit streams. This serial to parallel conversion divides the speed of the electronics of the switch by a factor N A parallel to serial conversion takes place at the output of the switch.

Figure 1-5
Input buffer and output buffer switch.

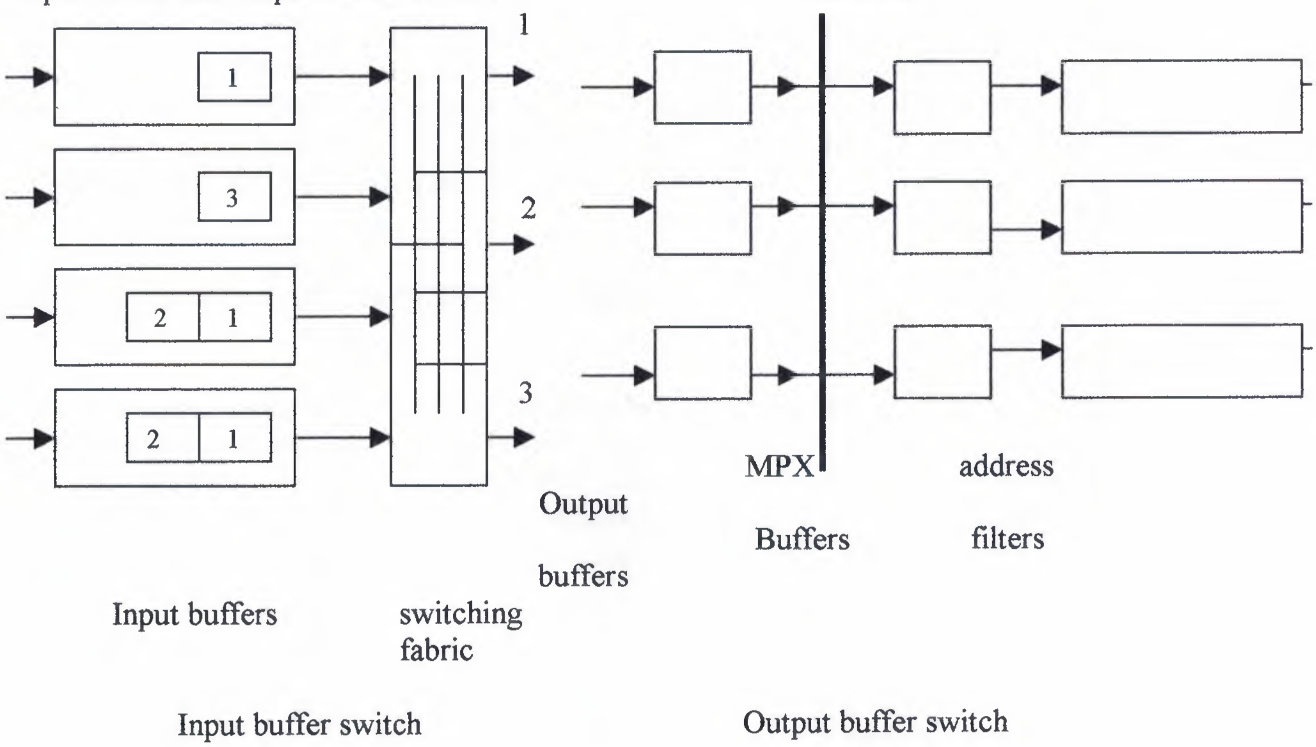

The disadvantage of the input butter design is that "head-of-line blocking" reduces the throughput or the switch. This blocking effect is illustrated in Figure 1.5. Three cells at the head of their input buffer arc destined to output link 1. The switching fabric can send only one of these three cells at a time to output link I. Assume that the switching fabric transmits the cell from input buffer 1 to output link 1 and that at the same time it sends the cell from the second input buffer the third output link In that case cells in the third and fourth output buffers that want to go to the second output link are blocked. The switching fabric could send one of these two cells if the cells at the head of heir waiting line did not block them

To reduce head-of-line blocking in an input buffer switch one may replace the single input buffer of each input line by a few parallel input buffers. For instance assume that in the input buffer switch in Figure 1.5, we put three buffers at each input line, one for each output line. As a cell arrives, we send it to the input buffer that corresponds to its output link. With this arrangement we can eliminate head-of-line blocking altogether The switching fabric must now choose which cells it should send to the output links At any given time, the switching fabric can choose among a potentially large number of such possibilities. The switch must make this choice in one cell transmission time. Researchers have designed fast scheduling algorithms that maximize the throughput of the switch and achieve a very low delay

1.5.2 Output Buffer Switch

The output buffer switch in Figure 1.5 transmits all the arriving cells on an internal bus. At each output link, an output buffer copies the cells destined for that output Such a design does not suffer from head-of-line blocking. However, since a number of cells may arrive at the same time for the same output link. the speed of the output buffers must be a multiple of the line rate.

1.5.3 Priorities and Multicast Switching

The input buffer and output buffer switches can be modified to handle multiple priorities For instance, each output buffer of an output buffer switch can be replicated for each priority level, and the transmitter then serves the cells in the corresponding priority order. The switches can also be designed to multicast cells, i.e., to copy a cell to different output links. In addition, the switches can be modified to copy one cell to a number of output links

We explain in Section 1.5 how to estimate the delay through an ATM switch. The upshot of the analysis is that the delay rarely exceeds a few cell transmission times provided that the sources race their transmissions suitably.

1.6 End-to-End Services

The main design objective of ATM is to provide a wide range of qualities of service. The ATM Forum specifies five service classes, These classes differ in the attributes of the quality of service and in the descriptors of the *traffic that* they carry. That is, a given service class carries traffic that must conform to some *traffic* descriptors. *The* service class then delivers the traffic with a quality of service that meets specific attributes. *We* explain the attributes of the quality of service and the traffic descriptors. We then describe the classes.

1.6.1 Quality of Service Attributes

The quality of service (QoS) attributes are the following:

•Cell loss ratio (CLR): fraction of cells lost during transmission.

•Cell delay variation (CDV) maximum difference between end-to-end cell delays.

•Maximum cell transfer delay *(max* CTD): maximum end-to-end cell delay.

•Mean cell transfer delay (mean CTD): average end-to-end cell delay.

•Minimum cell rate (MCR): minimum rate at which the network delivers cells,

Table 1.1 gives a few plausible attributes that representative applications might request.

1.6.2 Traffic Descriptors

The traffic descriptors are defined by an algorithm: the generalized cell rate algorithm (GCRA). The GCRA-***also called leaky bucket controls*** the cell transmission times. The figure shows two buffers with different transmission rates and buffer capacities. The cells that arrive are immediately duplicated and enter the two buffers, We say that the cell traffic that arrives conforms to the parameters (PCR, SCR, CDVT, BT) if the buffers never overflow. We describe these parameters below. The parameters are defined jointly for two different traffic types, which we discuss next,

TABLE 1.1 Representative QoS Attributes for Different Applications (Not Standardized)

	Attributes				
Applications	CLR	CDV	Max CTD	Mean CTD	MCR
Video conference	10^{-5}	50 ms	120 ms	100 ms	NA
Telephone	10^{-5}	2 ms	41 ms	40 ms	NA
Email	10^{-5}		Notspecified	NA	
File transfer	10^{-8}		Not specified	20 cells/s	

NA means not applicable.

Constant Bit Rate

*Constant-bit-rate (*CBR) traffic must not overflow the top buffer. PCR is the peak cell rate and CDVT is the cell delay variation tolerance of the traffic For instance, if CDVT =0, then the minimum time between two cell arrivals is 1/PCR. Indeed when the first cell arrives the buffer that can contain exactly one cell is full. That cell must be completely out of the buffer before the next cell arrives for the buffer not to overflow and the cell takes l/PCR to be transmitted. If CDVT * PCR = 0.1. then the minimum time between two cell arrivals is 0.9/PCR.

Essentially PCR is the maximum arrival rate of cells, and small fluctuations are possible when CDVT > 0. These fluctuations may be caused by the multiplexing or framing of the ATM cells by the physical layer. This is why the traffic descriptor allows for such fluctuations

Variable Bit Rate

Variable-bit-rate (VBR) traffic must not overflow either of the two buffers in Figure 1.12.The parameter SCR the sustained cell rate. is smaller than PCR SCR is an upper bound on the long-term arrival rate of cells. The parameter BT. burst tolerance, allows for the cells to arrive faster than at rate SCR for some time. The larger BT the larger this time.

Note that the source car use the traffic police shown to make sure that it does not send cells that do not conform to the descriptors Similarly, the network can use the policer verify that the source conforms to the descriptors. The ATM Forum specifies that the switch can set the CLP of nonconformant cells (that make one of the buffers overflow) and discard such cells as needed. We are now ready to define the service classes.

1.6.3 Service Classes

Here is the service class: CBR (constant-bit-rate service) transports CBR traffic with specified loss rare and delays VBR-RT (variable-bit-rate-real-time service) transports VBR traffic with specified loss rate and delays.

Table 1.2 Service Classes

Service class attributes	Definition: Traffic descriptors	Definition: Specified QoS
CBR	CBR	CLR, CDV, max CTD, mean CTD
VBR-RT	VBR	CLR, CDV, max CTD, mean CTD
VBR-NRT	VBR	CLR
ABR	NA	MCR
UBR	NA	None

NA means not applicable

- VBR-NRT (variable-bit-rate-non-real-time service) transports VBR traffic with specified loss rate.

- ABR (available-bit-rate service) delivers cells at a minimum rate; the understanding is that the network gets the cells from the source as fast as it can when it has spare capacity.

- UBR (unspecified-bit-rate service) is a best-effort service that attempts to deliver the cells without making any QoS commitment.

To provide CBR and VBR services, the network reserves resources for the connection when the service is setup. In ABR service, the network regulates the flow of cells by sending information to the sources indicating the rate at which they can send cells.

1.7 Internetworking with ATM

ATM is designed to complement Internet and LAN technologies. This complementarily is well illustrated by the use of ATM as a network to transport IP packets or LAN packets. We discuss these two important applications next.

1.7.1 I Power ATM

We want the computers to be able to run the TCP/IP-based applications as if the packets were delivered with TCP/IP In such a network, the ATM technology is used to take advantage of the flexibility in allocating quality of service to connections. Thus, a few of the computers may be using ATM in its “native mode”. For instance, a few computers may be set up for high-speed and low-latency data delivery for distributed computing applications. The other computers are set up to use their familiar TCP/IP applications. It is also possible to have some applications in one computer use IP while others use "native ATM”. We do not discuss here the pros and cons of this approach compared to an IP-only network. Let us simply remark that today IP is not capable of guaranteeing the quality of service that some applications might conceivably require and that ATM can provide that guarantee.

For simplicity. We look only at the TCP/IP computers. That is, we assume that all the computers are running IP over the ATM network. We explain the simplest scheme that can be used for transporting packets over ATM In this scheme. The network is divided logically into autonomous systems that are thought of as being interconnected by border gateways that is, the logical arrangement is identical to that of an IP network. The delivery of packets takes place as in an IP network, as we show next.

Consider that computers wants to send an IP packet to computer *D* These computers use the IP addresses S and *D* Computer's sends the packet to its router *R* Computer *R* consults its routing table and finds that the next hop is to router Q and that this router is at ATM address q Router *R* sets up an ATM virtual circuit to router Q and fragments the IP packet into ATM cells that it sends over the virtual circuit. When Q gets the cells, it reassembles the IP packet, looks at its destination address. finds that the next hop is to router *V,* sends the packet over to *V* by proceeding as *R* did. When it gets the packer, V sends it as an Ethernet packet to the destination *B.* In practice. the virtual circuits between the routers are set up permanently and the routers exchange information to maintain their routing tables.

The strategy we just described is called *classical (Rover ATW.* As you have noticed, a more efficient transfer can rake place by having *R* find out the ATM address of *V* and setup a virtual circuit directly from R to *V* Such a mechanism known as the *shortcut model-is* further away from the classical IP routing model but can be implemented.

1.7.2 LAN Emulation over ATM

We want to make it possible for all the computers to be able to assume that they are on a common Ethernet. That is, we want to make the ATM network transparent. In other words, we want to use the ATM network to emulate an Ethernet.

We describe a procedure that performs the emulation. The computers attached to the ATM network run software called LAN emulation (LANE). In addition, a computer Q attached to the LAN acts as the LANE server Consider a packet that computer S wants to send to computer *D*. Computer S assumes that *D* is on the same Ethernet and prepares a packet [s, d-data] where S and d are the Ethernet addresses of S and D. This packet is intercepted by the LANE in S. The LANE looks in a table to find the ATM address to which it should send the packet. If the table has no entry for d, the LANE asks the server Q for the ATM address that corresponds to ***d*** The server returns the address r; The LANE in S then sets up an ATM virtual circuit to r and sends the Ethernet packet *[s, d-data]* as a sequence of ***ATM*** cells, The router ***R*** reassembles the packet and sends it on the Ethernet ***E***. Permanent virtual circuits can be setup between ATM nodes, to reduce the processing and delay.

This LAN emulation strategy enables network managers to interconnect existing Ethernets with an ATM backbone, in addition, this strategy makes it possible to have these Ethernet coexist with native ATM connections. The need for a LANE server is viewed by network managers as a worrisome point of failure of the network. The arrangement can be made more reliable by duplicating this server role.

CHAPTER 2

Layered Protocols, the Architecture for ATM and SONET Networks

1.1 INTRODUCTION

This chapter introduces the concepts of layered protocols. The focus is on the open systems interconnection (OSI) model, because the architecture of ATM is based on OSI architecture. For this book, this model is examined in a general manner, with descriptions on the rationale for its use, and how it is used in networks. Other books treat the OSI model in more detail. Of course, in the spirit of the subject matter of this book, the OSI model is examined in relation to ATM.

The latter part of the chapter also provides an overview of several of the OSI and Internet protocols that operate at the network and transport layers in these two architectures. A knowledge of these system is important if the reader is to understand how ATM supports data communications systems.

2.1 PROTOCOLS AND THE OSI MODEL

Machines, such as computers and switches, communicate with each other through established conventions called protocols. Since computer systems provide many functions to users, more than one protocol is required to support these functions.

A convention is needed to define how different protocols of the systems interact with each other to support the end user. This convention is referred to by several names: network architecture, communications architecture, or computer-communications architecture. Whatever the term used, most systems are implemented with a set of protocols that are compatible, one hope, among the communicating machines.

The open systems interconnection (OSI) model was developed in the early 1980s by several standards organizations, principally led by the ITU-T and the OSI. It is now widely used for defining how communications protocols are standardized among different vendors' equipment. This model has provided a blue print for the design and implementation of computer-based networks.

The ITU-T publishes its OSI model specifications in the X.200-X.299 Recommendations. The X.200 documents contain slightly over 1100 pages. The ISO publishes its specifications for the OSI model in several documents, but does not use a numbering scheme that fits the mold of a simple "X.2xx" notation.

The model is organized into seven layers. Each layer contains several to many protocols that are invoked based on the specific needs of the user. However, the protocol entity in a layer need not be invoked, and the model provides a means for two users to negotiate the specific protocols needed for a session that takes place between them.

As suggested in figure 1-1, each layer is responsible for performing specific functions to support the end-user application. One should not think of a layer as monolithic code; rather each layer is divided into smaller operational entities. These entities are then invoked by the end user to obtain the services defined in the model.

An end user is permitted to negotiate services within layers in its own machine or layers at the remote machine. This capability allows, for example, a relatively low-function machine to indicate to a relatively high-capability machine that it may not support all the operational entities supported by the high-level machine. Consequently, the machines will still be able to communicate with each other, albeit at a lesser mode of service. Conversely, if two large-scale computers, each with the full OSI stack, wish to exchange traffic within a rich functional environment, they may do so by negotiating the desired services.

This concept holds true for networks as well. For example, OSI provides rules on how services can be negotiated between the user and the network. This concept is integral to ATM, because it uses this OSI idea of allowing the user to inform the network about its needs (such as maximum delay, minimum throughput, etc.). In turn, it allows the network to inform the user if it can meet these needs, or suggest a lesser quality of service (QOS).

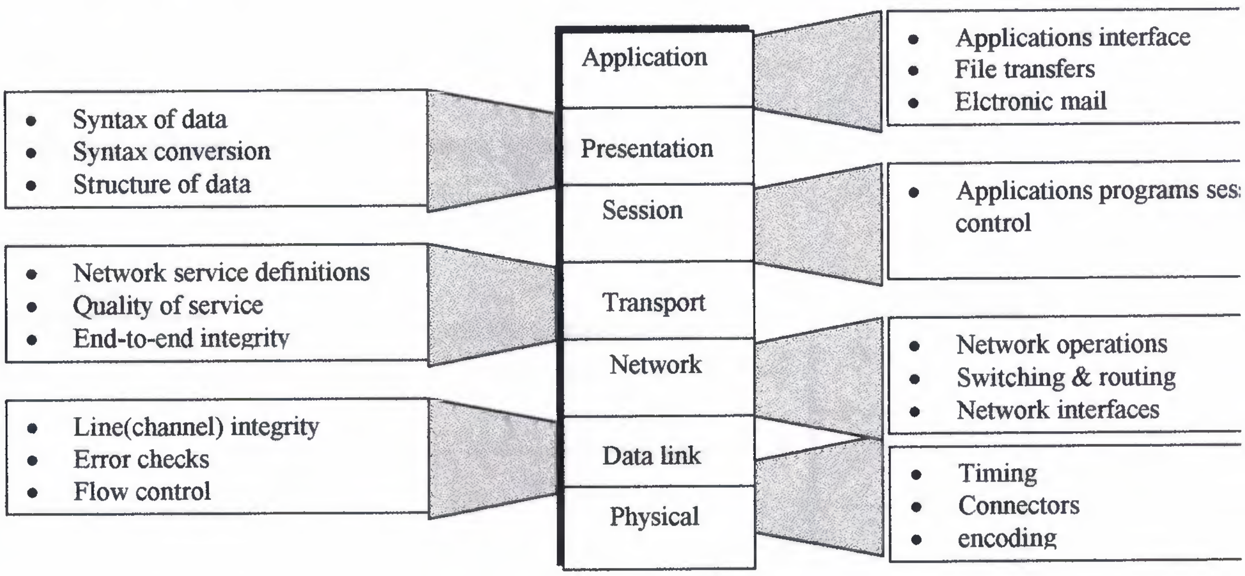

Figure 1-1
Functions of the OSI layers.

2.2 OSI Layer Operations

The layers of the OSI model and the layers of vendors' models (such as Apple's AppleTalk, IBM's SNA, etc.) contain communications functions at the lower three or four layers. From the OSI perspective, as demonstrated in figure 1-2, it is intended that the upper four layers reside in the host computers.

This statement does not mean to imply that the lower three layers reside only in the network. The hardware and software implemented in the lower three layers also exist at the host machine. End-to-end communications, however, occurs between the hosts by invoking the upper four layers, and between the hosts and the network by invoking the lower three layers. This concept is shown in figure 1-2 with arrows drawn between the layers in the hosts and the network. Additionally, the upper four layers may also reside in the network, for the network components to communicate with each other and obtain the services of these layers.

The end user rests on top (figuratively speaking) of the application layer. Therefore, the user obtains all the services of the seven layers of the OSI model.

Figure 1-2
Conventional position of host and network layers.

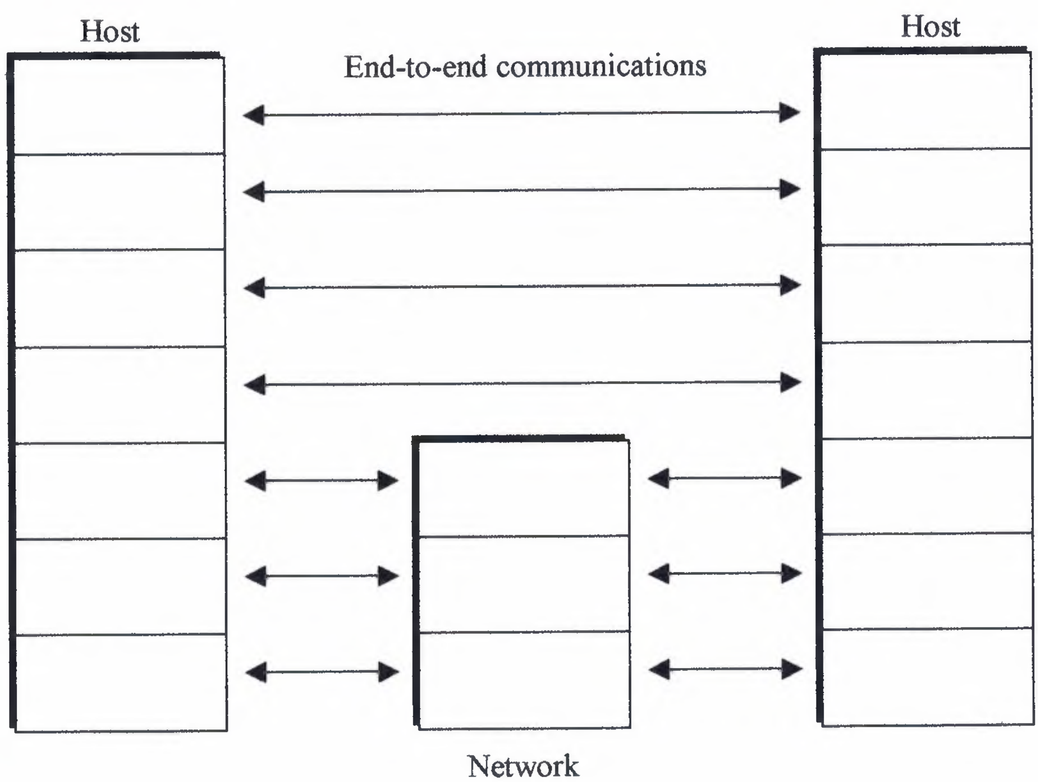

2.3 CONCEPT OF A SERVICE PROVIDER

In the OSI model, a layer is considered to be a service provider to the layer above it (Figure 1-3). This upper layer is considered to be a service user to its lower layer. The service user avails itself of the functions of the service provider by sending a transaction to the provider. This transaction (called a primitive) informs the provider as to the nature of the service to be provided (at least, requested). In so far as possible, the service provider does provide the service. It may also send a transaction to its user to inform it about what is going on.

At the other machine (B in this figure), the operation at A may manifest itself by the service provider B's accepting the traffic from service provider A, providing some type of service and informing user B about the operation. User B may be allowed to send a transaction back to service provider b, which may then forward traffic back to service provider A. In turn, service provider A may send a transaction to user A about the nature of the operations at site B.

Figure 1-3
The layer as a service provider.

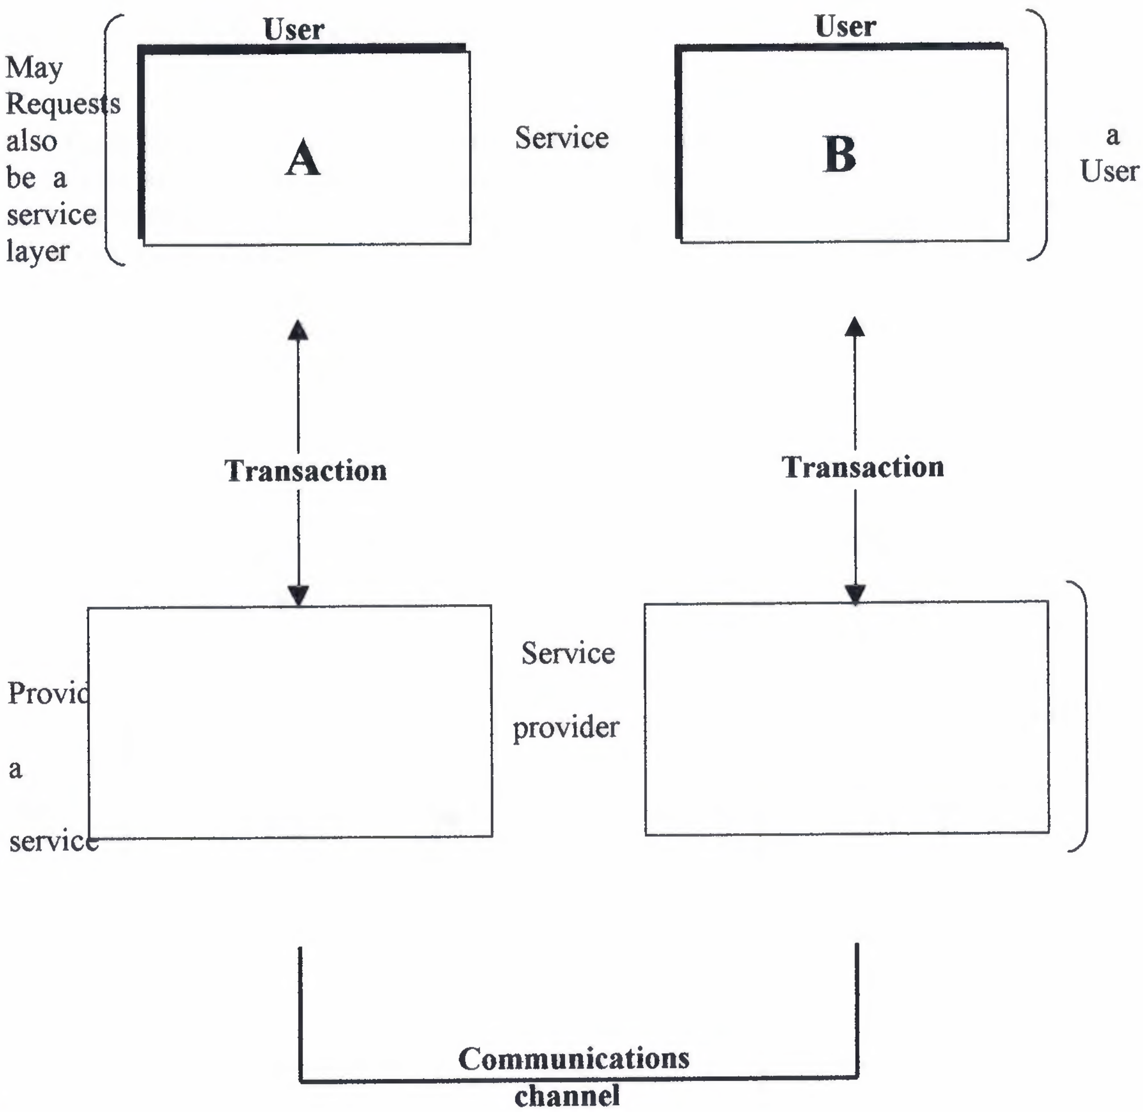

The OSI model provides several variations of this general scenario.

In accordance with the rules of the model, a layer cannot be bypassed. Even if an end-user does not wish to use the services of a particular layer, the user must still "pass through" the layer on the way to the next adjacent layer. This pass-through may only entail the invocation of a small set of code, but it still translates to overhead. However, every function in each layer need not be invoked. A minimum subset of functions may be all that is necessary to "conform" to the standard.

Layered network protocols allow interaction between functionally paired layers in different locations without affecting other layers. This concept aids in distributing the functions to the layers. In the majority of layered protocols, the data unit, such as a message or packet, passed from one layer to another is usually not altered, although the data unit contents may be examined and used to append additional data (trailers/headers) to the existing unit.

Each layer contains entities that exchange data and provide functions (horizontal communications) with peer entities at other computers. For example, in Figure 1-4, layer N in machine A communicates logically with layer N in machine B, and the N+1 layers in the two machines follow the same procedure. Entities in adjacent layers in the same computer interact through the common upper and lower boundaries (vertical communications) by passing parameters to define the interconnections.

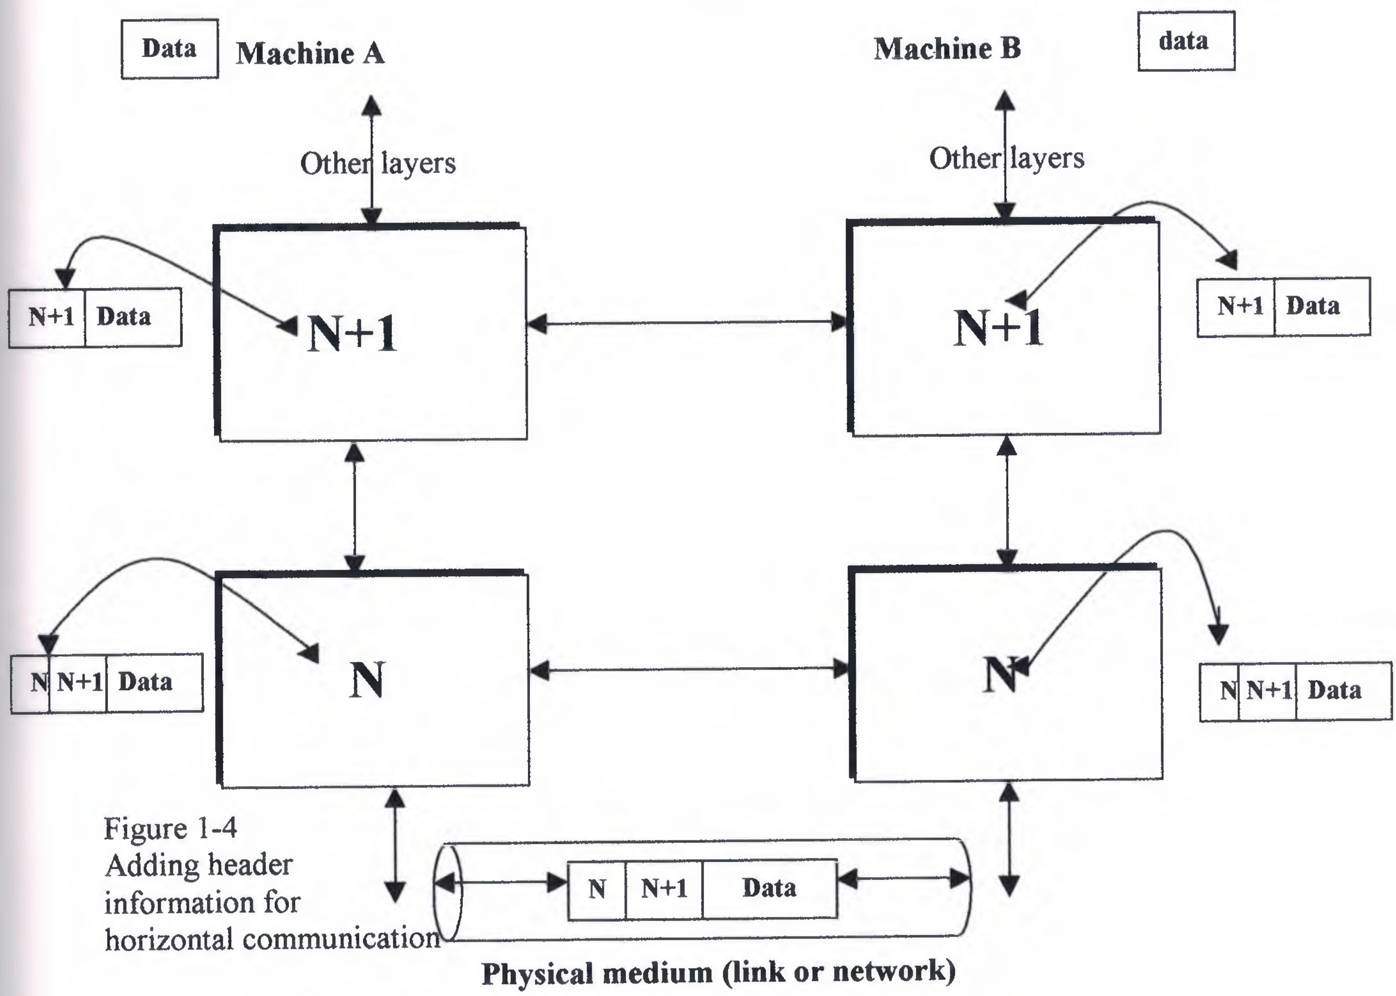

Figure 1-4
Adding header
information for
horizontal communication

Typically, each layer at a transmitting station (expect the lowest in most systems) adds "header" information to data. The headers are used to establish peer-to-peer session across nodes, and some layer implementations use headers to invoke functions and services at the N+1 or N adjacent layers.

The important point to understand is that, at the receiving site, the layer entities use the headers created by the peer entity at the transmitting site to implement predefined actions. For example, the ATM cell header created at the ATM layer at the sending machine is used by the receiving ATM layer to determine what actions it is to undertake.

Figure 1-5 shows an example of how machine A sends data to machine B. Data are passed from the upper layers or the user application to layer N+1. This layer adds a header to the data (labeled N+1 in the figure). Layer N+1also performs actions based on the information in the transaction that accompanied the data from the upper layer.

Figure 1-5
Machine A sends data to machine B.

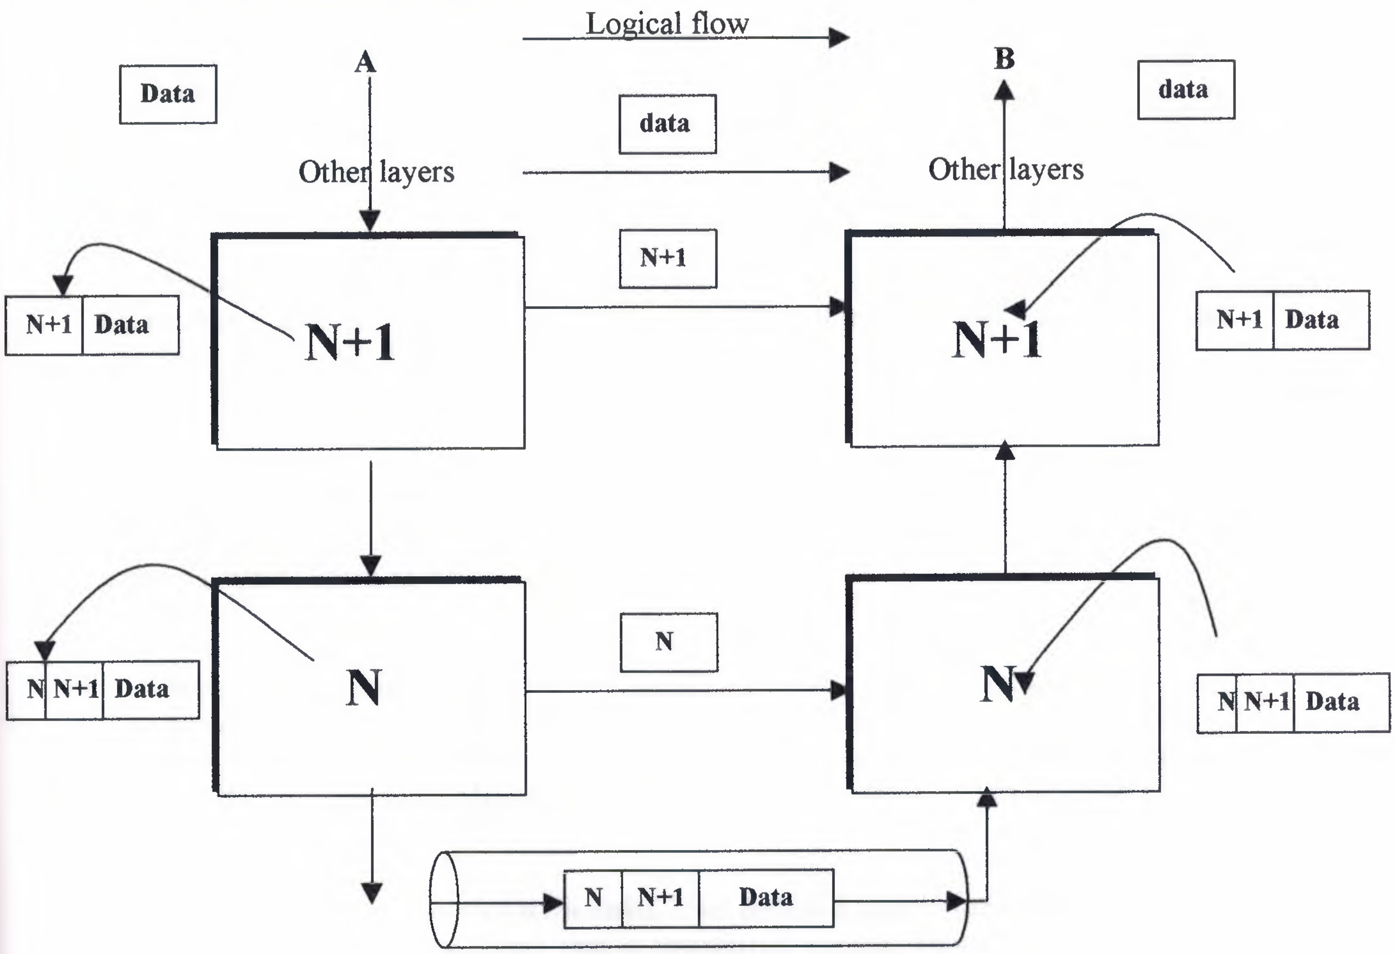

Physical Flow

Layer N+1 then passes the data unit with its N+1 header to layer N. Layer N performs the request actions, based on the information in the transaction , and adds its header N to N+1 traffic. This appended traffic is then passed across the communications line (or through a network) to the receiving machine B.

At B, the process is reversed. The headers that were created at A are used by the peer layers at B to determine that actions are to be taken. As the traffic is sent up the layers, the respective layer "removes" its header, performs defined actions, and passes the traffic on up to the next layer.

The user application at site B is presented only with user data which was created by the sending user application at site A. These user applications are unaware (one hopes) of the many operations in each OSI layer that were invoked to support the end-user data transfer.

The headers created and used at peer layers are not to be altered by any non-peer layer. As a general rule, the headers from one layer are treated as transparent "data" by any other layer.

There are some necessary exceptions to this rule. As examples, data may be altered by a nonpeer layer for the purposes of compression, encryption, or other forms of syntax changing. This type of operation is permissible, as long as the data are restored to the original syntax when presented to the receiving peer layer.

As an exception to the exception, the presentation layer may alter the syntax of the data permanently when the receiving application layer has requested the data in a different syntax (such as ASCII instead of BIT STRING).

Notwithstanding these exceptions, upper layer traffic and headers are usually sent transparently through a network. Another illustration should clarify the matter: Figure 1-6 depicts the relationship of the layers from the standpoint of how data are exchanged between them. Three terms shown in Figure 1-6(a) are important to this discussion.

- **SDU (service data unit).** User data and control information created at the upper layers that is transferred transparently through a primitive by layer (N+1) to layer (N) and subsequently to (N-1). The SDU identity (and as we just learned, its syntax) is preserved from one end of an (N) connection to the other.

- **PCI (protocol control information).** Information exchanged by peer (the same) entities at different sites on the network to instruct the peer entity to perform a service function. PCI is also called by the names headers and trailers (which are used in this book).

- **PDU (protocol data unit).** A combination of the SDU and PCI.

At the transmitting site, the PDU becomes larger as it passes (down) through each layer by adding that layer's header and/or trailer. At the receiving site, the PDU becomes smaller as it passes (up) through each layer by stripping away that layer's header and/or trailer.

Figure 1-6
Relationship of ATM to the layered operations.

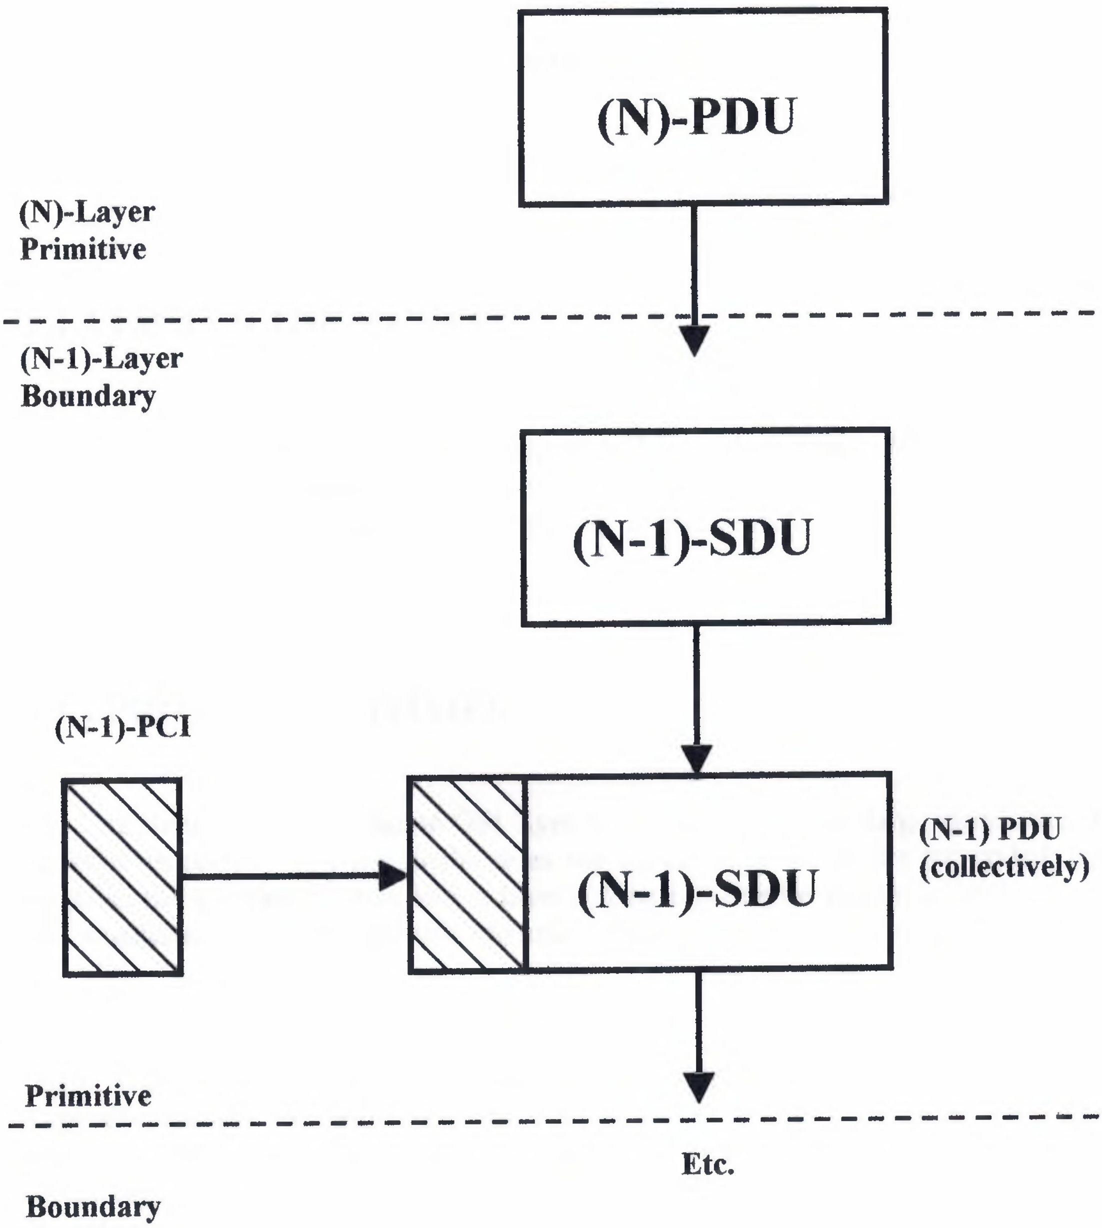

(a) OSI Process

Cell header (PCI)

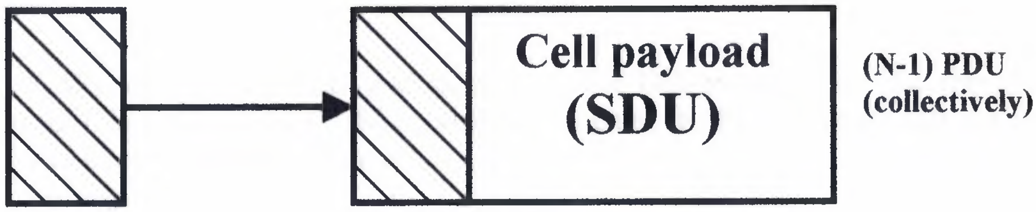

(b) Relationship to ATM

2.4 ATM AND THE MODEL

Figure 1-6(b) shows the relationship of ATM to OSI for their PCI, SDU, and PDU. ATM adds PCI to the cell payload. In OSI terms, the ATM cell header is PCI, and the cell payload is an SDU. Taken together, they form the PDU, which is called an ATM cell.

2.4.1 PROTOCOL ENTITIES

One should not think that an OSI layer is represented by one large monolithic block of software code. While the model does not dictate how the layers are coded, it does establish the architecture whereby a layer's functions can be structured and partitioned into smaller and more manageable modules. These modules are called entities.

The idea of the model is for peer entities in peer layers to communicate with each other. Entities may be active or inactive. An entity can be software or hardware. Typically, entities are functions or subroutines in a program. A user is able to tailor the universal OSI services by invoking selected entities through the parameters in the transactions passed to the service provider, although vendors vary on how the entities are actually designed and invoked.

2.4.2 SERVICE ACCESS POINTS (SAPs)

SAPs are OSI addresses and identifiers. The OSI model states: An (N+1)-entity requests (N)-services via an (N)-service access point (SAP), which permits the (N+1)-entity to interact with an (N)-entity.

Perhaps the best way to think of a SAP is that it is a software port (an identifier) that allows the two adjacent layers in the same machine to communicate with each other (Figure 1-7). SAPs may also be exchanged across machines in order to identify a process in the machine. In the OSI model, the SAP can identify a protocol entity that resides in a layer. For example, a SAP value could be reserved for E-mail, while a different SAP value could identify file server software. The reader may know about a UNIX socket-a concept similar to the OSI SAP.

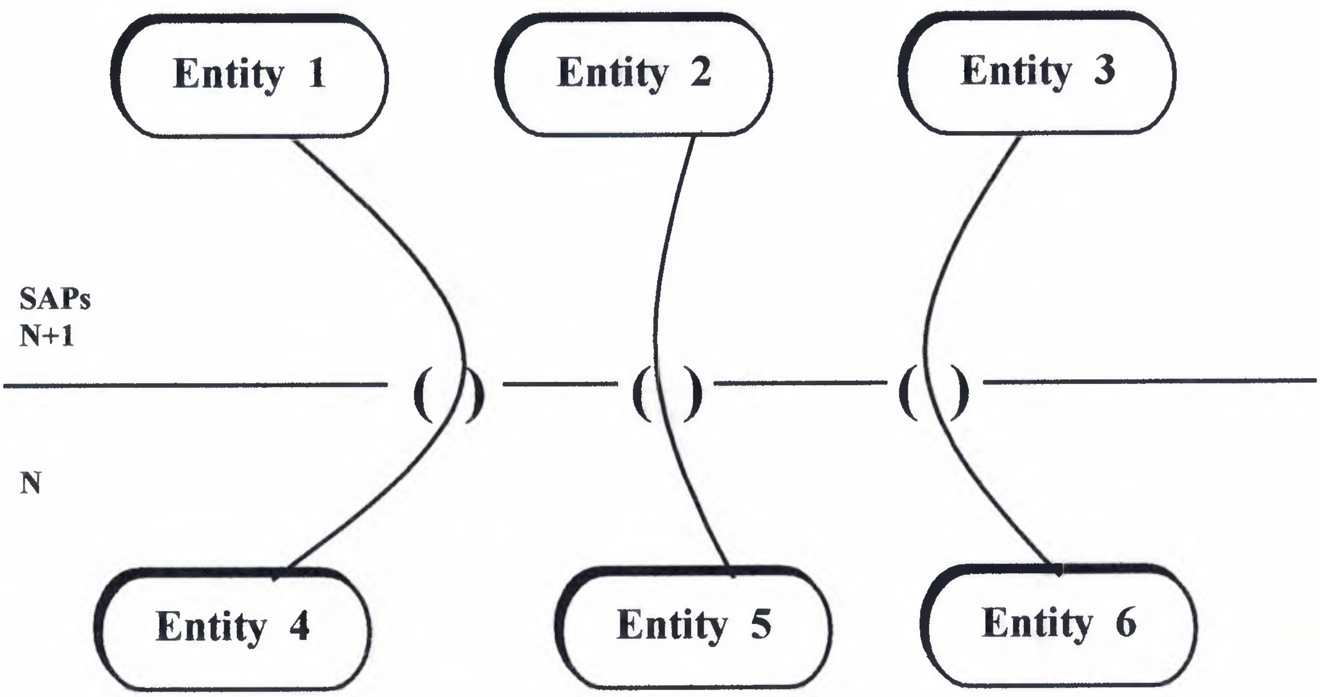

Figure 1-7
OSI service access points (SAPs).

2.5 ATM AND OSI LAYERS

Figure 1-8 introduces the ATM layers, and compares them to the OSI layered concept. The ATM operations reside in the ATM layer and the ATM adaptation layer (AAL). This chapter explains these layers in a general way. Subsequent chapters provide more detail.

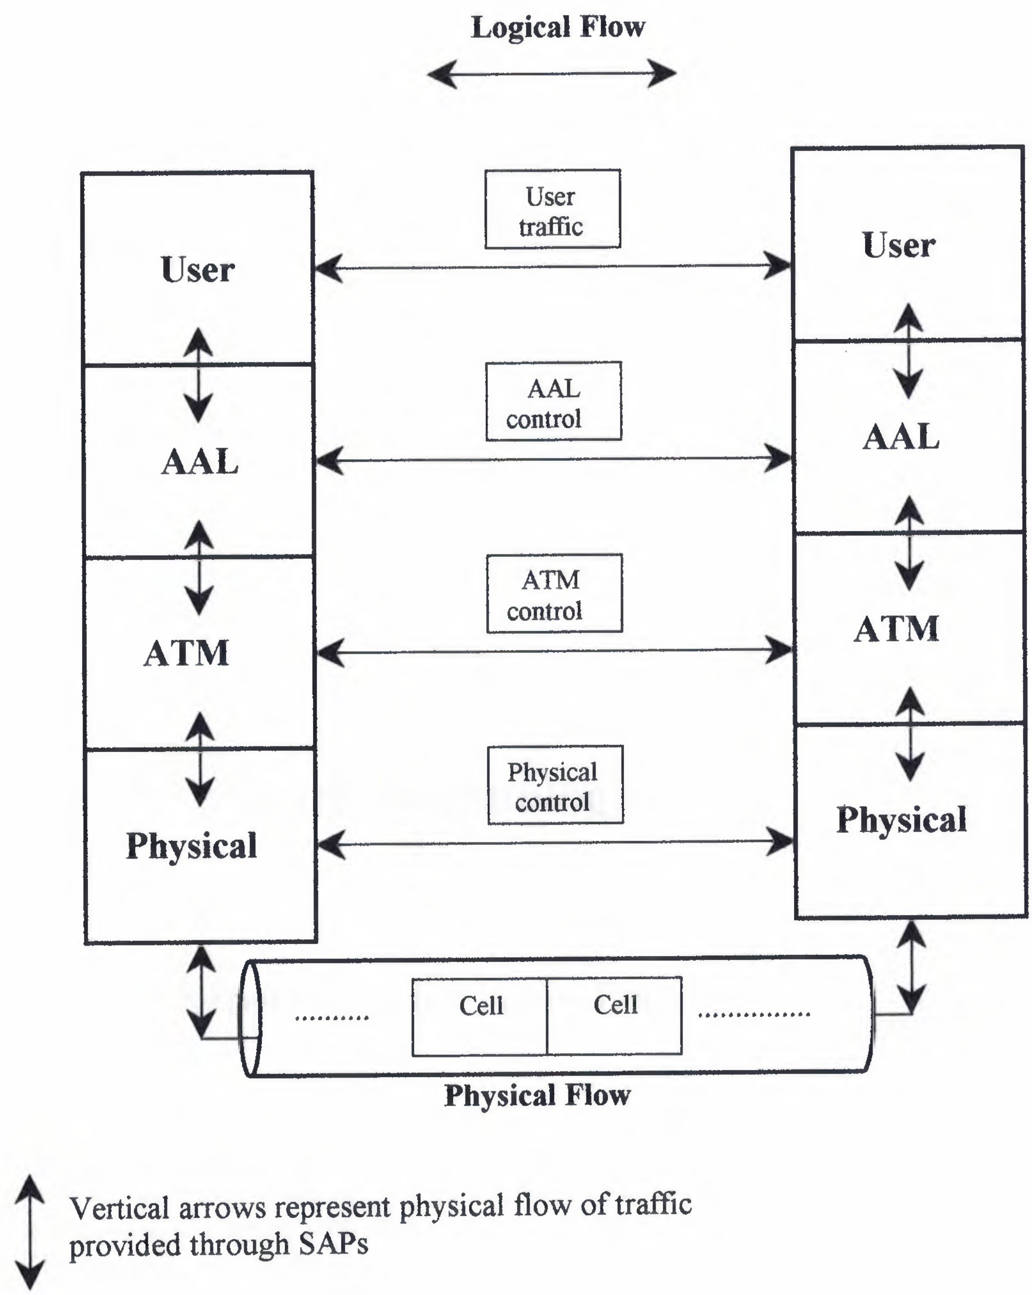

Figure 1-8
The relationship of ATM layers to the OSI model.

CHAPTER 3

Broadband Network Technologies

3.1 Broadband Network Technologies-On a Page

To be able to transport the MPEG compressed digital video signals through the network, from the encoder to the set top box for instance,a very powerful transmission mechanism is needed. Presently, the only technology that can fulfill all the requirements in terms of bandwidth, flexibility, and interactivity the digital video services have is the ATM protocol, the central element of the B-ISDN. This section will provide the background and the essential technical details of B-ISDN needed to enable a more detailed understanding of the key networking technologies, referred to in the later section dealing with "Video in Broadband Networks.' The following will be covered:

- The reasons for defining 'the B-ISDN-the history and background.
- The major interface points and network elements in the B-ISDN are defined in order to provide insight in the basic functional blocks and reference points.
- A detailed description of the organization of the different layers and protocols used in B-ISDN. The most cental aspects of these layers and protocols will then be dealt with in detail. This will cover the description of function and structure of:
- The physical layer and the transmission convergence layer, with special focus on some of the most central protocols, DS-3, E-3 PDH and, in particular. SONET OC-3/SDH STM-l.
- The ATM layer, including the ATM layer Quality of Service (QoS) aspects relevant for transporting digital video over ATM.
- The ATM adaptation layers with special focus on the protocols relevant for transporting MPEG-2 compressed digital video, AAL-1 and AAL-3, and finally the higher protocol layers with special focus on signaling, particularly the ITU protocol for UNI Signaling Q.2931.

3.2 History and Background of B-ISDN

The concept of B-ISDN was originally defined in the late 1980s. It is a collection of technologies with ATM as the "cornerstone," which is expected to form a universal network. The B-ISDN is essentially characterized by the ability to convey all present and future types of information, at very high speeds, and in a cost efficient manner. This is a contrast to the present situation, where a multitude of different net- works coexist to provide services of different kinds. The present structure, where the telephone network can transfer voice and data at low speeds, the X.23, or Frame Relay connection, which can transfer data at medium speeds, and broadcasting or cable TV, which can convey television, presents an inefficient use of resources. This is due to the fact that each and all of these networks must be installed and maintained separately Furthermore, when resources become available on one network, for instance, on the telephone network in the time interval after the 9-3 business peak hours, these resources cannot be made available for other networks or services, such as television.

An effort to cope with these drawbacks of the existing telecommunication infrastructure was made with the development and implementation of the Narrowband ISDN (N-ISDN). This network can transport, voice, data, and, in principle, video. The N-ISDN has some limitations, however.

- ISDN has a fairly limited bandwidth, namely 144 Kbps with the basic rate interface, and 2 Mbps with the primary rate interface. This is probably suitable for most uses presently, but within a short time normal bandwidth requirements are Likely to have risen beyond these rates. For example, a typical broadcast quality service carrying MPEG-2 compressed video; normally require at least 4 Mbps.

- N-ISDN is inflexible in many aspects. The user must for instance pay for the maximum bandwidth of the connection during a call, regardless of how much of the bandwidth capacity he actually uses.

- Only relatively few services supports the N-ISDN concept.

The next step in this evolution is B-ISDN. It has been the subject of much work, since the standardization body ITU in 1988 declared that ATM should be the basis of the future B-ISDN. Field trials have been done, and are still being done, all over the world in order to give telecom operators and equipment manufactures practical experience with B-ISDN. Presently, field trials involving B-ISDN networks with MPEG-2 compressed digital video services are done. The most important parts of B-ISDN, such as like ATM and SDH, have been finalized. Some standards may, however, still need to be developed further and refined, in order to enable a fully working B-ISDN that can cope with the multitude of potential new services.

3.3 Abilities and Benefits of B-ISDN

3.3.1 Introduction

There are a number of advantages of B-ISDN compared with the existing network structures. Some of the most important are the following six:

- Application independence

- Bandwidth efficiency

- LAN-MAN-WAN integration

- Bandwidth granularity
- Dynamic bandwidth
- Variable connection quality

It is mainly due to these qualities that ATM and B-ISDN have received so much attention in the Tele/datacommunication business over the last years, and now also in the video business. The details of these six overall advantage areas will here be dealt with in more detail.

3.3.2 Application Independence

With the B-ISDN only one network is needed to cover all the different possible services. It can convey information of different kinds, with different characteristics, namely;

- Data: Constant or variable bit rate data. Data transfer, such as LAN connections, often comes with bursts of traffic in short periods-typically in the range of milliseconds. Data transfer is normally not sensitive to delay in the network, unless real-time applications are involved. (A small delay in a file transfer is of no importance, for instance.) The band- width requirements of a data transfer is normally from 64 Kbps to 10 Mbps, but in principle there is no limit. Presently, 133 Mbps and even 622 Mbps connections to the desk top are being developed and commercially deployed

- Video: Video or television transmission, including every- thing from the relatively low VHS quality (or lower), all the way over HDTV quality to studio quality, is sensitive to delay, and especially to delay variation. It has constant (or in some cases varying) bandwidth demand of typically 180-270 Mbps in uncompressed format (e.g., MPEG-2), the bandwidth requirements typically lie in the range from 1.3-80 Mbps. The burst ness of variable bite rate video is normally lower than in the case of, let's say, data transfer connections.

- Voice: A constant 64 Kbps bandwidth requirement characterizes voice transfer from telephones, as we know it today. The transfer of voice is highly delay sensitive, since even small delays or interrupts are perceived to be annoying by the human ear. Longer delays can make communication virtually impossible.

- Combinations of the above mentioned or multimedia applications, which could, for instance, be interactive TV, such as tele-shopping, tele-education, or transmission of other multimedia services, including text, images, sound, video, and possibly other types of information.

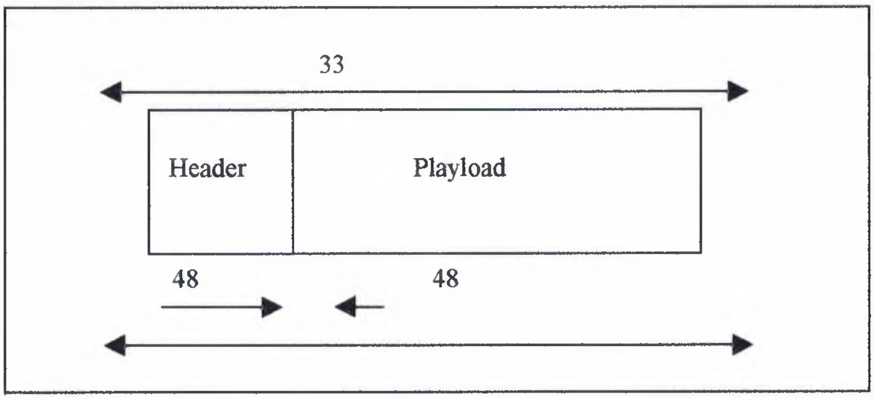

Figure 3.1 Structure of the ATM cell

The fact that one network can convey all these different types of information, with different demands to bandwidth, burst ness, etc., is primarily due to ATM technology It conveys information in small "cells" or packets of 33 bytes, 48 bytes for user information and 3 header bytes for different control purposes, as shown in Fig. 3.1. The information coming from the different services, such as file transfers or video transmissions, are segmented to fit into the ATM cell by the so called ATM adaptation layer (explained in detail later), and multiplexed into one stream of ATM cells.

The information flows from (or to) different users, for instance a specific type of video service can be collected in the network and transmitted via specific channels or paths, which can provide the required performance. The information of what channels or paths the information should be led via, is located in the 3 byte header of each ATM cell.

The fact that the ATM cell is of a fixed size instead of a varying size, enables simpler and much faster processing in all network components, such as interface cards and switches, as processing of the cells can be done in HW rather than SW. Furthermore, it is not necessary to calculate packet length, allocate varying buffer space, etc., as variable length packets require. Finally, the ATM cell is only 33 bytes long. Hereby delay sensitive applications will not be delayed significantly from the time it takes to "fill" a cell with information.

3.3.3 Bandwidth Efficiency

In contrast to "synchronous transfer mode" or "time division multiplexing," (TDM) as it is normally called, ATM uses the bandwidth of a connection in a fairly efficient way. TDM is found in a great deal of the existing communications network, for instance, in the existing Tele- phone network, utilizing the Plesiochronous Digital Hierarchy (PDH) technology. With TDM, the bandwidth of a given connection is shared in a fixed way among the users. Each user has a "time slot," which is available to him or her, whether it is needed or not. See Fig. 3.2. If, for instance, one user is inactive, this causes a 33 percent capacity waste.

TDM guarantees the required bandwidth with an acceptable, constant delay. But it is inefficient in its use of the transmission capacity. In ATM the access to the network is in principle unlimited, dependent only on the capacity available. This is also referred to as "statistical multiplexing."

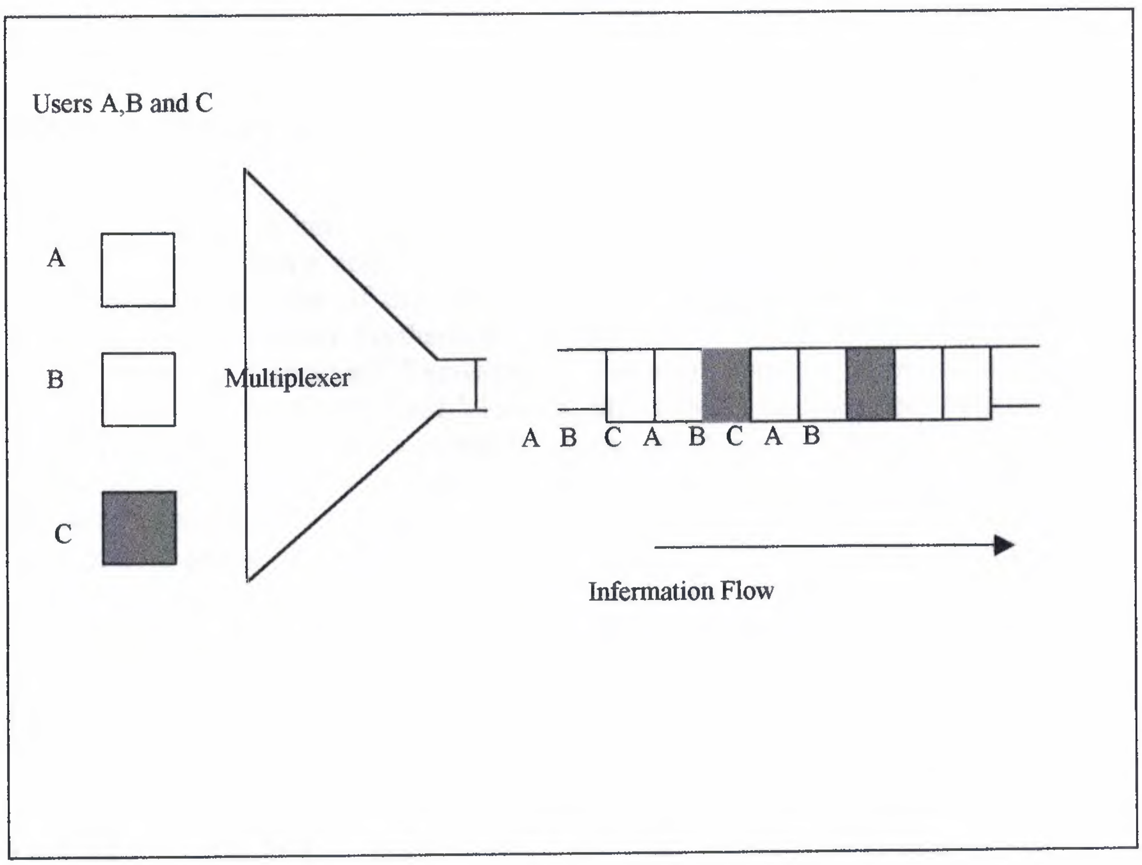

3.3.4 LAN-MAN-WAN Integration

B-ISDN is well suited for use in LANs, MANs, and WANs. Combined with the fact that ATM is independent of the physical layer protocols and cables, ATM cells can be transported on all the networks. This drastically reduces the need for protocol conversion, as we presently when we transport LAN traffic over a WAN connection. Furthermore, the independence of the physical layer means that the installed base of wiring can be re-used.

3.3.5 Bandwidth Granularity

Today, users and network element designers have to respect the discrete "blocks" in which bandwidth is available. These blocks are normally based on multiples of 64 Kbps, such as 2, 8, or 34 Mbps. With B-ISDN, it is the other way around. Here, the network can be tailored to closely fit to the bandwidth demands of the application, since the minimum unit of information is the 48 bytes of payload in the ATM cell. In practice, the bandwidth requirements, among other parameters, are specified at the time when the connection is established.

3.3.6 Dynamic Bandwidth

When a connection is initiated, the bandwidth requirements of the connection; in terms of traffic rate and burst ness, are "negotiated" with the network via the signaling procedures, as described later. If the demands change during the call, the connection parameters can be renegotiated. Furthermore, the user can be billed for the actual band-width he or she uses, if required. Expressed in a popular manner, he or she can be charged, "for the number of cells used." Finally, the B-ISDN has no upper limit with respect to bandwidth. This depends primarily on the capabilities of the physical layer components at a given time. Presently, the bandwidth available to the end user is typically lying in the range of T1/E1 (1.3 or 2 Mbps), T3/E3 (43 or 34 Mbps), or the ATM Forum 23.6 Mbps. For the core network; SDH STM-1 or SONET OC-3 (133 Mbps) is used, as well as SDH STM-4 or SONET OC-12 (622 Mbps).

3.3.7 Variable Connection Quality

It is possible for the user to specify what "Quality of service" (QoS) he wants from the network. In the B-ISDN, if, for instance, a moderate level of cell delay variation is a minor problem for the service used, this can be specified when the connection is initiated, and possibly renegotiated later as desired. In this way, it is possible for the user only to pay for what he needs, and chances are that the network is used more efficiently.

CHAPTER 4

Traffic management in ATM

INTRODUCTION

The ATM network must be able to adapt to unforeseen traffic patterns, for example, unusual bursts of traffic from the various end-user devices or applications. Also, the network must be able to shed traffic in certain conditions to prevent or read to congestion. In so doing, it ml's be able to enforce an allowable peak cell rate for each VPI/VCI connection, This means that when a user's traffic load is presented to the net work beyond a maximum peak rate, then it nay be discarded.

In addition to these responsibilities, the network must be able t monitor (police) traffic it must be able to monitor all VPI/VCI connections, and verify their correctness (that they are properly mapped into the network and operate effectively). The network must be able to detect problems, and emit alarms when certain troubling events are encountered.

An effective ATM network must be designed with an understanding that both the user and the network assume responsibility for certain QOS operations between them The user is responsible for agreeing to service contract with the ATM network that stipulates rules on the us' of the network, such as the amount of traffic that can be submitted in measured time period In turn, the network assumes the responsibility of supporting the user QOS requirements.

4.1 The Natural Bit Rate

All applications exhibit a natural bit rate [dePr91]the rate a which the application generates and/or receives a certain number of bit per second (bit/s) based on its "natural requirements". For example, digitized voice using conventional ITU-T standards exhibits a natural in rate of 32 kbit/s. Large data transfer systems have a fluctuating nature bit rate, depending on the nature of the traffic, but ranging from a few kbids to several hundred kbits/s. High definition television (HOW) ha a natural bit rate of approximately 100 to 150 kbit/s, depending on the coding schemes employed on the signals.

The challenge of the ATM network is to support the natural bi rates of all applications being serviced. Because of variable traffic profiles, it may be imperative to discard traffic from certain users if the network experiences congestion problems. This situation is illustrated in Figure 1-1(a). For a brief period of time the user exceeds the transfer rate permitted in the network with a burst of traffic. The shadowed Portion of the graph is the "burstyness" of the traffic, which exceeds, momentarily, the permitted rate of the network With ATM networks, the traffic will be tagged for possible discard

FİGURE 1-1

Bandwidth allocations. [Depr91]

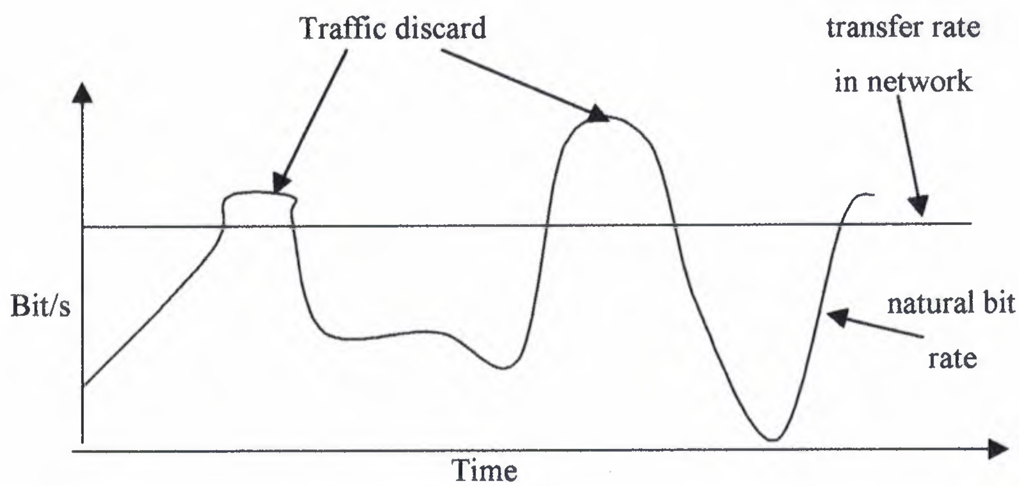

(a) User exceeds transfer rate

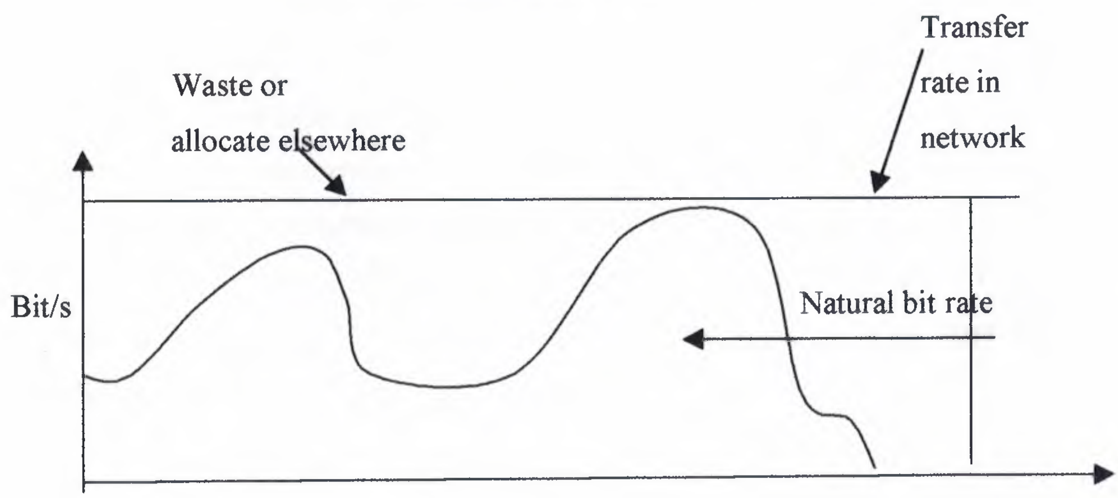

Time

(a) Transfer rate greater than user bit rate

An application's natural bit rate (because of its burstyness) may not use the full rate allocated by the network to this application. The possibility of wasted bandwidth occurring is shown in the striped area of Figure1-1 (b). Therefore, some systems (such as an ISDN terminal adapter) can stuff bits into the application's data stream such that it appears to be the same as the transfer rate in the network (or more likely the transfer rate of the communications link on which the traffic is to be sent). Or, the network can attempt to allocate to other applications. With these discussions in mind, the remainder of this chapter examines some approaches to the management of multiplication traffic.

4.2 TRAFFIC CONTROL AND CONGESTIOM CONTROL

In a B-ISDN, the terms traffic control and congestion control describe different aspects of ATM operations. We concern ourselves with the latter term first. Congestion is defined as a condition that exists at the ATM layer In the network elements (NEs) such as switches, trans-mission links, or cross connects where the network is not able to meet a stated and negotiated performance objective. In contrast, traffic control defines a set of actions taken by the network to avoid congestion; traffic control takes measures to adapt to unpredictable fluctuations In traffic flows and other problems within the network, The objectives of both traffic control and congestion control are to protect the network and at tile same time provide the user with its stated service contract objectives. For B-ISDN, this includes formally stated QOS objectives. ATM is not designed to rely on AAL to provide any type of traffic control or congestion control measures. While AAL may indeed perform these functions the design of the ATM network does not assume this service.

4.3 Functions to Achieve Traffic Control and Congestion Control

To meet the objectives of traffic control and congestion control, the ATM network must:

- Perform a set of actions called connection admission control (GAC) during a call setup to determine if a user connection will be accepted or rejected. These actions may include acquiring routes for the connection.

- Establish controls to monitor and regulate traffic at the UN'; these actions are called usage parameter control (UPC).
- Accept user input to establish priorities for different types of traffic, through the use of the cell loss priority (CLP) bit.
- Establish traffic shaping mechanisms to obtain a stated goal fat managing all traffic (with differing characteristics) at the UNI.

These concepts are described in the following sections.

4.4 ALLOCATION OF BANDWIDTH

Since an ATM network is expected to support a wide variety of applications, the network designer must answer the questions of how the natural bit rate of user applications can he accommodated vis-a-vis the transfer rate of the network. Or stated another way, How should the CAC and UPC be established and management effectively for each user connection? The term “effectively” has two aspects;

1. What is effective for an individual user may not be effective from the network perspective.
2. And, of course the opposite is true-what is effective for the network may not be perceived as effective for the user.

The ATM standards do not stipulate all the specific rules for the CAC and UPC operations. Notwithstanding, this section of the chapter summarizes several approaches that can lead to a fair and effective allocation or the network's bandwidth.

The network exhibits a finite transfer rate, as does each link in the network. Not only must ATM traffic management manage the limited bandwidth within the network, but it also must allocate traffic to *each* communications link within the network.

The examine the second issue, Figure 9-2 depicts a typical traffic stream that a ATM node might receive (assuming the machine is truly a multiplication node). Regardless of the capacity of the network, the aggregate transfer rates of all the input lines n (inlets) must not exceed the t' transfer rate of the output line m (outlet at least not for a prolonged period (in fractions of milliseconds). Notwithstanding, some buffering of traffic at the multiplexer (ATM node) allows n to exceed m for a few milliseconds

Each queue must be serviced by the multiplexer in a fair and equitable manner. Queue servicing operations should result in appropriate delays and acceptable traffic losses vis-a-vis the application As examples, (1) CBR video queues should be serviced between 1 and 2 ms; (2) VBR voice queues should not allow the loss of traffic to exceed 1 to 10 percent of the samples (the range depends on how traffic is dropped, see Chapter 7); (3) data should allow for no loss of traffic; and (4) the queue contains OAM information and should receive the highest consideration of any of the other queues.

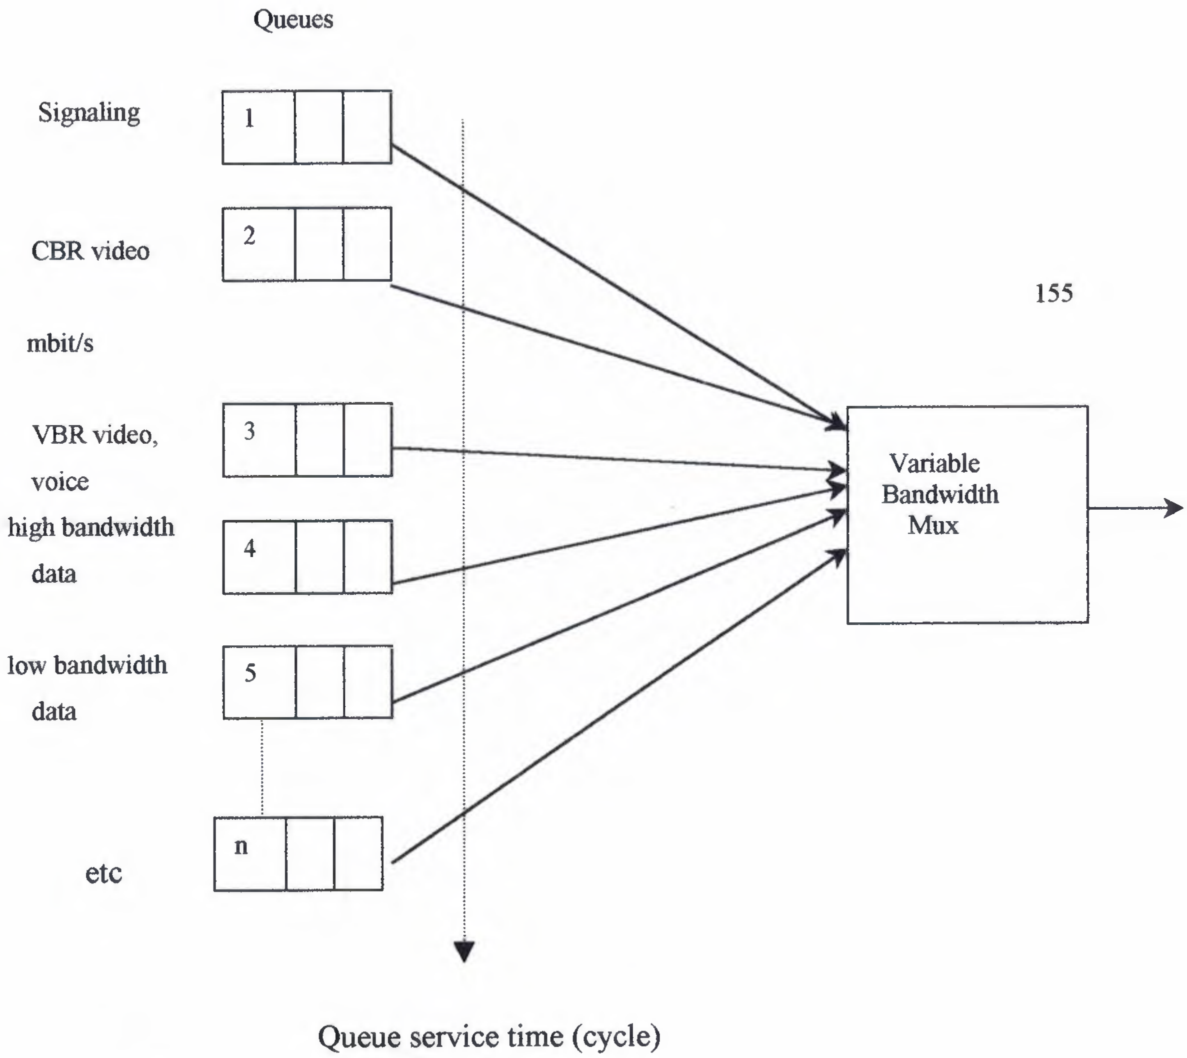

Figure 1-2

Allocating bandwidth using multiple queues

As figure 1-2 shows, the queues are serviced on a cyclical basis (the queue service cycle). Each queue is examined, payload is extracted and transported through the 155 Mbit/s output link. The next queue is examined pay1oad is extracted, and so on, until the last queue is examined and process starts over again on the first queue.

The manner in which the queues are set up is not defined in the ATM specifications. Typically a separate queue is established for each voice and video connection. This approach simplifies queue management.

The queue service cycle must assure that delay-sensitive traffic is serviced within 1 to 2 ms. Given this requirement, let us assume a cycle time of 1.5 ma, which translates to 666.6 service cycles per second (1 set ond/. 0015 = 666.6). Further, assume that the 155.52 output link can accept 353,207 cells per second (155,520,000 [less overhead of the 155.52 Mbit/s frame yields a rate of 149.760 Mbit/s]/424 bits in a 53-octet cell 353,207). Therefore, this configuration can service 529 cells per cycle (85:3,207/666.6 = 529).

Figure 1-3 shows that all queues are serviced during the service cycle. The multiplexer must adjust to the changes in the number of calls and the resultant queues, and vary the service times on the queues accordingly. Therefore, it must be able to add/delete queues and adjust its queue extractions accordingly

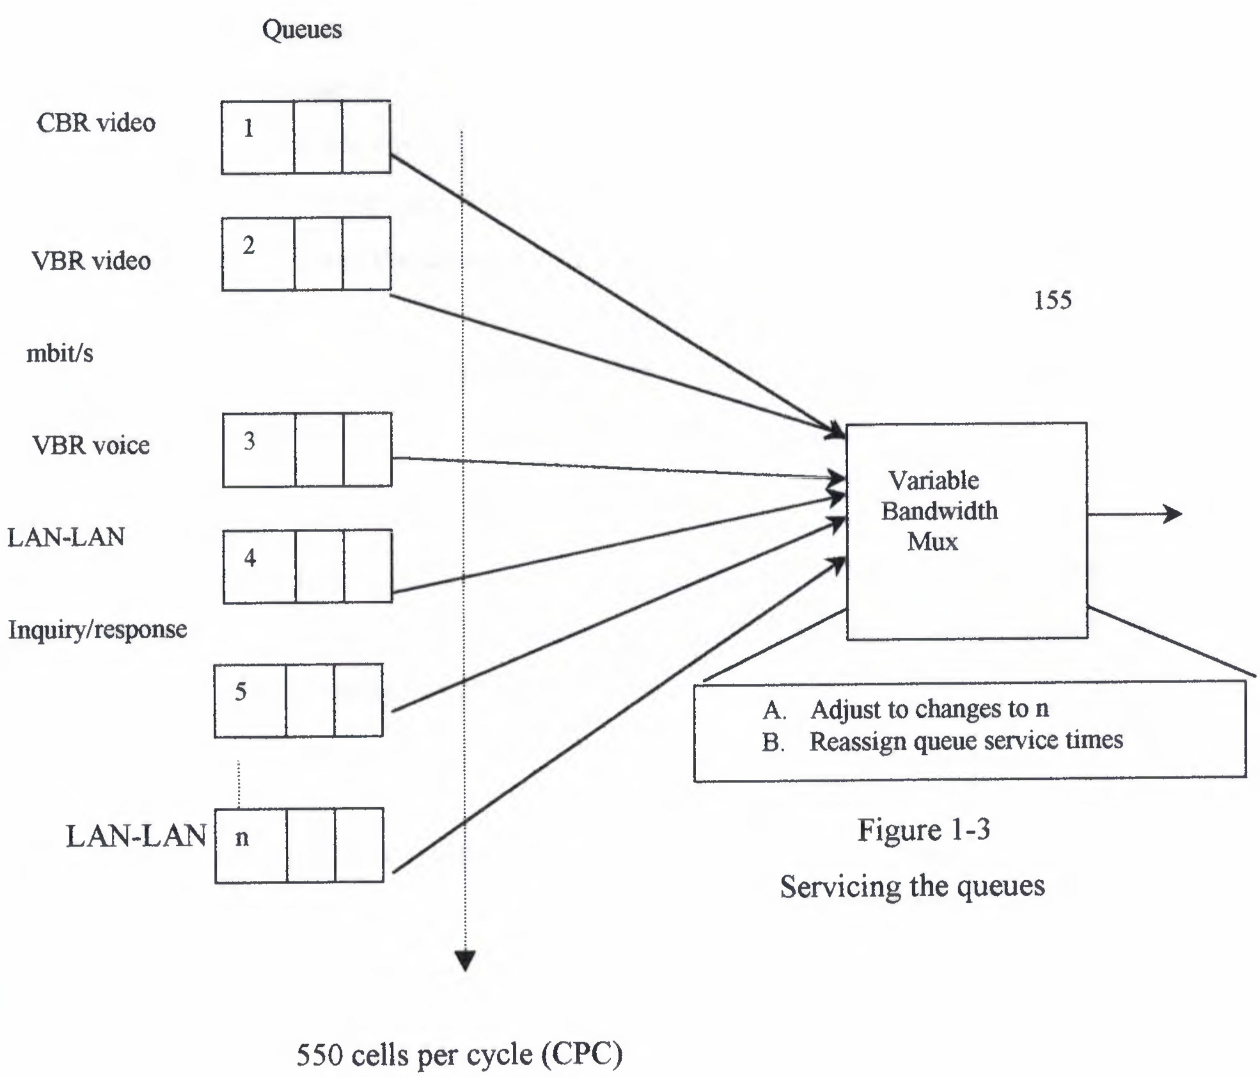

Figure 1-3

Servicing the queues

The example in Figure 1-2 is based on the use of a SONET STS-3c 155.52 Mbit/s link. Other links with different bit rates obviously affect the value of the cells per service cycle. For example, if the output link were a 44.736 DS3 link, the cells could use 40.704 Mbit/s of this capacity-the other bits are overhead. Therefore, *40,704,0001424* = 86,914 for the number of cells that can be transported per second across the DS3 link. This configuration supports 130 cells per service cycle (86,914/666.6 = 180) The task now is to determine which queues are to have their cells withdrawn during the service cycle, and at what rate the cells will he withdrawn from the queues. The general strategy is simple [SRIR9Oaj and [SlnR9ob]:

- The signaling queue is given the highest priority, and should experience little or no delay and experience no loss.
- Delay-sensitive queues are serviced next, if cells are in the queues, for T_1 ms, or until these queues are empty
- Next, delay-insensitive queues are serviced for T_2 ins.
- In some installations, the absence of talk spurts allows data queues to be serviced more frequently
- When the signaling queue is serviced, either T1 or T2 is suspended, and resumed (not restarted) when this queue service is finished.

Thus, all queues are guaranteed an established and minimum bandwidth of (T1/{T1+T2}) B for delay sensitive traffic and (T2/{T1+T2}) B delay insensitive traffic, where B is the transfer rate of the output link.

The manner in which the queues are serviced depends on how many cells must be withdrawn per service cycle. Be aware that the number cells serviced per cycle varies, not only between voice, video, and data but within those applications as well. For example, a conventional 64 kbit/s voice call operates at about 166 cells per second. In marked contrast, a 32 kbit/s voice call using selective cell discard and compression of silent periods needs only about 38 cells per second. Therefore, the variable bandwidth mux must know how to service each queue in regard to the cell withdrawal rate. Once again, these decisions are vendor specific.

4.5 DEALING WITH VARIABLE DELAY

Traffic management operations occur not only at the source UNI but also at the destination node itself The manner in which traffic management is implemented depends on the type of traffic Assuming that the traffic has been granted admission at the source UNI into the network upon its arrival at the end-user device, further traffic management decisions on how to "outplay" the traffic to the user application must be made.

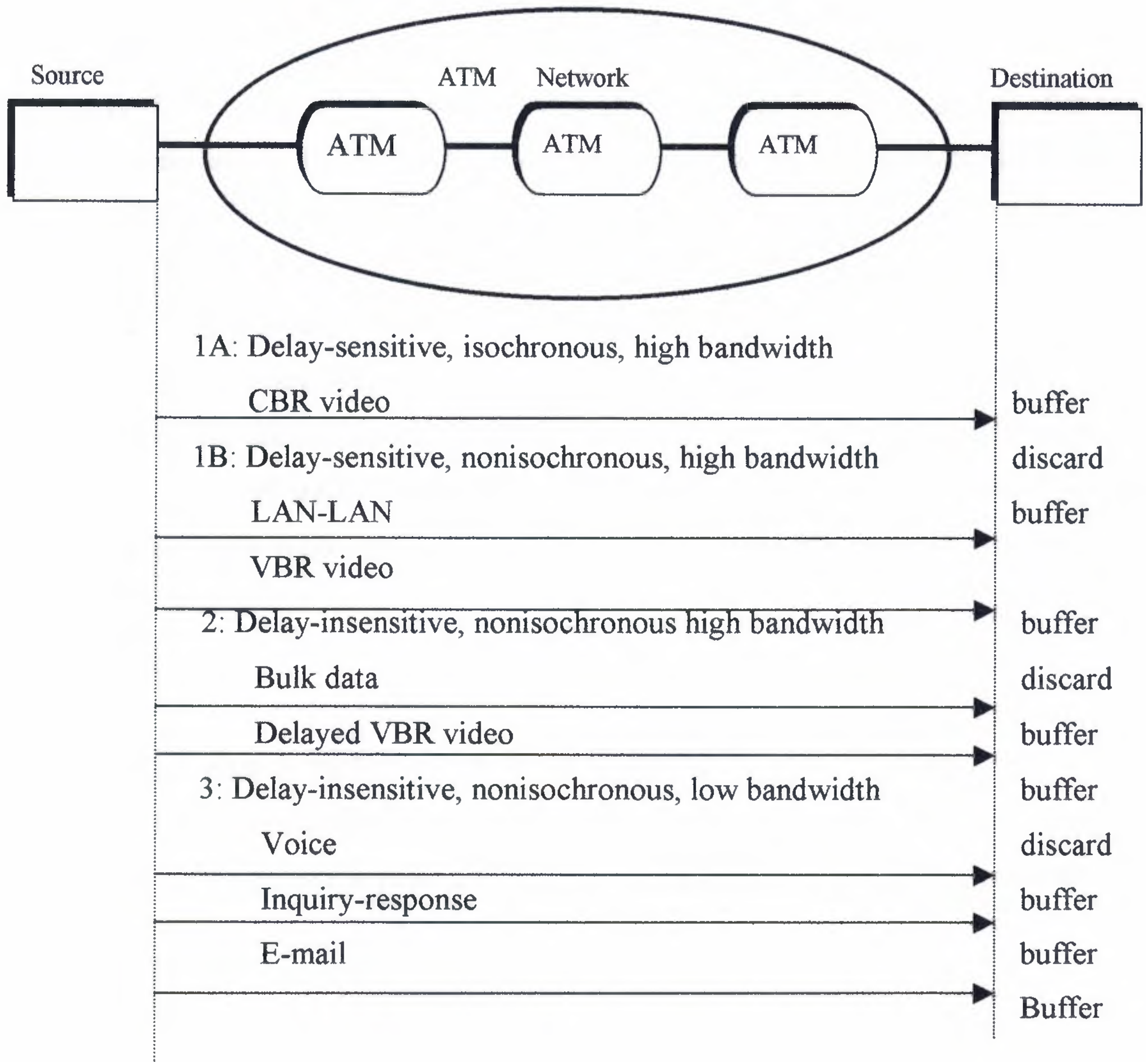

Figure 1-4

Processing the payload at the receiver.[SRIR93]

As Figure 1-4 shows, different types of traffic must be handled in different ways. The classes of traffic, 1A, 1B, 2, and 3, represent. One method of categorization. The traffic with high bandwidth requirement using isochronous timing is buffered at the receiver if the traffic arrives sooner than a predetermined time. If it arrives after this time, it is discarded. Type lB traffic, which is delay-insensitive nonisochronous, high--bandwidth traffic, is handled depending on the specific kind of traffic For example, LAN-to-LAN traffic would be buffered, whereas variable hi-rate video would be buffered if the traffic arrives early and discarded if the traffic arrives late. Type 2 traffic, which is delay-insensitive, nonisochronous, high-bandwidth type traffic, again is handled differently depending of the specific subtype. Data traffic is buffered and delayed VER video traffic (video that must be delivered at a later time) is handled like any type or VBR video traffic that, if it is early, is buffered, but if it arrives late is discarded.

Finally, class 3 traffic, which is delay insensitive not isochronous, is handled differently for voice than for data. As the reader might now expect, early arriving voice packets are buffered until a payout time , is reached, and then late arriving packets are discarded.

In essence, any type of traffic involving data is buffered and held in buffer for quite sometime to prevent discarding it. In contrast, video and voice will be discarded if it arrives too late and it will be buffered (to nice a standard arrival time) if it arrives too early.

The reader may wish to review "Voice Packetization" in chapter 7 for examples of queue management and outplay operations.

4.6 USAGE PARAMETER CONTROL (UPC)

After a connection is granted, and the network has reserved n sources for the connection, each user's session 'S monitored by the network. This is the UPC operation, which is designed to monitor and control traffic, and to check on the validity of the traffic entering the network. UPC maintains the integrity of the network and makes sure that only valid VIPs and VCIs are entering the network

According to the ATM Forum, several other features are desirable for UPC:

- The ability to detect noncompliant traffic
- The ability to vary the parameters that are checked
- A rapid response to users violating their contracts
- To keep the operations of noncompliant users transparent to con pliant users

4.7 PERFORMANCE PARAMETERS AT THE UNI

ITU-T Recommendation I.35B

ITU-T Recommendation I.35B defines a number of performance parameters for ATM-based networks. Figure 1-7 summarizes these categories, which are: (1) successfully delivered cells, *(2)* lost cells, (3) inserted cells, and (4) severely damaged cells.

As shown in Figure 1-7(a), any error-free cell that arrives (Δt) before its maximum allowed time (T) is considered to be a successfully delivered cell ($\Delta < T$). This cell must be "conformant", in that it is intelligible to the receiver. In Figure 9-

7(b), cell loss occurs when a cell arrives later than time T ($\Delta t > T$). Even though the cell arrived safely, it is still considered lost, because in some applications, such as voice and video, a late-arriving cell cannot be used in the digital-to-analog conversion process. In Figure 9-7(c), the cell is defined as lost when it is either lost or discarded in the network. Cell loss can occur if an error in the VPI/VCI is not detected, which results in a new or duplicate VPI/VCI being injected into the network, which of course results in the original cell disappearing from the network and the correct virtual circuit. It is also possible to experience a cell loss when a single-bit error results in the incorrect rebuilding of the header; which also results in a new VPI/VCI being injected in to the network, and a cell disappearing from the network and the virtual circuit. In Figure 1-7(d), the cell is defined as lost when the cell has an error of more than one bit in the header The header error correction function cannot correct errors of more than one bit. In Figure 1-7(e), an inserted cell error occurs when a cell arrives at the destination from a source other than the virtual connection source. In Figure 1-7(f), a severely damaged cell is one in which the user I field contains bit errors.

4.8 TRAFFİC MANAGEMENT AT THE UNI-BASIC CONCEPTS

I.

A wide number of alternatives exist for the management of traffic in an ATM network. This section provides a review of several proposals.

Eckberg Scheme

Figure 1-8 shows one proposal that has met with general approval from the ITU-T. This analysis was provided by Bell Laboratories and its Teletraffic Theory Group, which is supervised by A E. Eckberg

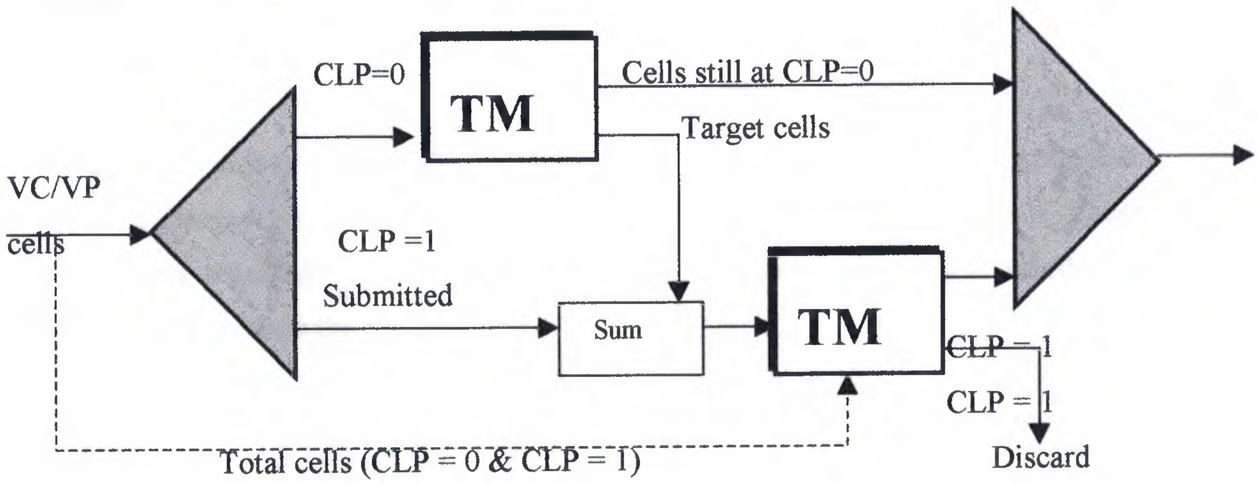

Figure 1-8
Proposed scheme for traffic management (from ITU paper CCITT 371). [ECKB92]

The scheme shown in this figure assumes that the user has a 5cr-vice contract with the ATM network wherein certain QOS parameters have been specified. Upon submitting traffic to the network, the user has the option of identifying individual cells with certain precedence. This decision by the user allows the ATM network service provider to determine how to treat the traffic in the event of congestion problems.

For traffic that is considered to be essential, the network then assumes the responsibility to assure that this traffic is treated fairly vis-a-vis the type of traffic. As an illustration, if the type of traffic is more tolerant to loss, the network then assumes the responsibility of tagging this traffic as possible loss traffic, and perhaps discarding this traffic if network congestion becomes a problem.

The major task of the ATM network is to make certain that the total cells presented to the network are consistent with the total cells processed by the network; this entails balancing the traffic submitted with CLP 0 and CLP 1 to that of the user's service contract.

On the left side of Figure 1-8, the user submits its traffic to the network node with the virtual channel and virtual path identifiers residing in the submitted cell. These identifiers are matched against the user's contract and decisions are made whether to grant the user access to the network.

A transmission-monitoring machine (TMM1) is responsible for keep-track of all cells submitted by the user with CLP 0 (those that not be lost, such as data). In the event of unusual problems, or (more likely) if the user violates the contract with the network, the can change the CLP of 0 to a 1. The cell that is tagged by the is called an excessive traffic tag. TMM1 does not tag all cells.

CHAPTER 5

ATM Layer Services

The ATM layer provides for the transparent transfer of fixed size ATM layer Service Data Units (ATM-SDUs) between communicating upper layer entities (e.g., AAT-entities). This transfer occurs on a pre-established ATM connection according to a traffic contract. A traffic contract is comprised of a QoS class, a vector of traffic parameters, a conformance definition and other items as specified in section 3.6. Each ATM end-point is expected to generate traffic, which conforms to these parameters. Enforcement of the traffic contract is optional at the Private UNI. The Public Network is expected to monitor the offered load and enforce the traffic contract.
Two levels of virtual connections can be supported at the UNI:

- A point-to-point or point-to-multipoint Virtual Channel Connection (VCC) which consists of a single connection established between two ATM VCC end-points.

- A point-to-point or point-to-multipoint Virtual Path Connection (VPC) which consist of a bundle of VCCs carried transparently between two ATM VPC end-points.

Note: For VPC at the Public UNI, traffics monitoring and throughput enforcement will be performed across all cells carried on the same VPI independently of the VCI values.

(R) From a single source the relay of cells within a VPC/VCC must preserve cell sequence integrity.

This layer performs no retransmission of lost or corrupted information. Flow control over ATM connections is for further study. The ATM layer also provides its users with the capability to indicate the loss priority of the data carried in each cell. The information exchanged between the ATM layer and the upper layer (e.g., the AAT,) across the ATM-SAP includes the following primitives:

Primitive	**Request**	**Indicate**	**Confirm**	**Respond**
ATM-DATA	X	X		

Figure 3-1 ATM Service Access Point (SAP) Primitives

These primitives make use of the following parameters:

Parameter	**Associated Primitives**	**Meaning**	**Valid values**
ATM-SDU	ATM-DATA.request ATM-DATA.indication	48 byte pattern for transport	Any 48 byte pattern
SDU-type	ATM-DATA.request ATM-DATA.indication	End-to-end cell type indicator	0 or 1
Submitted Loss-Priority	ATM-DATA.request	Requested Cell Loss Priority	High or Low Priority
Received Loss-Priority	ATM-DATA.indication	Received Cell Loss Priority	High or Low Priority
Congestion experienced	ATM-DATA.indication	EFCN indication	True or False

Figure 3-2 ATM-SAP Parameters

The primitives provide the following services:

ATM-DATA.request: Initiates the transfer of an ATM-SDU and its associated SDU-type to its peer entity over an existing connection. The loss priority parameter and the SDU-type parameter are used to assign the proper CLP and PTI fields to the corresponding ATM-PDU generated at the ATM layer.

ATM-DATA.indication: Indicates the arrival of an ATM-SDU over an existing connection, along with a congestion indication and the received ATM-SDU type. In the absence of errors, the ATM-SDU is the same as the ATM-SDU sent by the corresponding remote peer upper layer entity in an ATM-DATA.request.
The following parameters are passed within one or more of the previous primitives:

ATM-SDU: This parameter contains 48 bytes of ATM layer user data to be transferred by the ATM layer between peer communicating upper layer entities.

Submitted Loss Priority: This parameter indicates the relative importance of the requested transport for the information carried in the ATM-SDU. It can take only two values, one for high priority and the other for low priority.

Received Loss Priority: This parameter indicates the relative importance of the transport given to the information carried in the ATM-SDU. It can take only two values, one for high priority and the other for low priority.
Congestion indication: This parameter indicates that the received ATM-SDU has passed through one or more network nodes experiencing congestion.
SDU-type: the ATM layer user to differentiate two types of ATM-SDUs associated with an ATM connection only uses this parameter.

5.1 Service Expected from the Physical Layer

The ATM layer expects the Physical layer to provide for the transport of ATM cells between communicating ATM-entities. The information exchanged between the ATM layer and the Physical layer across the PHY SAP includes the following primitives:

Primitive	Request	Indicate	Confirm	Respond
PHY-UNITDATA[1]	X	X		

Figure 3-3 PHY-SAP Services required b y the ATM Layer

5.2 ATM Cell Structure and Encoding at the UNI

(R) Equipment supporting the UNI shall encode and transmit cells according to the structure and field encoding convention defined in T 1 LB 310 [7].

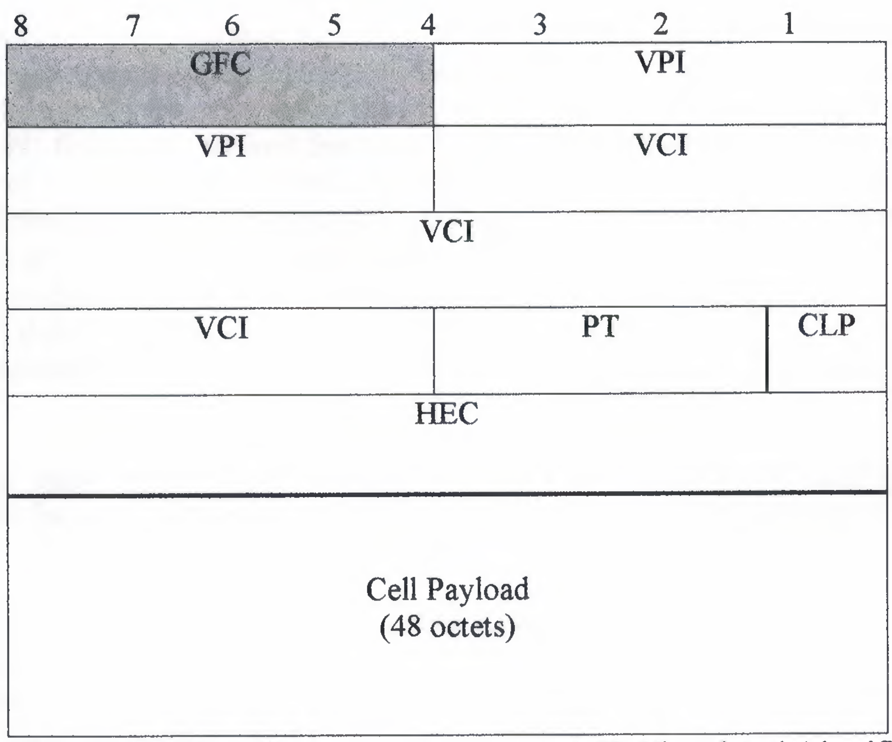

GFC: general Flow Control
VPI: Virtual Path Identifier
VCI: Virtual Channel Identifier
PT: Payload Type
CPL: Cell Loss Priority
HEC: Header Error Check

Figure 3-4 ATM Cell Structure at the UNI

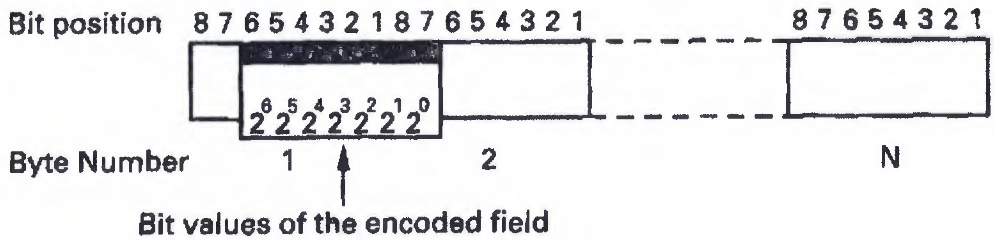

Figure 3-5 ATM Field Encoding Convention

The structure of the ATM cell is shown in Figure 3-4. It contains the following fields:

Generic Flow Control (GFC): This field has local significance only and can be use to provide standardized local functions (e.g. flow control) on the customer site. The value encoded in the GFC is not carried end-to-end and will be overwritten by the ATM switches. Two modes of operation have been defined for operation of the GFC field. These are "uncontrolled access" and "controlled access". The "uncontrolled access" mode of operation is used in early ATM environment. This mode has no impact on the traffic, which a host generates. Each host transmits the GFC field set to all zeros (0000). In order to avoid unwanted interactions between this mode and the "controlled access" mode where hosts are expected to modify their transmissions according to the activity of the GFC field, it is required that all CPE and public network equipment monitor the GFC field to ensure the attached equipment is operating in "uncontrolled mode". A count of the number of non-zero GFC fields should be measured for non-overlapping intervals of 30,000 +/- 10,000-cell times. If ten (10) or more non-zero values are received within this interval, an error is indicated to Layer Management.
(R) CPE at the UNI shall encode the GFC value to all zeros (0000).
(R) Public network equipment at the public UNI shall encode the GFC value to all zeros (0000).

(O) CPE shall inform Layer Management if a count of the non-zero GFC fields measured for non-overlapping intervals of 30,000 +/- 10,000 cell times reaches ten (10) or more.

(O) Public network equipment shall inform Layer Management if a count of the non-zero GFC fields measured for non-overlapping intervals of 3 0,000 +/- 10,000 cell times reaches ten (10) or more.

Virtual Path/Virtual Channel (VPI/VCI) Identifier: The actual number of routing bits in the VPI and VCI subfields use for routing is negotiated between the user and the network, e.g. on a subscription basis. This number is determined on the basis of the lower requirement of the user or the network.

Note: The number of VCI routing bits used in a user-to-user VP is negotiated between the users of the VP.
(R) The bits within the VPI and VCI fields used for routing are allocated using the following rules:

- The allocated bits of the VPI subfield shall be contiguous;
- The allocated bits of the VPI subfield shall be the least significant bits of the VPI subfield, beginning at bit S of octet 2;
- The allocated bits of the VCI subfield shall be contiguous;
- The allocated bits of the VCI subfield shall be the least significant bits of the VCI subfield, beginning at bit 5 of octet 4;

(R) Any bits of the VPI subfield that are not allocated are set to 0. For a given VP, any bits of the VCI subfield that are not allocated are set to 0.

Payload Type (PT): This is a 3 -bit field used to indicate whether the cell contains user information or Connection Associated Layer Management information (F5 flow). It is also used to indicate a network congestion state or for network resource management.

Cell Loss Priority (CLP): This is a 1-bit field, which allows the user or the network to optionally indicate the explicit loss priority of the cell. More details on the use of the CLP bit are given in section 3.4.5.
Header Error Control (HEC): The HEC field is used by the physical layer for detection/ correction of bit errors in the cell header. It may also be used for cell delineation.

5.3 ATM Layer Functions Involved at the UNI (U-plane)

This section describes ATM layer functions that need to be supported at the User-Network Interfaces (Figure 3-6). It does not cover those ATM functions that are described in standards but have no impact on the UNI specification.

Functions	Parameter
Multiplexing among different ATM connections	VPI/VCI
Cell rate decoupling (unassigned cells)	Pre-assigned header field values
Cell discrimination based on pre-defined header fields values	Pre-assigned header field values
Payload type discrimination	PT field
Loss priority indication and Selective cell discarding	CLP field, Network congestion state
Traffic shaping	Traffic descriptor

Figure 3-6 Functions supported at the UNI (U-plane)

Multiplexing among different ATM connections

This function multiplexes ATM connections with different QoS requirements. ATM connections may have either a specified or the unspecified QoS. The QoS class is the same for all cells belonging to the same connection, and remains unchanged for the duration of the connection.
(R) Network equipment supporting the public UNI shall support at least the Specified QoS Class 1
(O) Network Equipment supporting the public UNI may support the unspecified QoS class defined in section 4.2 of Appendix A.
(O) Network equipment supporting the private UNI may support either one or more specified QoS classes and/or the unspecified QoS class.

Cell rate decoupling

The cell rate decoupling function at the sending entity adds unassigned cells to the assigned cell stream (cells with valid payload) to be transmitted, transforming a non-continuos stream of assigned cells into a continuous stream of assigned and unassigned cells. At the receiving entity the opposite operation is performed for both unassigned and invalid cells. The rate at which the unassigned cell is inserted/extracted depends on the bit rate (rate variation) of assigned cell and/or the physical layer transmission rate. The unassigned and invalid cells are recognized by specific header patterns, which are shown in Figure 3-7.

Physical layers that have synchronous cell time slots generally required. Cell rate decoupling (e.g. SONET, DS3 and 8B/lOB block-coded interfaces) whereas physical layers that have asynchronous cell times slots do not require this function (e.g. 4B/SB block-coded interface) since no continuous flow of cells needs to be provided. Therefore the requirements in this section only apply to physical layers that require continuous cell streams at the Physical-ATM layers boundary.
(R) Equipment supporting the UNI shall generate unassigned cells in the flow of cells passed to the physical layer in order to adjust to the cell rate required by the payload capacity of the physical layer.
(R) Equipment supporting the UNI shall encode the header fields of unassigned cells in accordance with the pre-assigned header field values defined in T 1 LB 310 [7] and CCI'TT Recommendation 1.3.6.1.
(R) The receiving ATM entity shall extract and discard the unassigned and invalid cells from the flow of cells coming from the physical layer.
Note: The cell rate governing the flow between physical and ATM layer will be extracted from Physical layer (e.g. SONET) timing information if required.

Cells discrimination based on pre-defined header field values

The pre-defined header field values defined at the UNI are given in Figure 3-7 (Ref. T 1 LB 310).

Use	Value[1,2,3,4]			
	Octet 1	Octet 2	Octet 3	Octet 4
Unassigned cell indication	00000000	00000000	00000000	0000xxx0
Meta-signalling (default)[5,7]	00000000	00000000	00000000	00010a0c
Meta-signalling[6,7]	0000yyyy	yyyy0000	00000000	00010a0c
General Broadcast signalling (default)[5]	00000000	00000000	00000000	00100aac
General broadcast signalling[6]	0000yyyy	yyyy0000	00000000	00100aac
Point-to-point signalling (default)[5]	00000000	00000000	00000000	01010aac
Point-to-point signalling[6]	0000yyyy	yyyy0000	00000000	01010aac
Invalid Pattern	xxxx0000	00000000	00000000	0000xxx1
Segment OAM F4 flow cell[7]	0000aaaa	aaaa0000	00000000	00110a0a
End-to-End OAM F4 flow cell[7]	0000aaaa	aaaa0000	00000000	01000a0a

1: "a" indicates that the bit is available for use by the appropriate ATM layer function
2: "x" indicates "don't care" bits
3: "y" indicates any VPI value other than 00000000
4: "c" indicates that the originating signalling entity shall set the CLP bit to 0. The network may change the value of the CLP bit
5: Reserved for user signalling with the local exchange
6: Reserved for signalling with other signalling entities (e.g. other users or remote networks)
7: The transmitting ATM entity shall set bit 2 of octet 4 to zero. The receiving ATM entity shall ignore bit 2 of octet 4.

Figure 3-7 Pre-Defined Header Field Values

Meta-signalling cells are used by the meta-signaling protocol for establishing and releasing signaling virtual channel connections. For virtual channels allocated permanently (PVC), meta-signaling is not used.
(R) Equipment not supporting meta-signaling protocol at the UNI shall discard any cells received with VCI value = 1.
General broadcast signaling cells are used by the ATM network to broadcast signaling information independent of service profiles. For permanent virtual channel (PVC) service, the general broadcast-signaling channel is not used since there is no control-plane process involved above the ATM layer.
(R) Equipment not supporting general broadcast signaling at the UNI shall discard any cells received with VCI value = 2

The Virtual Path Connection (VPC) operation flow (F4 flow) is carried via specially designated OAM cells. F4 flow OAM cells have the same VPI value as the user-data cell transported by the VPC but are identified by two unique pre-assigned virtual channels within this VPC. At the UNI the virtual channel identified by a VCI value = 3 is used for VP level management, functions between ATM nodes on both sides of the UNI (i.e., single VP link segment) while the virtual channel identified by a VCI value = 4 can be used for VP level end-to-end (User <-> User) management functions.

The detailed layer management procedures making use of the F4 flow OAM cells at the UNI and the specific OAM cells format will be covered in section 3.5.
(R) Equipment supporting VP level management functions at the UNI shall encode the VCI field of the F4 flow OAM cells with the appropriate values as defined in T 1 LB.
(R) Equipment supporting VP level management functions at the UNI shall have the capability to identify F4 flow OAM cells within each VPC.
(R) Equipment not supporting VP level management functions at the UNI shall not transmit cells with VCI values 3 and 4 and shall discard any cells received with VCI value = 3 or 4.

A default header field value has been defined for the carriage of ILMI messages across the UNI the specific encoding is shown in Figure 3-8.

Use	**Value**			
	Octet 1	Octet 2	Octet 3	Octet 4
Carriage of ILMI message	00000000	00000000	0000001	0000aaa0

Figure 3-8 Default Header Field Value

1: "a" indicates that the bit is available for use by the appropriate ATM layer function 2: The transmitting ATM entity shall set the CLP bit to 0. The receiving ATM entity

Shall process MMI cells with CLP=1 as ILMI cells and as any other CLP=1 Ce.
(R) Equipment supporting the UNI shall support VPI = 0 and VCI = 16 as the default values for the carriage of ILMI messages across the UNI.

Cell discrimination based on Payload Type (PT) Identifier field values

The main purpose of the PT Identifier is to discriminate between user cells (i.e., cell carrying user information) from non-user cells (Figure 3-9). Code points 0 to 3 are used to indicate user cells. Within these code points, values 2 and 3 are used to indicate that congestion has been experienced in the network (see §3.6). Code points 4 and S are used for VCC level management functions. The PT value of 4 is used for identifying OAM cells communicated within the bounds of a VCC segment (i.e., single link segment across the UNI) while the PT value of S is used for identifying end-to-end OAM cells

The detailed layer management procedures making use of the F5 flow of OAM cells at the UNI and the specific OAM cells format is covered in section 3.5.

PTI Coding (MSB first)	Interpretation
000	User data cell, congestion not experienced, SDU-type = 0
001	User data cell, congestion not experienced, SDU-type = 1
010	User data cell, congestion experienced, SDU-type = 0
011	User data cell, congestion experienced, SDU-type = 1
100	Segment OAM F5 flow related cell
101	End-to-end OAM F5 flow related cell
110	Reserved for future traffic control and resource management
111	Reserved for future functions

Figure 3-9 Payload Type Indicator Encoding

(R) Equipment supporting VC level management functions at the UNI shall encode the PT field of the F5 flow OAM cells with the appropriate code points as defined in T 1 LB 310 [7].
(R) Equipment not supporting VC level management functions, via OAM cells, at the UNI shall ignore PT code points 100, 101.

Loss priority indication and selective cell discarding

The CLP field may be used for loss priority indication by the ATM end point and for selective cell discarding in network equipment. In a given ATM connection and for each user-data cell in the connection, the ATM equipment that first emits the cell can set the CLP bit equal to zero or one. The CLP bit is used to distinguish between cells of an ATM connection: A CLP bit equal to zero indicates a higher priority cell and a CLP bit equal to one indicates a lower priority cell. Upon entering the network, a cell with CLP value = 1 may be subject to discard depending on network traffic conditions.
The treatment of cells with CLP bit set (low priority cells) by the network traffic management functions is covered in section 3.6.
(O) User equipment supporting the UNI may use the CLP header field to indicate lower priority traffic (cells). ,

(O) ATM switches may tag CLP=0 cells detected by the UPC to be in violation of the Traffic Contract by changing the CLP bit from 0 to 1 (see 3.6).

Traffic Shaping

Traffic Shaping is expected to be an important function of ATM end-points in order to achieve the desired QOS. Traffic Shaping is covered in section 3-. 6.3 .2.5.

5.4 ATM Layer Management Specification (M-plane)

This section identifies the ATM Layer Management functions and procedures at the User Network Interface. Management functions at the UNI require some level of cooperation between customer premise equipment and network equipment. To minimize the coupling required between equipment on both sides of the UNI, the functional requirements have been reduced to a minimal set. The ATM Layer Management functions supported at the
UNI are grouped into the following categories (see Figure 3-10):

Functions	Parameter
Fault Management	
• *Alarm surveillance (VP)*	*OAM cells*
• *Connectivity Verification (VP,VC)*	*OAM cells*
• *Invalid VPI/VCI detection*	*VPI/VCI*

Figure 3-10 ATM Layer Management Functions at the UNI

- Fault Management contains Alarm Surveillance and Connectivity verification functions. OAM cells are used for exchanging related operation information. ATM cells with invalid VPI/VCI values are discarded and Layer Management is informed.

ATM Layer Management Information Flows

Figure 3 -11 shows the OAM flows defined for the exchange of operations information between nodes (including customer premises equipment). At the ATM layer, the F4-F5 flows will be carried via OAM cells. The OAM cell flow used for end-to-end management functions may be carried transparently through the private ATM switch and made available to the user. The F4 flow is used for segment2 or end-to-end

(VP termination) management at the VP level using VCI values 3 and 4. The F5 flow is used
UNI

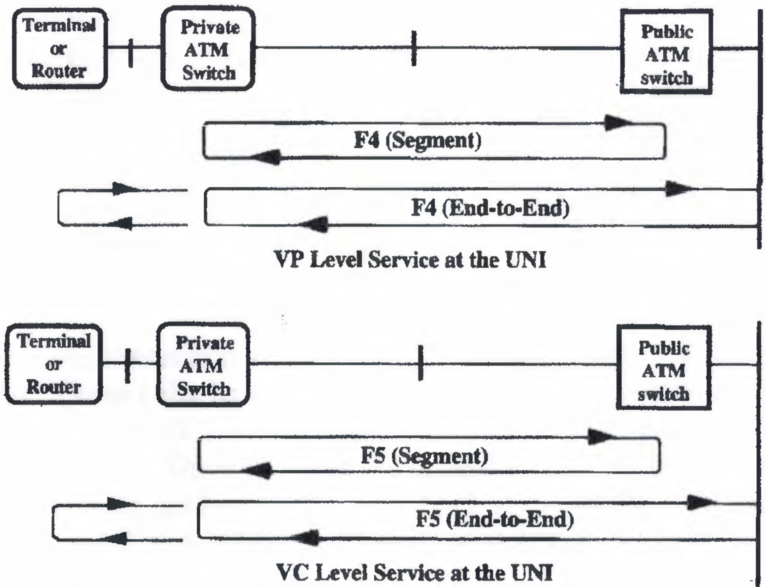

Figure 3-11 ATM Layer OAM flows at the UNI

for segment3 or end-to-end (VC termination) management at the VC level using PT code points 4 and 5. A detailed explanation on OAM cell flow mechanism is given in T1 S 1.5/92-029R3, CCITT Recommendation. I.610.

In case of only Virtual Path (VP) visibility (i.e., VPC service at the Public Network Interface) the OAM operation information exchange will be limited to the F4 flow. Under this scenario, any VC level OAM functions and information exchange (F5 flow) are user-specific and ignored by the network. It is however possible to have VP level service at the Public UNI while maintaining VC visibility at the Private UNI. In this case, the private ATM switch would terminate the F4 flow but could carry transparently the user end-to-end F5 flow. For Virtual Channel (VC) visibility (VCC service), the OAM operation information exchange specified at the Public UNI could be limited to the F5 flow or could invoke both F4 and F5 flows.

(R) Equipment requiring/offering VP level service at the UNI shall support the F4 management flow as defined in T 1 S 1. 5/92 -02 9R3 and CCITT Recommendation

I.610 for the functions defined in 3.5.3. Equipment should also be capable of transparently passing/carrying the F5 management flow.
(R) Equipment requiring/offering VC level service at the UNI shall support the FS management flow as defined in T 1 S 1. S/92 -02 9R3 and CCITT Rec. I.610 for the required functions defined 3.5.3.
(R) Equipment requiring/offering VC level service only at a UNI shall either 1) process the F4 flow in accordance with T 1 S 1. 5/92 002 9R3 and ITU-T Recommendation I.610, or 2) discard any F4 flow cells (i.e., cells with VCI = 3 and VCI = 4 are considered as cells with invalid VCI value).
The definition of "zero bandwidth" for a particular direction of any connection does not prohibit the transmission of OAM cells (ref. 3.6.3.2.3.7).
(R) For a point-to-multi-point connection, the only allowed use of F4 and FS OAM flows is segment flows. Neither the user nor the network shall send end-to-end OAM cells.

ATM OAM Cell Format

The virtual path connection (VPC) operational information is carried via the F4 flow OAM cells. These cells have the same VPI value as the user-data cells but are identified by pre-assigned VCI values. Two unique VCI values are used for every VPC as shown in Figure 3-12 a. The VCI value = 3 is used to identify the connection between ATM layer management entities (LMEs) on both sides of the UNI (i.e., single link segment) and VCI value = 4 is used to identify connection between end-to-end ATM LMEs.

The virtual channel connection (VCC) operation information is carried via the FS flow OAM cells. These cells have the same VPI/VCI values as the user-data cells but are identified by pre-assigned code points of the Payload Type (PT) field. Two unique PT values are used for every VCC as shown in Figure 3-12 a. The PT value = " 100" (4) is used to identify the connection between ATM layer management entities (LMEs) on both sides of the UNI (I.e., single link segment) while the PT value = "101" (5) is used to identify connection between end-to-end ATM LMEs.

End-to-end OAM cells must be passed unmodified by all intermediate nodes. Any node in the path may monitor the contents of these cells. These cells are only to be removed by the endpoint of the VPC (F4 flow) or VCC (FS flow). Segment OAM cells shall be removed at the end of a segment where, for the purposes of this specification, segment is defined as a single VP or VC link across the UNI.
The format of the Functions-Specific fields of the Fault Management OAM cell is defined in T 1 S 1. 5/92 -02 9R3 and CCITT Rec. I.610 and shown in Figure 3-12 b.
(R) Equipment supporting the F4 and/or FS flow at the LTM shall encode /interpret the OAM cells according to the format and encoding rules defined in T 1 S 1. 5/92 -02 9R 1 and CCITT Rec. I.610 for the functions defined in 3.5.3.

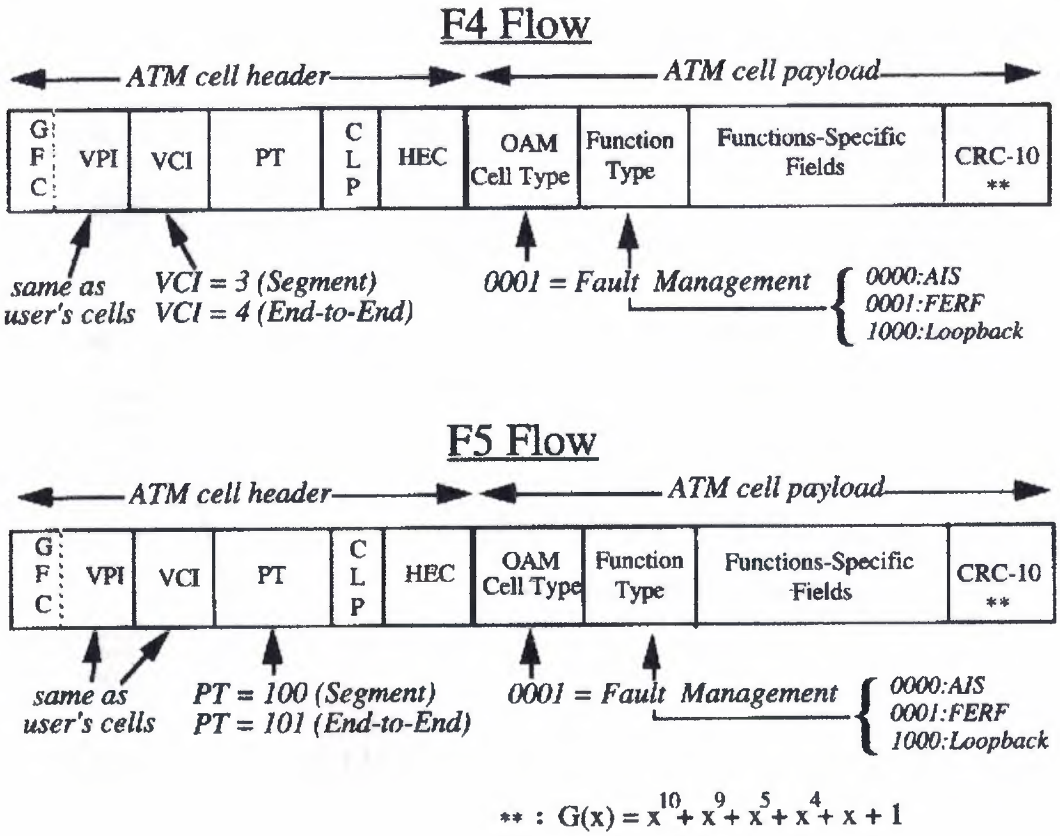

Figure 3-12a Common OAM Cell Format

Note: OAM Fault Management functions and cell formats may continue to evolve and software implementations may be advisable.

ATM Fault Management functions at the UNI

The Fault Management functions at the UNI are grouped in two categories: Alarm Surveillance and Connectivity Verification.

Alarm Surveillance

Alarm surveillance at the public UNI involves detection, generation and propagation of VPC/VCC failures (failure indications). In analogy with SONET physical layer, the failure indication signals are of two types: Alarm Indication Signal (AIS) and Far End Receive Failure (FERF). These signals are carried via OAM cells as defined in section 3.5.2. A VPC/VCC node at a connecting point to alert the downstream VPC/VCC nodes

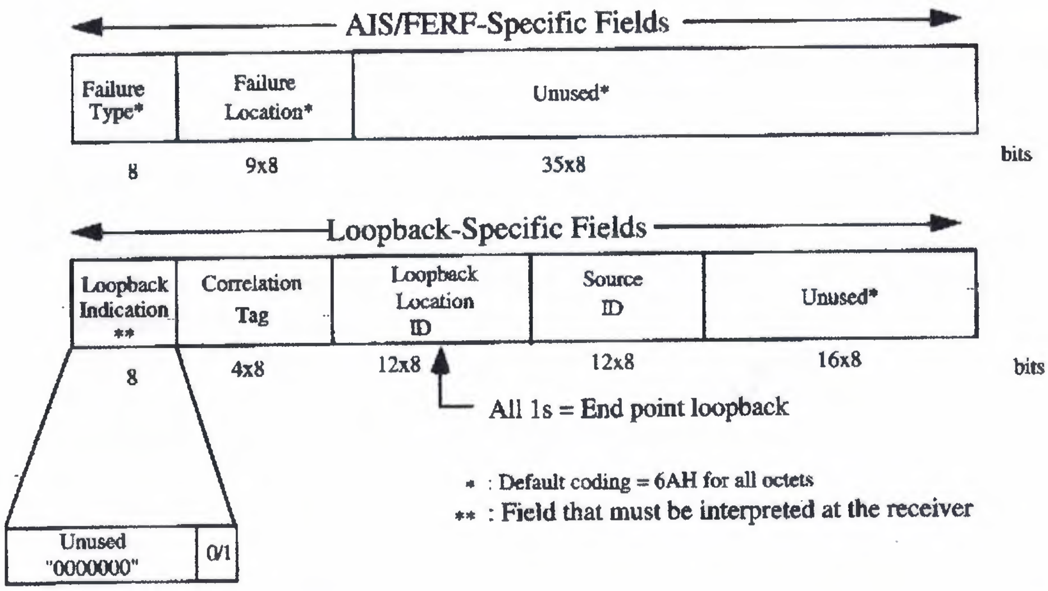

Figure 3-12b OAM Cell Fault Management-Specific Fields

that a failure has been detected upstream generates the VP/ VC Alarm Indication Signal (VP-AIS/VC-AIS). The VP-AIS/VC-AIS can be caused by the detection of a VPC/VCC failure or by the notification of a physical layer failure.
(R) ATM End-Point at the public UNI shall detect all incoming VP-AIS and generate a VP-FERF in the upstream direction (toward the public network) to alert the ATM nodes about the failure.
(R) Public network equipment supporting the UNI, acting as VP intermediate node, shall generate VP-AIS upon detection of a VPC failure or upon receiving a physical layer failure notification.
(O) End-Point of VCCs traversing the public network may detect an incoming VC-AIS and generate a VC-FERF in the upstream direction (toward the public network).
(O) Public network equipment supporting the UNI may generate VC-AIS upon detection a VCC failure or upon receiving a physical layer failure notification.
Equipment inserting Alarm Surveillance cells will do so at a rate low enough to insure that Alarm Surveillance cells amount to less than one percent of the capacity of any link in the connection.
The duration of the condition and the rate associated with the generation and removal of alarm signals (VP-AIS, VP-FERF) is to be defined.

Connectivity Verification

Connectivity verification is supported by the use of the OAM loopback capability for both VP and VC connections. More complete details on this loopback function can be found in the modified text of I.610 [37]. The VCC or VPC being checked can remain in service while this loopback function is being performed. The OAM Cell Loopback function supported at the UNI uses the following three fields:

- Loopback indication - This eight-bit field identifier for the endpoint receiving the OAM cell, whether the incoming cells is to be looped back. A value of (00000001) indicates that the cell should be looped back. All other values indicate that the cell is to be discarded. Before the cell is looped back, the end point should decrement the value c the loopback indication field.
- Correlation Tag - At any given time multiple OAM Fault Management cells may be inserted in the same virtual connection. As a result, the OAM cell loopback mechanism requires a means of correlating transmitted OAM cells with received OAM cells. Maybe the node inserting the OAM cell put any value in this 32 -bit field and the endpoint looping back the cell ought to not modify it.
- Loopback Location ID (optional) - This 96-bit field identifies the point(s) along a virtual connection where the loopback is to occur. The transmitter to indicate the end point uses the default value of all ones. The receiver is not required to decode this field.
- Source ID (optional) - This 96-bit field can be used to identify the originator of the Loopback cell so the originator can identify the looped back cell when it returns. This may be encoded any way that enables the originating point to be certain that it has received the cell it transmitted. The default value is all ones.

(R) Endpoints receiving OAM cells with a loopback indication value other than (00000001) and Function Type = 1000 shall discard the cell.
(R) Endpoints receiving OAM cells with a loopback indication value of (00000001) and Function Type = 1000 shall decrement the loopback indication value and then loopback the cell within one second.
For connections that do not terminate in the public network, public network equipment will only insert end to end loopback cells when attempting to verify or isolate a fault. Equipment inserting loopback cells will do so at a rate low enough to insure that loopback cells amount to less than one percent of the capacity of any link in the connection. No requirement should be made that loopback cells support delay measurement.

End to End Loopback

End to end loopback cells is only looped back by the end point of a VPC or VCC. These cells may be inserted by any node in the connection (intermediate or end point) and may be monitored by any node. However, only end points may remove these cells. End to end loopback cells is indicated by a Payload Type value of (101) for VCCs and a VCI value of (4) for VPCs. An example of how the end to end loopback cell is used is shown in Figure 3-13.

UNI Loopback

UNI loopback is performed using segment loopback cells, which are looped back by the end point of a VPC or VCC segment. The segment is defined as the link between the ATM nodes on either side of the UNI. Segment end points must either remove these cells or loop them back depending on the value in the loopback indication field. That is, these cells must not travel beyond the segment in which they are generated. UNI loopback cells are indicated by a Payload Type value of (100) for VCCs and a VCI value of (3) for VPCs. An example of how the segment loopback cell is used is shown in Figure 3-13.

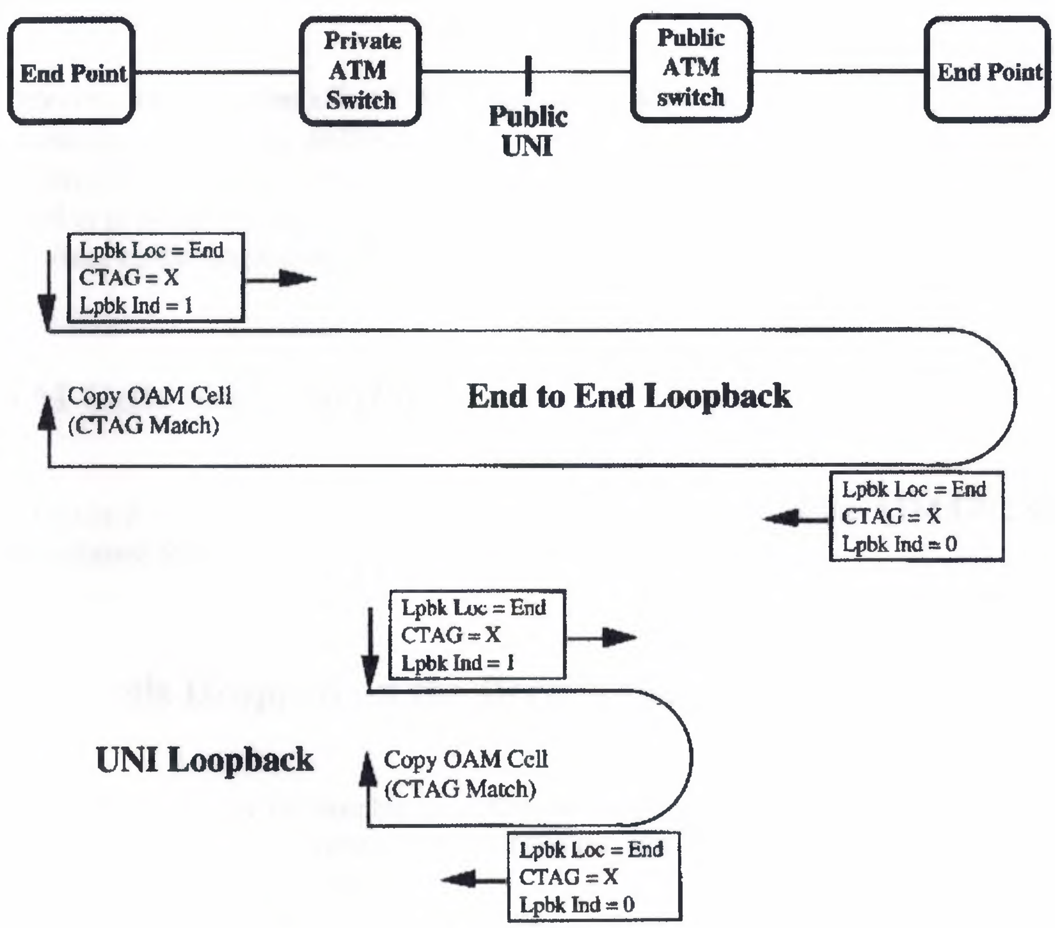

Figure 3-13 Loopback Function

ATM Layer Statistics (O)

These attributes are located in the ATM Statistics Group. This group is indexed by the interface index. MIB information at this level includes;

- Interface Index
- ATM Cells Received
- ATM Cells Dropped on the Receive Side
- ATM Cells Transmitted

Certain counters make sense only at the ATM Layer for an entire physical interface, while others are rolled up across all VPCs and VCCs at an interface. These counters are thus aggregated across the entire ATM UNI, and accumulated over time. All counters are 3 2 bits which wrap around.

The optional ATM layer statistics can be used for the following purposes.

- Identify problems that affect UNI performance, such as:
- Loss due to bit errors, and/or
- Unavailability of link connections.
- Aid in problem diagnosis and troubleshooting.
- Collect traffic engineering data.

ATM Cells Received (O)

This is a count of the number of ATM cells received across the ATM UNI, which were assigned and not dropped.

ATM Cells Dropped on the Receive Side (O)

This is a count of the number of ATM cells received across the ATM UNI, which were dropped for the following specific reasons. The reasons for which an ATM cell can increment this counter shall be:

The Header Error Control (HEC) processing either detected an error, or identified an uncorrectable error once cell delineation has been achieved.
An invalidity formatted cell header was received.

This device was not configured to process the received VPI/VCI (i.e., it is unknown). For example, a switch translation table entry was not defined for the received VPI/VCI value.

ATM Cells Transmitted (O)

This is a count of the number of ATM cells transmitted across the ATM UNI, which were assigned. This is a count of cells carrying actual user data across the ATM UNI.

Per-Virtual Path UNI MIB Attributes (R)

These attributes are located in the Virtual Path Group (ATMs VpcGroup). The interface index (atmfVpcPortIndex) and the VPI value (atmfVpcVpi) index this group. MIB information at this level includes:

- Interface Index
- VPI Value
- Transmit Traffic Descriptor
- Receive Traffic Descriptor
- Operational Status
- Transmit QoS Class
- Receive QoS Class

Interface Index (R)

The Interface Index is the same as that for the Physical interface as defined in section 4.2.1.1
It is implicitly the local ATM UNI, unless the optional mode of explicit identification is supported.

VPI (R)

This is the value of the Virtual Path Identifier (VPI) for this VPC.

Transmit Traffic Descriptor (R)

This is a specification of the conformance definition and associated source traffic descriptor parameter values described in section 3.6.2.5 that are applicable to the transmit side of this interface for this VPC.

Receive Traffic Descriptor (R)

This is a specification of the conformance definition and associated source traffic descriptor parameter values described in section 3.6.2.5 that are applicable to the receive side of this interface for this VPC.

Operational Status (R)

This object represents the state of the VPC as known by the local device. If the end-to-end status is known then a value of end2 endue (2) or end2 endDown (3) is used. If only the local status is known then a value of localUpEnd2endUnknown (4) or localDown (5) is used.

Transmit (QoS) Class (R)

This is the QoS Class defined in section 4 of Appendix A, for the transmit side of this interface for this VPC.

Receive (QoS) Class (R)

This is the QoS Class defined in section 4 of Appendix A, for the receive side of this interface for this VPC.

ATM Call States Relating to Interworking Requirements

Not supported in this Implementation Agreement.

States Associated with the Global Call Reference

This section defines the states that the protocol may adopt using the global call reference. The procedures for use of the global call reference for restart procedures are contained in §5.5.5. There is only one global call reference per interface.

Call States at the User Side of the Interface

The states which may exist on the user side of the user-network interface are defined in this section.

Null (Rest 0)

No transaction exists.

Restart Request (Rest 1)

This state exists for a restart transaction when the user has sent a restart request but has not yet received an acknowledgment response from the network.

Restart (Rest 2)

This state exists when a request for a restart has been received from the network and responses have not yet been received from all locally active call references.

Call States at the Network Side of the Interface

The states which may exist on the network side of the user-network interface are defined in this section.

Null (Rest 0)

No transaction exists.

Restart Request (Rest 1)

This state exists for a restart transaction when the network has sent a restart request but has not yet received an acknowledgment response from the user.

Restart (Rest 2)

This state exists when a request for a restart has been received from the user and a response has not yet been received from all locally active call references.

5.6 Message Functional Definitions and Contents

This section provides an overview of the message structure, which highlights the functional definitions and information content (i.e., semantics) of each message. Each definition includes: A brief description of the message direction and use, including whether the message has: 1

a) Local significance, i.e., relevant only in the originating or terminating access;
b) Access significance, i.e., relevant in the originating and terminating access, but not. In the network,
c) Dual significance, i.e., relevant in either the originating or terminating access and in the network; or
d) Global significance, i.e., relevant in the originating and terminating access and in the network.

Note - Messages of access significance and dual significance are not supported in this Implementation Agreement

2. A table listing the code set 0 (ITU-TS [CCI'T'~I`] standardized) information elements. For each information element the table indicates:

a) the section of this Implementation Agreement describing the information element;

b) the direction in which it may be sent; i.e., user to network (`u -> n'), network to user (`n -> u'), or both;

Note - The user-network terminology in this section refers to the interface structures between ATM terminal equipment and ATM public or private network ~ (TE-LCRF/CN) and between ATM customer network and ATM public network (CN-LCRF); the terms TE, CN, and LCRF being used as defined in CCITT Recommendation I. 3 2 7.

c) Whether inclusion is mandatory (`M') or optional (`O'), with a reference to notes explaining the circumstances under which the information element shall be included; and

d) The length of the information element (or permissible range of lengths), in octets.

3. Further explanatory notes, as necessary.

Messages for ATM point-to-point call and connection control

Table 5-2 summarizes the messages for ATM point-to-point call and connection control.

Table 5-2 Messages for ATM Call and Connection Control

Message	Reference
Call establishment messages:	
CALL PROCEEDING	5.3.1.2
CONNECT	5.3.1.3
CONNECT ACKNOWLEDGE	5.3.1.4
SETUP	5.3.1.7
Call clearing messages:	
RELEASE	5.3.1.5
RELEASE COMPLETE	5.3.1.6
Miscellaneous messages:	
STATUS	5.3.1.8
STATUS ENQUIRY	5.3.1.9

ALERTING

Not supported in this Implementation Agreement.

CALL PROCEEDING

The called user sends this message to the network or the network to the calling user to indicate that the requested call establishment has been initiated and no more call establishment will accept information.
The sending of this message is optional; the receiving of this message is required (see § 5.5)

Message Type: CALL PROCEEDING
Significance: local
Direction: both

Information Element	Reference	Direction	Type	Length
Protocol discriminator	5.4.2	both	M	1
Call reference	5.4.3	both	M	4
Message type	5.4.4.1	both	M	2
Message length	5.4.4.2	both	M	2
Connection identifier	5.4.5.1.6	both	O(1)	4-9
Endpoint reference	5.4.8.1	both	O(2)	4-7

Note 1 - Mandatory in the network-to-user direction if this message is the first message in response to a SETUP message. It's mandatory in the user to-network direction if this message is the first response to a SETUP message, unless the user accepts the connection identifier indicated in the SETUP message. Note 2 - Mandatory if an Endpoint reference was included in the SETUP message. Figure 5-2 CALL PROCEEDING Message Content

CONNECT

This message is sent by the called user to the network and by the network to the calling user to indicate call acceptance by the called user.

Message Type: CONNECT
Significance: global
Direction: both

Information Element	Reference	Direction	Type	Length
Protocol discriminator	5.4.2	both	M	1
Call reference	5.4.3	both	M	4
Message type	5.4.4.1	both	M	2
Message length	5.4.4.2	both	M	2
AAL parameters	5.4.5.5	both	O(1)	4-11
Broadband low layer information	5.4.5.9	both	O(2)	4-17
Connection identifier	5.4.5.16	both	O(3)	4-9
Endpoint reference	5.4.8.1	both	O(4)	4-7

Figure 5-3 CONNECT Message Contents

Note 1 - Included in the user to network direction when the called user wants to pass ATM adaptation layer parameters information to the calling user, and the ATM adaptation layer parameters information element was present in the SETUP message. Included in the network to user direction if the called user included an ATM adaptation layer parameters information element in the COTJNECT message. The ATM adaptation layer parameters information element shall not be present when the endpoint reference information element was present in the SETUP message and contained a non-zero value.
Note 2 - Included in the user-to-network direction when the answering user wants to return low layer information to the calling user. Included in the network-to-user direction if the user awarded the call included a Broadband low layer information element in the CONNECT message. Optionally included for Broadband low layer

information negotiation, but some networks may not transport this information element to the calling user (see Annex C)

. Note 3 - Mandatory in the network-to-user direction if this message is the first message in response to a SETUP message. It's mandatory in the user-to-network direction if, this message is the first response to a SETUP message, unless the user accepts the connection identifier indicated in the SETUP message.

Note 4 - Mandatory if the Endpoint reference was included in the SETUP message.

CONNECT ACKNOWLEDGE

This message is sent by the network to. The called user to indicate the user has been awarded the call. It is also sent by the calling user to the network to allow symmetrical call control procedures.

Message Type: CONNECT ACKNOWLEDGE
Significance: local
Direction: both

Information Element	Reference	Direction	Type	Length
Protocol discriminator	5.4.2	both	M	1
Call reference	5.4.3	both	M	4
Message type	5.4.4.1	both	M	2
Message length	5.4.4.2	both	M	2

Figure 5-4 CONNECT ACKNOWLEDGE Message Contents

RELEASE

This message is sent by the user to request the network to clear the end-to-end connection (if any) or is sent by the network to indicate that the end-to-end connection is cleared and that the receiving equipment should release the virtual channel and prepare to release the call reference after sending a RELEASE COMPLETE.

Message Type: RELEASE
Significance: global
Direction: both

Information Element	Reference	Direction	Type	Length
Protocol discriminator	5.4.2	Both	M	1
Call reference	5.4.3	Both	M	4
Message type	5.4.4.1	Both	M	2
Message length	5.4.4.2	both	M	2
Cause	5.4.5.15	Both	M	6-13

Figure 5-5 RELEASE message content

RELEASE COMPLETE

This message is sent by the user or the network to indicate that the equipment sending the message has released the virtual channel (if any) and call reference, the virtual channel is available for reuse, and the receiving equipment shall release the call reference.
Message Type: RELEASE COMPLETE
Significance: local (l)
Direction: both

Information Element	Reference	Direction	Type	Length
Protocol discriminator	5.4.2	Both	M	1
Call reference	5.4.3	Both	M	4
Message type	5.4.4.1	Both	M	2
Message length	5.4.4.2	both	M	2
Cause	5.4.5.15	Both	O(2)	4-34

Figure 5-6 RELEASE COMPLETE message content

SETUP

This message is sent by the calling user to the network arid by the network to the called user to initiate call establishment.
Message Type: SETUP
Significance: global
Direction: both

Information Element	Reference	Direction	Type	Length
Protocol discriminator	5.4.2	both	M	1
Call reference	5.4.3	both	M	4
Message type	5.4.4.1	both	M	2
Message length	5.4.4.2	both	M	2
AAL parameters	5.4.5.5	both	O(1)	4-20
ATM user cell rate	5.4.5.6	both	M	12-30
Broadband bearer capability	5.4.5.7	both	M	6-7
Broadband high layer information	5.4.5.8	both	O(2)	4-13
Broadband repeat indicator	5.4.5.19	both	O(3)	4-5
Broadband low layer information	5.4.5.9	both	O(4)	4-17
Called party number	5.4.5.11	both	M	(5)
Called party sub-address	5.4.5.12	both	O(6)	4-25
Calling party number	5.4.5.13	both	O(7)	4-26
Calling party sub-address	5.4.5.14	both	O(8)	4-25
Connection identifier	5.4.5.16	N -> U	M	9
QoS parameter	5.4.5.18	both	M	6
Broadband sending complete	5.4.5.21	both	O(9)	4-5
Transit network selection	5.4.5.22	U -> N	O(10)	4-8
Endpoint reference	5.4.8.1	both	O(11)	4-7

Figure 5-7 SETUP message content

Note 1 - Included in the user-to-network direction when the calling user wants to pass ATM adaptation layer parameters information to the called user. Included in the network-to-user direction if the calling user included an ATM adaptation layer parameters information element in the SETUP message.

Note 2 - Included in the user-to-network direction when the calling user wants to pass broadband high layer information to the called user. Included in the network-to-user direction if the calling user included a Broadband high layer information element in the SETUP message.

Note 3 - Included when two or more Broadband low layer information information elements are included for Broadband low layer information negotiation.

Note 4 - Included in the user-to-network direction when the calling user wants to pass broadband low layer information to the called user. Included in the network-to-user direction if the calling user included a Broadband low layer information information element in the SETUP message. Two or three information elements may be included in descending order of priority, i.e., highest priority first, if the. Broadband low layer information negotiation procedures are used (see Annex C)

Note 5 - Minimum length depends on the numbering plan. Maximum length is 2 5 octets. Note 6 - Included in the user-to-network direction when the calling user wants to indicate the called party sub addresses. Included in the network-to-user direction if the calling user included a Called party sub address information element in the SETUP message.

Note 7 - May be included by the calling user or by the network to identify the calling user.

Note 8 - Included in the user-to-network direction when the calling user wants to indicate the calling party sub addresses. Included in the network-to-user direction if the calling user included a Calling party sub address information element in the SETUP message.

STATUS

This message is sent by the user or the network in response to a STATUS ENQUIRY message or at any time to report certain error conditions listed in § 5 .5.

Information Element	Reference	Direction	Type	Length
Protocol discriminator	5.4.2	both	M	1
Call reference	5.4.3	both	M	4
Message type	5.4.4.1	both	M	2
Message length	5.4.4.2	both	M	2
Call state	5.4.5.10	both	M	5
Cause	5.4.5.15	both	M	6-14
Endpoint reference	5.4.8.1	both	O(1)	4-7
Endpoint state	5.4.8.2	both	O(2)	4-5

Note 1 - Included when responding to a status inquiry about a party state or at any time to report certain error conditions in the point-to-multipoint procedures.
Note 2 - Included when the Endpoint reference information element is included.

Figure 5-8 STATUS message content

STATUS ENQUIRY

The STATUS ENQUIRY message is sent by the user or the network at any time to solicit a STATUS message from the peer layer 3 entity. Sending a STATUS message in response to a STATUS ENQUIRY message is mandatory.

Message type: STATUS ENQUIRY
Significance: local
Direction: both

Information Element	Reference	Direction	Type	Length
Protocol discriminator	5.4.2	both	M	1
Call reference	5.4.3	both	M	4
Message type	5.4.4.1	both	M	2
Message length	5.4.4.2	both	M	2
Endpoint reference	5.4.8.1	both	O(1)	4-7

Figure 5-9 STATUS ENQUIRY message content

Messages for the Support of 64 kbit/s based ISDN Circuit Mode Services

Not supported in this Implementation Agreement.

Messages Related to Release 1 Supplementary Services

Not supported in this Implementation Agreement.

Messages Used with the Global Call Reference

Table 5-3 Messages Used with the Global Call Reference

Message	Reference
RESTART	5.3 .4.1
RESTART ACKNOWLEDGE	5.3.4.2
STATUS	5.3.1.8

RESTART

This message is sent by the user or the network to request the recipient to restart (i.e., release all resources associated with) the indicated virtual channel or all virtual channels controlled by the Signaling Virtual Channel.

Message type: RESTART
Significance: local
Direction: both

Information Element	Reference	Direction	Type	Length
Protocol discriminator	5.4.2	both	M	1
Call reference	5.4.3	both	M(1)	4
Message type	5.4.4.1	both	M	2
Message length	5.4.4.2	both	M	2
Connection identifier	5.4.5.16	both	O(2)	4-9
Restart indicator	5.4.5.20	both	M	5

Note 1 - This message is sent with the global call reference defined in §5.4.3.
Note 2 - Included when necessary to indicate the particular virtual channel to be restarted.

Figure 5-10 RESTART message content

RESTART ACKNOWLEDGE

This message is sent to acknowledge the receipt of a RESTART message and to indicate that the requested restart is complete.

Message type: RESTART ACKNOWLEDGES.
Significance: local
Direction: both

Information Element	Reference	Direction	Type	Length
Protocol discriminator	5.4.2	Both	M	1
Call reference	5.4.3	Both	M(1)	4
Message type	5.4.4.1	Both	M	2
Message length	5.4.4.2	both	M	2
Connection identifier	5.4.5.16	Both	O(2)	4-9
Restart indicator	5.4.5.20	Both	M	5

Note 1 - This message is sent with the global call reference defined in § 5 .4. 3 .
Note 2 - Included when necessary to indicate the particular virtual channel which has been restarted.

Figure 5-11 RESTART ACKNOWLEDGE message content

Messages for Point-to-multipoint call and connection control

Table 5-4 Messages Used with ATM Point-to-multipoint call and connection control

Message	Reference
ADD PARTY	5.3 .5.1
ADD PARTY ACKNOWLEDGE	5.3.5.2
ADD PARTY REJECT	5.3.5.3
DROP PARTY	5.3.5.4
DROP PARTY ACKNOWLEDGE	5.3.5.5

ADD PARTY

This message is sent to add a party to an existing connection (see §5.6).

Message type: ADD PARTY

Significance: global

Direction: both

Information Element	Reference	Direction	Type	Length
Protocol discriminator	5.4.2	both	M	1
Call reference	5.4.3	both	M	4
Message type	5.4.4.1	both	M	2
Message length	5.4.4.2	both	M	2
AAL parameters	5.4.5.5	both	O(1)	4-20
Broadband high layer information	5.4.5.8	both	O(2)	4-13
Broadband low layer information	5.4.5.9	both	O(3)	4-17
Called party number	5.4.5.11	both	M	(4)
Called party sub-address	5.4.5.12	both	O(5)	4-25
Calling party number	5.4.5.13	both	O(6)	4-26
Calling party sub-address	5.4.5.14	both	O(7)	4-25
Broadband sending complete	5.4.5.21	both	O(8)	4-5
Transit network selection	5.4.5.22	U -> N	O(9)	4-8
Endpoint reference	5.4.8.1	both	O(10)	7

Figure 5-12 ADD PARTI~ message content

Note 1 - Included in the user-to-network direction when the calling user wants to pass ATM adaptation layer parameters information to the called user. Included in the network-to-user direction if the calling user included an ATM adaptation layer parameters information element in the ADD PARTY message. Must be the same as in the initial SETUP of the call, but is not checked by the network.

Note 2 - Included in the user-to-network direction when the calling user wants to pass broadband high layer information to the called user. Included in the network-to-user direction if the calling user included a Broadband high layer information information element in the ADD PARTY message. Must be the same as in the initial SETUP of the call, but is not checked by the network.

Note 3 - Included in the user to-network direction when the calling user wants to pass broadband low layer information to the called user. Included in the network-to-user direction if the calling user included a Broadband low layer information information element in the ADD PARTY message. Must be the same as the one negotiated during the initial SETUP of the call, but is not checked by the network. Only one Broadband low layer information element is permitted in the ADD PARTY message.

Note 4 - Minimum length depends on the numbering plan. Maximum length is 2 5 octets.

Note 5 - Included in the user-to-network direction when the calling user wants to indicate the called party sub addresses. Included in the network-to-user direction if the calling user included a Called Party Sub address information element in the ADD PARTY message.

Note 6 - May be included by the calling user, or by the network to identify the calling user.

Note 7 - Included in the user-to-network direction when the calling user wants to indicate the calling party sub address. Included in the network-to-user direction if the calling user included a Calling Party Sub address information element in the ADD PARTY message.

Note 8 - It is optional for the user to include the Broadband sending complete information element when enbloc sending procedures (i.e., complete address information is included) are used; its interpretation by the network is optional. It is optional for the network to include the Broadband sending complete information element when enblock receiving (i.e., complete address information is included) are used.

Note 9 - Included by the calling user to select a particular transit network (see Annex D.)

Note 10 - The endpoint reference must be unique within a given call reference on a given link.

ADD PARTY ACKNOWLEDGE

This message is sent to acknowledge that the ADD PARTY request was successful.

Message type: ADD PARTY ACKNOWLEDGES
Significance: global
Direction: both

Information Element	Reference	Direction	Type	Length
Protocol discriminator	5.4.2	both	M	1
Call reference	5.4.3	both	M	4
Message type	5.4.4.1	both	M	2
Message length	5.4.4.2	both	M	2
Endpoint reference	5.4.8.1	both	M(1)	7

Figure 5-13 ADD PARTY ACKNOWLEDGE message content

ADD PARTY REJECT

This message is sent to acknowledge that the ADD PARTY request was not successful.
Message type: ADD PARTY REJECTS
Significance: global
Direction: both

Information Element	Reference	Direction	Type	Length
Protocol discriminator	5.4.2	both	M	1
Call reference	5.4.3	both	M	4
Message type	5.4.4.1	both	M	2
Message length	5.4.4.2	both	M	2
Cause	5.4.5.15	both	M	6-34
Endpoint reference	5.4.8.1	both	M(1)	7

Figure 5-14 ADD PARTY REJECT message content

DROP PARTY

This message is sent to drop (clear) a party from an existing point-to-multipoint connection.
Message type: DROP PARTY
Significance: global
Direction: both

Information Element	Reference	Direction	Type	Length
Protocol discriminator	5.4.2	both	M	1
Call reference	5.4.3	both	M	4
Message type	5.4.4.1	both	M	2
Message length	5.4.4.2	both	M	2
Cause	5.4.5.15	both	M	6-34
Endpoint reference	5.4.8.1	both	M	7

Figure 5-15 DROP PARTY message content

DROP PARTY ACKNOWLEDGE

This message is sent in response to a DROP PARTY message to indicate that the party was dropped from the connection.
Message type: DROP PARTY ACKNOWLEDGES
Significance: local
Direction: both

Information Element	Reference	Direction	Type	Length
Protocol discriminator	5.4.2	both	M	1
Call reference	5.4.3	both	M	4
Message type	5.4.4.1	both	M	2
Message length	5.4.4.2	both	M	2
Cause	5.4.5.15	both	O(1)	4-34
Endpoint reference	5.4.8.1	both	M	7

Note 1 - Mandatory when DROP PARTY ACKNOWLEDGE is sent as a result of an error condition.

Figure 5-16 DROP PARTY ACKNOWLEDGE message content

CONCLUSION

As the 21st century begins, telecommunications and data communications are converging. Both customer-premises equipment and transmission facilities are increasingly using digital technology. Digital PBXs are being used to integrate voice and data with digital phones. LANs and PBXs can be connected to packet switching networks, frame relay network, and IP networks allowing communications among devices attached to any of them. In addition, transmission of still images, using fax, or moving images, using video, once performed only by analog transmission methods can now be done digitally. Digital transmission of voice is possible by using a digitizing method known as pulse code modulation, or PCM. In PCM, an analog signal is typically converted to a 64kbps digital signal, using 8000 samples per second and eight bits per sample.

High-bandwidth digital transmission facilities are playing a major role in communications today. The T-1 standard, developed by AT&T, provides a 1.544 Mbps 24 path now available from almost all of the common carriers. T-1 carrier can provide 24 channels, each able to carry 64 kbps of digitized voice or 56kbps of data over two twisted pairs; analog transmission of 24 voice channels would require 24 pairs of wires.

Even with the added expense of repeaters and digital transmission equipment, T-1 is an economical alternative for users with high-volume communications traffic. There are several different methods for carrying signaling information using T-1 carrier. The T-1 channels can also be broken down into several lower speeds channels by submultiplexing. In addition, voice and data compression can further increase the effective bandwidth of T-1 circuits. Finally, there are higher speed digital transmission facilities than T-1, such as T-1C and T-3, providing 48 and 672 channels, respectively.

In ATM, the bits transmitted synchronously, although the cells may not always arrive at a fixed rate. Synchronous Optical Network, SONET, is an ANSI standard commonly discussed as a transmission method for ATM. Line speeds in SONET are usually designated by the term Optical Carrier-N or OC-N.

ATM establishes the route for cells at the beginning of each conversation. When cells arrive at an ATM switch, the switching can be performed mostly by hardware, since the routing decisions has been made in advance and stored in the table residing in each switch. Based on the header of the incoming cell, the ATM switch looks up the correct port on which to transmit the cell on the next leg of its journey.

ATM provides service at least at layers 1 and 2 of the OSI model, through some higher layer activities, such as routing, are included. The ATM standard itself describes its functions in three layers: the Physical Layer, the ATM Layer, and the ATM Adaptation Layer, or AAL. A cell relay network such as ATM presents a variety of capacity, traffic, and tariff issues far different from those encountered with a circuit-switched network. ATM and IP networks will continue to coexist for years to come.

References

1. "ATM User Network Interface Specification Version 3.0", ATM Forum, PTR Prentice Hall, September 10, 1993.
2. "Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements", Bellcore, Issue 1, August 1992.
3. "Synchronous Optical Networks (SONET) Transport System: Common Generic Critaria", Bellcore, Issue 2, December 1991.
4. "Broadband ISDN ATM Layer Functionality and Specification", January 1993.
5. "Fiber Channel Physical Layer (FC-PH)", Revision 2.1, Working Draft, May 25, 1991.
6. "Broadband ISDN Transport Network Elements Framework Generic Criteria", December 1990.
7. "Broadband ISDN Switching System Generic Requirements", August 1992.
8. "B-ISDN User Network Interface-Physical Layer Specification", June 1990.
9. "OAM principles of B-ISDN access", Geneva, June 1992.
10. "General Aspects of Quality of Service and Network Performance in Digital Networks, Including ISDN".
11. "Broadband ISDN Performance".
12. "ATM&MPEG 2" Integrating digital video into broadband networks
13. "ATM" Foundation for broadband networks UYLESS BLACK
14. "ISDN explained" John M. Griffith