Near East University

Engineering Faculty

Computer Engineering Department

COM 400

GRADUATION PROJECT

ATM NETWORKING

Submitted by : 940153, Burak Özgören

Submitted to : Prof. Fakhreddin Mamedov

INDEX

PART 1

PART 2

PART 3

PART 1 - STANDARDS AND SPECIFICATIONS

This part summarises the standards and specifications that have been approved by the CCITT/ITU-T, ANSI, the ATM Forum, the IETF, and the Frame Relay Forum as of early 1994.

CCITT/ITU-T Standards

The current ITU-T standards relating to B-ISDN are listed in Table 1. Starting in the 1988, the CCITT published Recommendations I.113 and I.121 which defined vocabulary, terms, principles, and basic objectives for broadband aspects of ISDN, called B-ISDN. These recommendations were revised and approved in November 1990, and published in 1991. Also in 1991, eleven more recommendations I.150, I.211, I.311, I.321, I.327, I.361, I.362, I.363, I.413, I.432, and I.610 were published, further detailing the functions, service aspects, protocol layer functions, Operations, Administration, and Maintenance (OAM), and user-to-network and network-to-network interfaces. In 1992, the following additional recommendations were approved: I.364, I.371, and I.580. In 1993, I.113 and I.321 were revised, and new Recommendations I.356, a new section of I.363 for AAL5, I.365, and I.555 were approved. A number of signalling B-ISDN signalling standards is up for approval in 1994. A brief description of each standard follows:

Table 1. CITT/ITU-T B-ISDN Standards

Number	*Title*
I.113	Vocabulary for B-ISDN
I.121	Broadband aspects of ISDN
I.150	B-ISDN Asynchronous Transfer Mode Functional Characteristics
I.211	General Service Aspects of B-ISDN
I.311	B-ISDN General Network aspects
I.321	B-ISDN Protocol Reference Model and Its Application
I.327	B-ISDN Functional Architecture
I.350	General Aspects of Quality of Service and Network Performance in Digital Networks, including ISDN
I.356	B-ISDN ATM Layer Cell Transfer Performance
I.361	B-ISDN ATM Layer Specification
I.362	B-ISDN ATM Adaptation Layer (AAL) Functional Description
I.363	B-ISDN ATM Adaptation Layer (AAL) Specification

I.364	Support of Connectionless Data Service on a B-ISDN
I.365.1	Frame Relaying Bearer Service Specific Convergence Sub-layer (FR-SSCS)
I.371	Traffic Control and Congestion Control in B-ISDN
I.413	B-ISDN User-Network Interface
I.432	B-ISDN User-Network Interface - Physical Layer Specification
I.555	Frame Relay Bearer Service Interworking
I.580	General Arrangements for Interworking between B-ISDN and 64 kb/s ISDN
I.610	B-ISDN OAM Principles and Functions
G.804	ATM Cell Mapping Into Plesiochronous Digital Hierarchy (PDH)

Recommendation I.113, Vocabulary for B-ISDN - is a glossary of terms and acronyms used in B-ISDN and ATM.

Recommendation I.121 CCITT, Broadband Aspects of ISDN - defines basic principles and characteristics of B-ISDN, and how it can be evolved from TDM and ISDN networks.

Recommendation 1.150 CCITT, B-ISDN Asynchronous Transfer Mode Functional Characteristics - defines functional characteristics of ATM, such as the establishment of and signalling for virtual paths and channels, cell level multiplexing, per virtual connection Quality of Service (QoS) and Generic Flow Control (GFC).

Recommendation 1.211 CCITT, General Service Aspects of B-ISDN - defines interactive and distribution service classes, the types of information needing support, example applications, and some possible attributes such as bit rate, QoS, synchronisation, responsiveness, and source characteristics.

Recommendation 1.311 CCITT, B-ISDN General Network Aspects - peals back the multilayered ATM onion and begins to detail the concepts behind ATM, such as the physical and ATM sublayers and the way Virtual Path (VP) and Virtual Channel (VC) connections are constructed from smaller links, and defines the notions of VC switching and VP cross-connects using a number of examples. It then covers the control and management of B-ISDN, including the physical network management architecture and general principles of signalling.

***Recommendation 1.321 CCITT, B-ISDN Protocol Reference Model and Its Application* -** defines the layered model that is used as the road.

***Recommendation 1.327 CCITT, B-ISDN Functional Architecture* -** defines a basic architectural model for B-ISDN, and how it relates to ISDN and connectionless services.

***Recommendation 1.350 CCITT, General Aspects of Quality of Service and Network Performance in Digital Networks, including ISDN* -** defines the terms Quality of Service (QoS) as the user's perception and Network Performance (NP) as the network operator's observation. It defines specific performance parameters in terms of a number of generic categories.

***Recommendation 1.356 CCITT, B-ISDN ATM Layer Cell Transfer Performance* -** defines the reference events, the definitions, and how the detailed ATM layer performance parameters identified in I.350 can be theoretically calculated.

***Recommendation 1.361 CCITT, B-ISDN ATM Layer Specification* -** defines the bits and bytes of the ATM cell format, what they mean, and how they are to be used.

***Recommendation 1.362 CCITT, B-ISDN ATM Adaptation Layer (AAL) Functional Description* -** defines basic principles and sublayering. It also defines service classification in terms of constant or variable bit rate, timing transfer requirement, and whether the service is connection-oriented or connectionless.

***Recommendation 1.363 CCITT, B-ISDN ATM Adaptation Layer (AAL) Specification* -** defines the specifics of the AALs 1, 2, 3/4, and 5. AAL1 is for connection-oriented, continuous bit rate service that requires timing transfer. AAL2 is for connection-oriented, variable bit-rate service that requires timing transfer. AAL3/4 and AAL5 can be used for connection-oriented or connectionless, variable bit-rate service that does not require timing transfer.

***Recommendation 1.365.1 CCITT, Frame Relaying Bearer Service Specific Convergence Sublayer (FR-SSCS)* -** defines the specifics for interworking frame relay and ATM.

***Recommendation 1.364 CCITT, Support of Connectionless Data Service on a B-ISDN* -** defines an approach for support of connectionless services, such as 802.6/DQDB, over AAL3/4.

Recommendation 1.371 CCITT, Traffic Control and Congestion Control in B-ISDN - defines terminology for traffic parameters, a traffic contract, conformance checking, resource management, connection admission control, prioritization, and implementation tolerances.

Recommendation 1.413 CCITT, B-ISDN User-Network Interface - defines the reference configurations and terminology used in the B-ISDN standards.

Recommendation 1.432 CCITT, B-ISDN User-Network Interface - Physical Layer Specification - defines the details of how ATM cells are mapped into the Synchronous Digital Hierarchy (SDH) TDM structure, how the ATM Header Error Control (HEC) is generated, and how bit errors impact HEC and cell delineation time.

Recommendation 1.555 CCITT, Frame Relay Bearer Service Interworking - defines how frame relay interworks with a number of other services, including B-ISDN.

Recommendation 1.580 CCITT, General Arrangements for Interworking between B-ISDN and 64 kb/s ISDN - defines in general terms how the narrowband ISDN can be interworked with the Broadband-ISDN in support of user data transfer, control and management.

Recommendation 1.610 CCITT, B-ISDN OAM Principles and Functions - covers the initial principals and functions required for Operation, Administration, and Maintenance (OAM) of primarily the ATM layer.

Recommendation G.804, ATM Cell Mapping Into Plesiochronous Digital Hierarchy (PDH) - defines how ATM cells are mapped into various TDM structures, such as E1, D51, E3, and D53.

ANSI Standards

ANSI committee T1 adapts CCITT/ITU-T standards to the competitive environment and the unique physical layer transmission requirements of North America. The standards approved to date are listed and briefly described below.

T1.624-1993, BISDN UNI: Rates and Formats Specification - defines the mapping's of ATM cells into DS3 and SONET payloads and how fault management is performed.

***T1.627-1993, BISDN ATM Functionality and Specification* -** defines the ATM layer following I.361, adding additional explanations of the protocol model, further interpretations of traffic management, and some examples describing functions required in an implementation as annexes.

***T1.629-1993, BISDN ATM Adaptation Layer 314 Common Part Functionality and Specification* -** defines the AAL3/4 functionality based on I.363, expanding on the protocol model and giving an example state machine in an annex.

***T1.630-1993, BISDN - Adaptation Layer for Constant Bit Rate Services Functionality and Specification* -** defines AAL1 based upon I.363, but in considerably more detail through more explanatories, detailed requirements and a number of good technical annexes. T1.630 also defines the specifics of emulating the North American DS1 circuit function, interface, and alarm management.

***T1.633, Frame Relay Bearer Service Interworking* -** is based very closely on recommendation I.555.

***T1.634, Frame Relay Service Specific Convergence Sublayer (FR - SSCS)* -** is based upon a draft CCITT recommendation that will likely be part of I.365 in the future.

***T1.635, BISDN ATM Adaptation Layer Type 5* -** defines AAL5 based precisely upon I.363.

ATM Forum Specifications

The ATM Forum has produced several implementation specifications that are summarised below.

***ATM User-Network Interface (UNI) Specification Version 2.0* -** defined a PVC capability for the ATM UNI, added physical layers based upon FDDI and fiber channel technology for the local area, defined an SNMP-based Interim Local Management Interface (ILMI), and adopted a subset of standardised ATM OAM functions.

***ATM User Network Interface (UNI) Specification Version 3.0* -** supersedes version 2.0, correcting errors and adding new functions. The major new functions were definition of an initial signalling protocol defined as a subset of the ITU-T standard, definition of traffic control beyond the peak rate in an unambiguous manner, and the specification of scrambling for the DS3 rate in order to allow operation over current transmission systems based upon implementation experience.

ATM Data eXchange Interface (DXI) Specification Version 1.0 - defines a frame-based interface that allows a DTE to pass the ATM VPI/VCI in the frame address. It is similar in nature to the SMDS DXI specification.

ATM Broadband-InterCarrier Interface (B-ICI) Specification Version 1.0 - defined characteristics of service for PVC connection between carrier networks. The OAM functions, frame relay network interworking, and transport of SMDS-ICI over ATM are defined in great detail.

IETF RFCs Related to ATM

The IETF has completed several documents to date in support of ATM, with several others in progress.

RFC 1483 MultiProtocol Encapsulation Over ATM - defines how higher layer protocols, such as IP, are encapsulated for bridging and routing over an ATM network. Interworking at the protocol encapsulation level with frame relay networks is also defined.

RFC 1577 Classical IP Over ATM - defines how the current Internet Protocol (IP) can utilise the ATM Switched Virtual Connection (SVC) capability.

Future

The ITU-T, ANSI, and ETSI continue to refine and expand upon the set of standards introduced earlier. Significant activity is occurring in the areas of the control and management planes. The ATM Forum has announced that it planned to work on the following nine areas, starting in the fall of 1993. Some results from these ATM Forum activities are expected in 1994 and 1995, even though specific documents and schedules have not been announced. This is a very ambitious charter, and it will likely take years to complete specifications in all of these areas.

- Signalling
- Traffic Management
- System Aspects and Applications (SAA)

- SMDS Access over ATM
- Frame Relay Interworking
- Video Support
- Circuit Emulation
- Application Program Interface (API)
- LAN Emulation
- Testing and Interoperability
- Private Network-Network Interface (NNI)
- Broadband InterCarrier Interface (B-ICI)
- Physical Layer
- Network Management

The Frame Relay Forum and ATM Forum have announced that they will define service interworking with ATM and address signalling interworking in 1994.
The SMDS Interest Group (SIG) and ATM Forum have announced a joint effort to specify how users can access SMDS service over ATM in 1994.
The IETF has several activities ongoing in the area of ATM, including how routing should be performed, what initial steps can be taken to support classical IP over ATM, and definition of how the IP Maximum Transfer Unit (MTU) can be negotiated.

PART 2 - INTRODUCTION TO ATM

This part introduces the reader to the basic principles of ATM. ATM in a most basic sense is a technology, defined by protocols standardised by the ITU-T, ANSI, ETSI, and the ATM Forum introduced in the previous part. The coverage begins with the building blocks of ATM - transmission paths, virtual paths, and virtual channels. Next a look is taken at the ATM cell and its transmission through a series of simple examples. The fact that the 53-octet cell size turned out to be a compromise between a smaller cell size optimised for voice and a larger cell size optimised for data is presented. This part concludes with a discussion of how ATM means many things to many people, such as a method of integrated access, a public virtual data service, hardware and software implementation, or a network infrastructure.

OBJECTIVES OF ATM

Asynchronous Transfer Mode (ATM) is a cell-based switching and multiplexing technology designed to be a general-purpose, connection-oriented transfer mode for a wide range of services. ATM is also being applied to the LAN and private network technologies as specified by the ATM Forum. ATM handles both connection-oriented traffic directly or through adaptation layers, or connectionless traffic through the use of adaptation layers. ATM virtual connections may operate at either a Constant Bit Rate (CBR) or a Variable Bit Rate (VBR). Each ATM cell sent into the network contains addressing information that establishes a virtual connection from origination to destination. All cells are then transferred, in sequence, over this virtual connection. ATM provides either Permanent or Switched Virtual Connections (PVCs or SVCs). ATM is asynchronous because the transmitted cells need not be periodic as time slots of data are in Synchronous Transfer Mode (STM). ATM offers the potential to standardise on one network architecture defining the multiplexing and switching method, with SONET/STM providing the basis for the physical transmission standard for very high-speed rates. ATM also supports multiple Quality of Service (QoS) classes for differing application requirements on delay and loss performance. Thus, the vision of ATM is that an entire network can be constructed using ATM and ATM Application Layers (AALs) switching and multiplexing principles to support a wide range of all services, such as:

- Voice
- Packet data (SMDS, IP, FR)
- Video
- Imaging
- Circuit emulation

ATM provides bandwidth-on-demand through the use of SVCs, and also support LAN-like access to available bandwidth.

THE ATM CELL AND TRANSMISSION

The primary unit in ATM is the cell. This section defines the basics of the ATM cell.

ATM Cell

ATM standards define a fixed-size cell with a length of 53 octets (or bytes) comprised of a 5-octet header and a 48-octet payload as shown in Figure 1. The bits in the cells are transmitted over the transmission path from left to right in a continuous stream. Cells are mapped into a physical transmission path, such as the North American DSI, D53, or SONET; European, E1, E3 and E4; or ITLT-T STM standards; and various local fiber and electrical transmission payloads. This is only a brief introduction to the ATM cell format.

All information is switched and multiplexed in an ATM network in these fixed-length cells. The cell header identifies the destination, cell type, and priority. The Virtual Path Identifier (VPI) and Virtual Channel Identifier (VCI) hold local significance only, and identify the destination. The Generic Flow Control (GFC) field allows a multiplexer to control the rate of an ATM terminal. The Payload Type (PT) indicates whether the cell contains user data, signalling data, or maintenance information: The Cell Loss Priority (CLP) bit indicates the relative priority of the cell. Lower priority cells are discarded before higher priority cells during congested intervals.
Because of its critical nature, the cell Header Error Check (HEC) detects and corrects errors in the header. The payload field is passed through the network intact, with no error checking or correction. ATM relies on higher layer protocols to perform error checking and correction on the

Figure 1. ATM Cell Transmission and Format

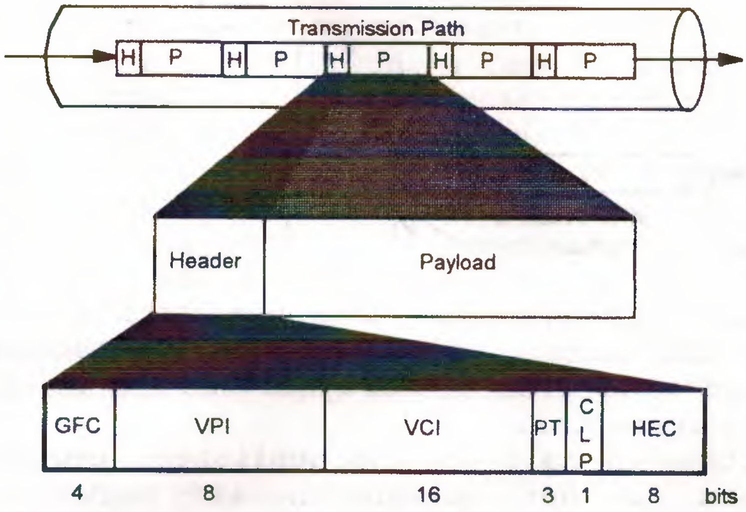

GFC = Generic Flow Control
VPI = Virtual Path Identifier
VCI = Virtual Channel Identifier
PT = Payload Type
CLP = Call Loss Priority
HEC = Header Error Check

payload. The fixed cell size simplifies the implementation of ATM switches and multiplexers and enables implementations at very high speeds.

When using ATM, longer packets cannot delay shorter packets as in other packet switched implementations because long packets are chopped up into many cells. This enables ATM to carry Constant Bit Rate (CBR) traffic such as voice and video in conjunction with Variable Bit-Rate (VBR) data traffic, potentially having very long packets within the same network.

Cell Segmentation Example

ATM switches take a user's data, voice, and video and chops it up into fixed length cells, and multiplexes it into a single bit stream that is transmitted across a physical medium. An example of multimedia application is that of a person needing to send an important manuscript for a book to his or her publisher. Along with the letter, this person would like to show his or her joy at receiving a contract to publish the book.

Figure 2 illustrates the role of ATM in this real-life example, where Jeanne is sitting at her workstation. Jeanne's workstation has an ATM interface card, soundboard

Figure 2. Multimedia Communications Example Using ATM

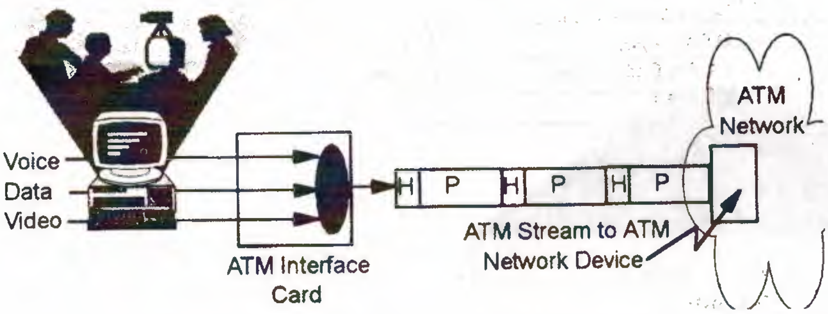

with microphone, and videocamera. The workstation is connected to a local ATM switch, which in turn is attached to a public ATM-based wide area network service to which the publisher is also connected.

Jeanne places a multimedia call to the publisher, begins transmitting the data for her manuscript, and begins a conversation with the publisher, with Jeanne and the publisher able to see each other's faces - providing text, voice, and video traffic, respectively, in real time. The publisher is looking through the manuscript at its workstation all the while and having an interactive dialogue with Jeanne. Let's take this scenario one piece at a time.

Figure 3. Virtual Channels Supporting Multiple Applications

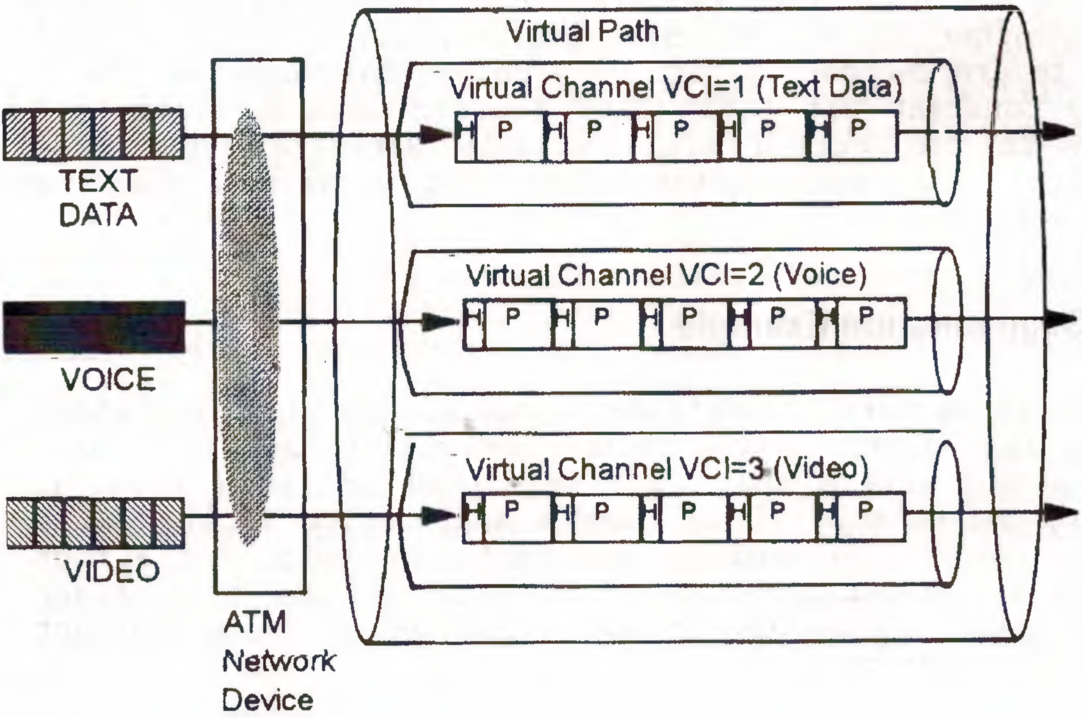

Video and voice are very time-sensitive; the information cannot be delayed for more than a blink of the eye, and the delay cannot have significant variations. Disruptions in the video image of Jeanne's face or distortion of the voice destroy the interactive, near real-life quality of this multimedia application. Data can be sent in either connection-oriented or connectionless mode. In either case, the data is not nearly as delay-sensitive as voice or video traffic. Data traffic, however, is very sensitive to loss. Therefore, ATM must discriminate between voice, video, and data traffic, giving voice and video traffic priority and guaranteed, bounded delay, simultaneously assuring that data traffic has very low loss.

Examining this example in further detail, a virtual path is established between Jeanne and the publisher, and over that virtual path three virtual circuits are defined for text data, voice, and video. Figure 3 shows how all three types of traffic are combined over a single ATM Virtual Path (VP), with Virtual Circuits (VCs) being assigned to the text data (VCI=1), voice (VCI=2), and video (VCI=3).

THEORY OF OPERATION

This section presents two examples of how user traffic is segmented into ATM cells, switched through a network, and processed by the receiving user.

A Simple ATM Example

Let's look at the last example, where Jeanne is simultaneously transmitting text, voice, and video data traffic from her workstation, in more detail. The workstation contains an ATM interface card, where the chopper "slices and dices" the data streams into 48-octet data segments as shown in Figure 4. In the next step the postman addresses the payload by prefixing it with the VPI, VCI, and the remaining fields of the 5-octet header. The result is a stream of 53-octet ATM cells from each source: voice, video, and text data. These cells are generated independently by each source, such that there may be contention for cell slot times on the interface connected to the workstation. The text, voice, and video are each assigned a VCC: VCI=1 for text data, VCI=2 for voice, and VCI=3 for video, all on VPI=0. This example has been greatly simplified, as there would normally be many more than just three active VCI values on a single VPI.

Figure 4. Asynchronous Transfer Mode Example

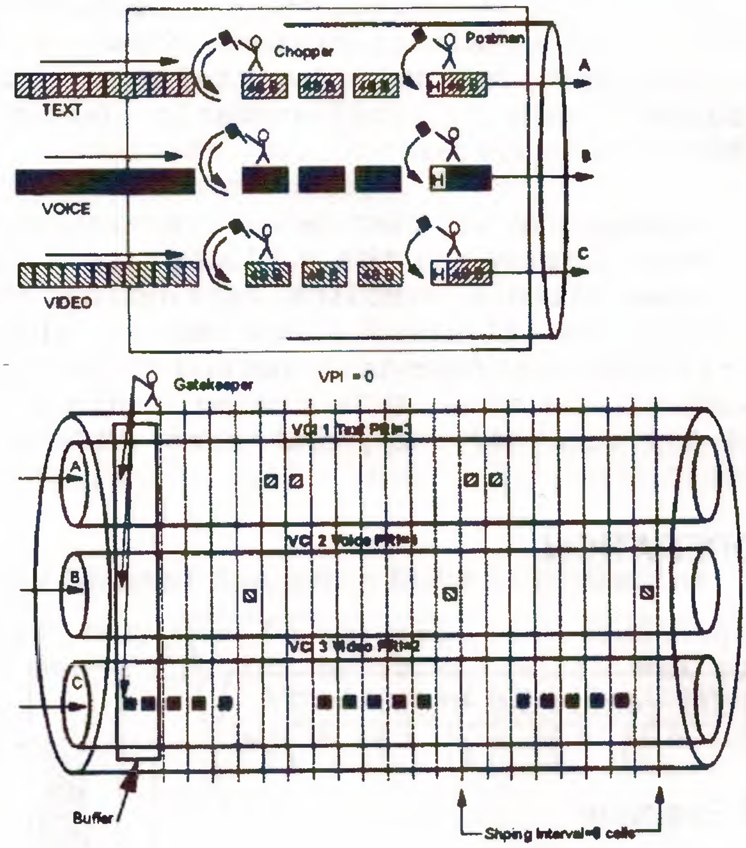

Figure 4 shows an example of how Jeanne's terminal sends the combined voice, video, and text data. A gatekeeper in her terminal shapes the transmitted data in intervals of eight cells (about 80 µs at the DS3 rate), normally allowing one voice cell, then five video cells, and finally what is left - two text data cells - to be transmitted. This corresponds to about 4 Mbps for high-fidelity audio, 24 Mbps for video, and 9 Mbps for text data. All data sources (text, voice, and video) contend for the bandwidth each shaping interval of eight cell times, with the voice, video, and then text data being sent in the above proportion. The gatekeeper retains cells in the buffer in case all of the cell slot times were full in the shaping interval. A much larger shaping interval is used in practice to provide greater granularity in bandwidth allocation.

An ATM Switch Example

An illustration of an ATM switch is shown in Figure 5. A continuous video source is shown as input to a packeting function, with logical destination VPI/VCI address D. The continuous bit stream is broken up into fixed-length cells comprised of a header and a payload field (indicated by the shading). The rate of the video source is greater than the continuous DS3 bit stream with logical destination address A, and the high-speed computer directly packeted input addressed to B. These sources are shown time division multiplexed over a transmission path, such as SONET or DS3.

Figure 5. Asynchronous Transfer Mode Example

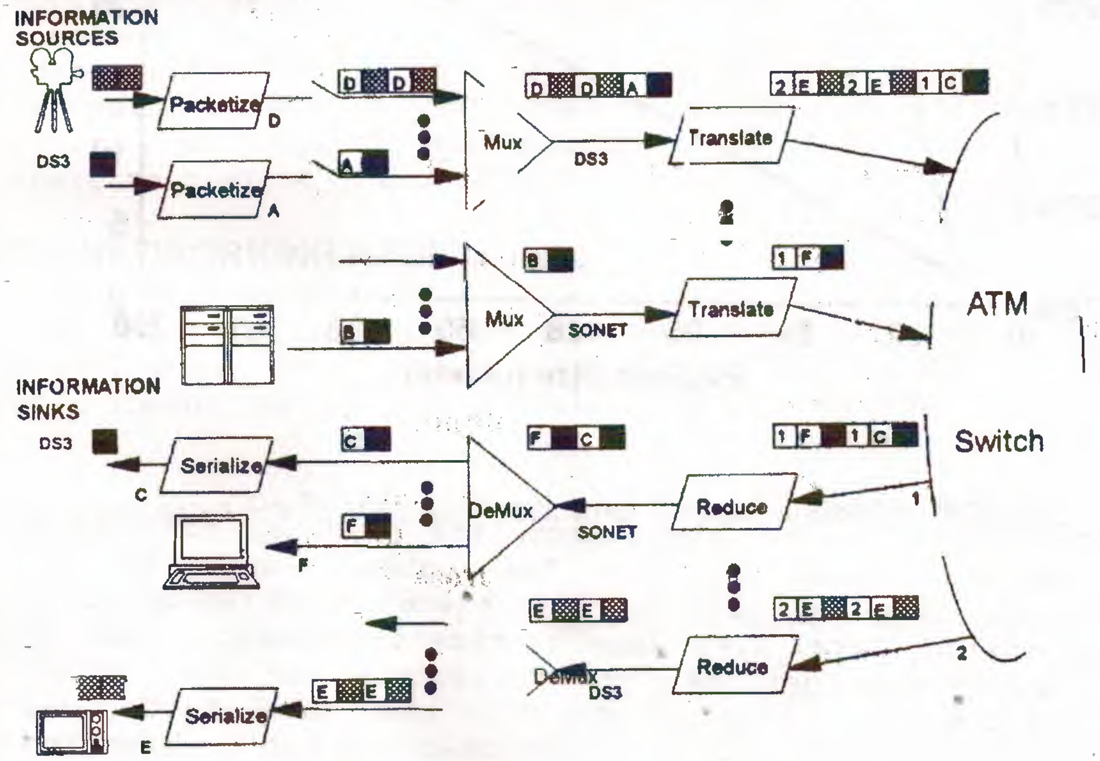

The initial function of the ATM switch is to translate the logical address to a physical outgoing switch port address and to an outgoing logical VPI/VCI address. This additional ATM switch header is prefixed to every input ATM cell as shown previously. There are three point-to-point virtual connections in the figure. The DS3 has address that is translated into C destined for physical port 1. The video source has address D that is translated into address E destined for port 2. The computer source has address B that is translated to address F destined for port 1.

The ATM switch utilises the physical destination address field to deliver the ATM cells to appropriate physical switch port and associated transmission link.

Figure 6. Delay versus Cell Size Trade off

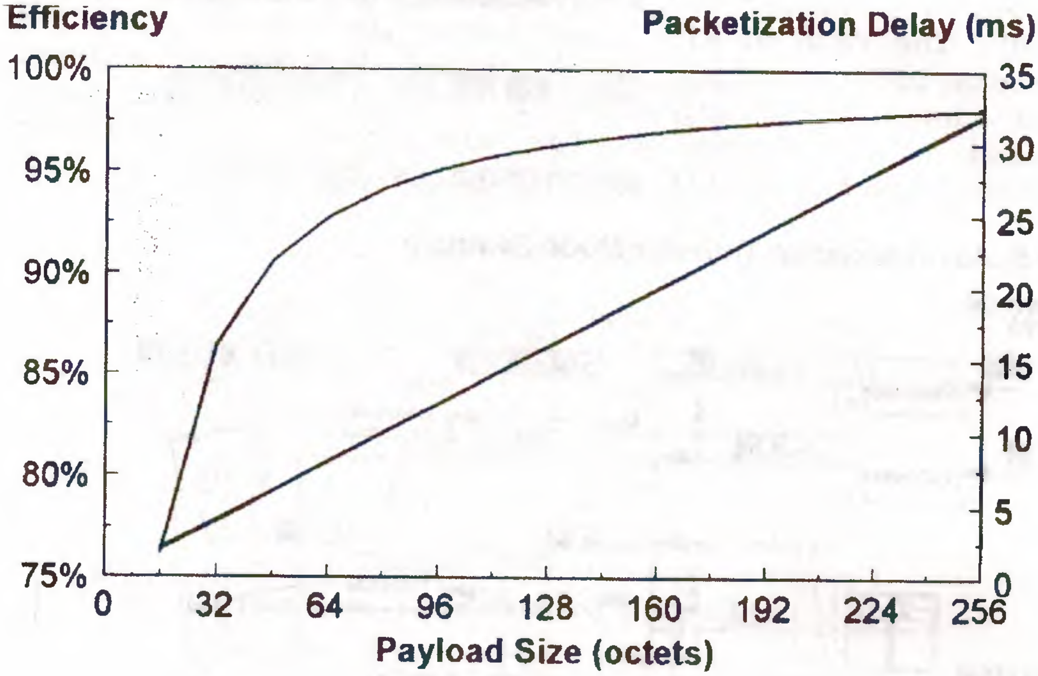

At the output of the ATM switch, the physical address is removed by a reduce function. The logically addressed ATM cells are then time division multiplexed onto the outgoing transmission links. Next these streams are demultiplexed to the appropriate devices. The Continuous Bit Rate (CBR) connections (ie, video and the DS3) then have the logical addresses removed, and are reclocked to the information sink via the serialise function. Devices, such as workstation, can receive ATM cells directly.

CHOICE OF PAYLOAD SIZE

When a standard cell size was under discussion by the ATM Forum, there was a raging debate between a 32-octet versus a 64-octet payload size. The decision on the 48-byte payload size was the compromise between these positions. The choice of the 5-octet header size was a separate trade off between a 3-octet header and an 8-octet header.

There is a basic trade off between efficiency and packetization delay versus cell size illustrated in Figure 6. Efficiency is computed for a 5-octet cell header. Packetization delay is the amount of time required to fill the cell at a rate of 64 kbps, that is, the rate to fill the cell with digitised voice samples. Ideally high efficiency and low delay are both desirable, but cannot be achieved simultaneously. As seen from the figure, better efficiency occurs at large cell sizes at the expense of increased packetization delay. In order to carry voice over ATM and interwork with two-wire analog telephone sets, the total delay should be less than about 12 ms, otherwise echo cancellation must be used. Two TDM to ATM conversions are required in the round-trip echo path. Allowing 4 ms for propagation delay and two ATM conversions, a cell size of 32 octets avoids the need for echo cancellation. Thus, the ITU-T adopted the fixed-length 48-octet cell payload as a compromise between a long cell size for time-insensitive traffic (64 octets) and smaller cell sizes for time-sensitive traffic (32 octets).

ATM NETWORKING BASICS

Three major concepts in ATM are the transmission path, the Virtual Path (VP), and, optionally, the Virtual Channel (VC). These form the basic building blocks of ATM.

Transmission Path, Virtual Path, and Virtual Channel Analogy

Let us look at a simple example of these concepts in relation to vehicle traffic patterns. These analogies are not intended to be exact. Think of cells as vehicles, transmission paths as roads, virtual paths as a set of directions, and virtual channels as a lane discipline on the route defined by the virtual path. Figure 7 illustrates the example described in this section.

Three transmission paths form the set of roads between three cities: Dallas, Fort Worth, and Houston. There are many interstates, highways, and back roads between the two cities which create many possibilities for different routes - but the primary routes, or virtual paths, are the interstate (VP1) from Dallas to Houston, the highway from

Dallas to Fort Worth (VP2), and a back road (VP3) from Fort Worth to Houston. Thus, a car (cell) can travel from Dallas to Houston either over the highway to Fort Worth and then the back road to Houston, or take the direct interstate. If the car chooses the interstate (VP 1), it has the choice of three lanes: car pool or High Occupancy Vehicle (HOV) (VCCI), car lane (VCC2), or the truck lane (VCC3). These three lanes have speed limits of 65 mph, 55 mph, and 45 mph, respectively, which will cause different amounts of delay in reaching the destination. In our analogy, vehicles strictly obey this lane discipline (unlike on real highways).

Figure 7. Transportation Example of ATM Principles

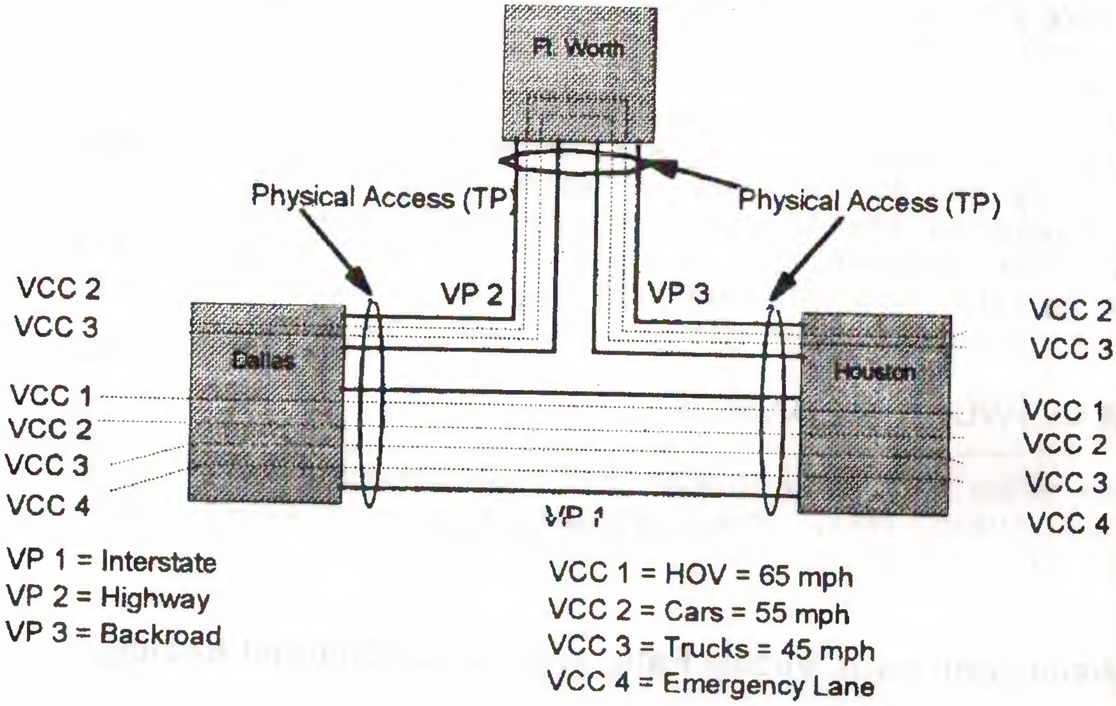

In our example, the interstate carries high-speed traffic: tractor trailers, buses, tourists, and business commuters. The highway can carry car and truck traffic, but at a lower speed. The back roads carry locals and traffic avoiding backups on the interstate (spillover traffic), but at an even slower speed.

Note that our example of automotive traffic (cells) has many opportunities for missequencing. Vehicles may decide to pass each other, there can be detours, and road hazards (like stalled cars in Texas!) may cause some vehicles (cells) to arrive out of sequence or vary in their delay.

This is evident in normal transportation when you always seem to leave on time, but traffic causes you to be delayed. Automotive traffic must employ an Orwellian discipline where everyone follows the traffic routes exactly (unlike any real traffic) in order for the analogy to apply.

The routes also have different quality. When you get a route map from the American Automobile Association (AAA), you have a route selected based on many criteria: least driving (routing) time, most scenic route, least cost (avoids most toll roads), and avoid known busy hours. The same principles apply to ATM.

Figure 8. Transportation Example – STM Versus ATM

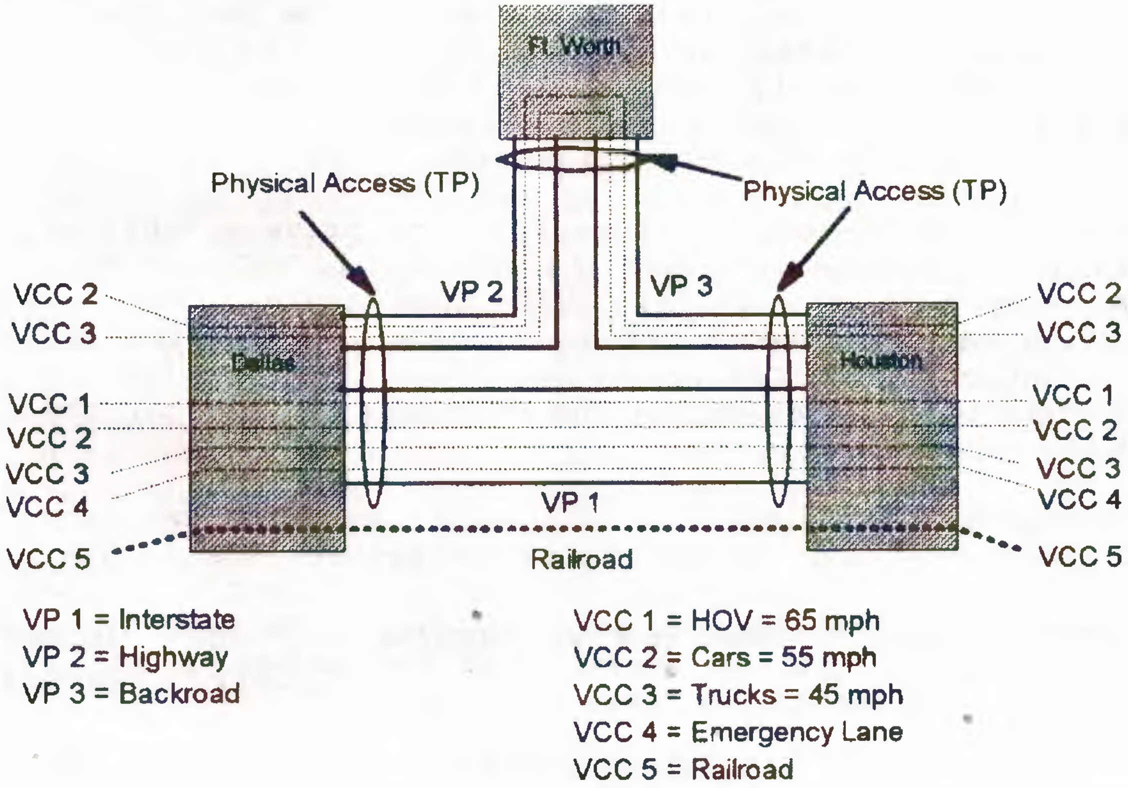

Now, let's give each of the road types (VPs) and lanes (VCCs) a route choice. A commuter from Dallas to Houston in a hurry would first choose the VPI, the interstate. A sightseer would choose the highway to Fort Worth (VP2) to see the old cow town, and then the back road to Houston

(VP3) to take in Waxahachie and Waco on the way. When commuters enter the interstate toward Houston, they immediately enter the HOV lane (VCC 1) and speed toward their destination.

Figure 8 adds a railroad (VCC5) running from Dallas to Houston along the same interstate route (VP1) in the previous example. Assuming no stops between Dallas and Houston, the railroad maintains the same speed from start to finish, with one railroad train running after another according to a fixed schedule. This is like the Synchronous Transfer Mode (STM) or Time Division Multiplexing (TDM). Imagine there are passengers and cargo going between Dallas and Houston, each having to catch scheduled trains. The arriving passengers and cargo shipments originating at Dallas must wait for the next train. Trains travel regardless of whether there is any passenger or cargo present. If there are too many passengers or cargo for the train's capacity, the excess must wait for the next train. If you were a commuter, would you wait to rely on the train always having capacity, or would you prefer to have a car and statistically have a better chance of making it to Houston in an even shorter time period using ATM?

Studying this analogy, observe that the private vehicles (and their passengers) travelling over VCC 1, VCC2, or VCC3 have much more flexibility (ATM) than trains (STM) in handling the spontaneous needs of travel. The trains are efficient only when the demand is accurately scheduled and very directed such as during the rush hour between suburbs and the inner city.

Note that the priorities, or choice, of each VCC can vary throughout the day, as can priorities between VPs in ATM. An additional VCC can be configured on a moment's notice (VCC) and assigned a higher priority, as in the case of an ambulance attempting to travel down the median during a traffic jam to get to the scene of an accident.

Transmission Path, Virtual Path, and Virtual Channels

Bringing our last analogy forward into ATM transmission terms, Figure 9 depicts graphically the relationship between the physical transmission path, Virtual Path (VP), and Virtual Channel (VC). A transmission path contains one or more virtual paths, while each virtual path contains one or more virtual channels. Thus, multiple virtual channels can be trunked a single virtual path. Switching can be performed on a transmission path, virtual path, or virtual circuit level.

Figure 9. Relationship of VC, VP, and transmission Path

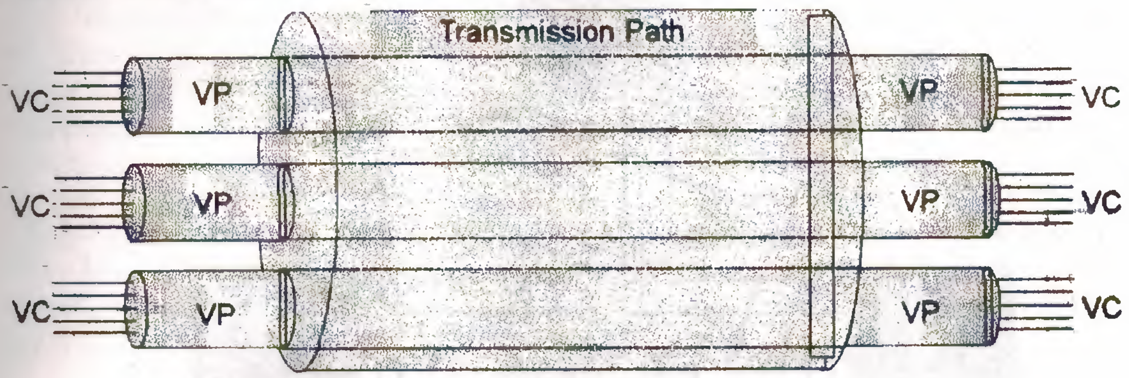

This capability to switch down to a virtual channel level is similar to the operation of a Private or Public Branch eXchange (PBX) or telephone switch in the telephone world. In the PBX/switch, each channel within a trunk group (path) can be switched. Figure 10 illustrates this analogy. In the literature devices which perform VC connections are commonly called VC switches because of this analogy with telephone switches. Transmission networks use a cross-connect, which is basically a space division switch, or effectively an electronic patch panel. ATM devices that connect VPs are commonly often called VP cross-connects in the literature by analogy with the transmission network. These analogies are useful for those familiars with TDM/STM and telephony to understand ATM, but should not be taken literally. There is little reason for an ATM cell-switching machine to restrict switching to only VCs and cross-connection to only VPs.

Virtual Path Connections (VPCs) and Virtual Channel Connections (VCCs)

At the ATM layer, users are provided a choice of either a VPC or a VCC, defined as follows:

Virtual Path Connections (VPCs) are switched based upon the Virtual Path Identifier (VPI) value only. The users of the VPC may assign the VCCs within that VPI transparently since they follow the same route.

Virtual Channel Connections (VCCs) are switched upon the combined VPI and Virtual Channel Identifier (VCI) value.

Figure 10. Switch and Cross-Connect Analogy

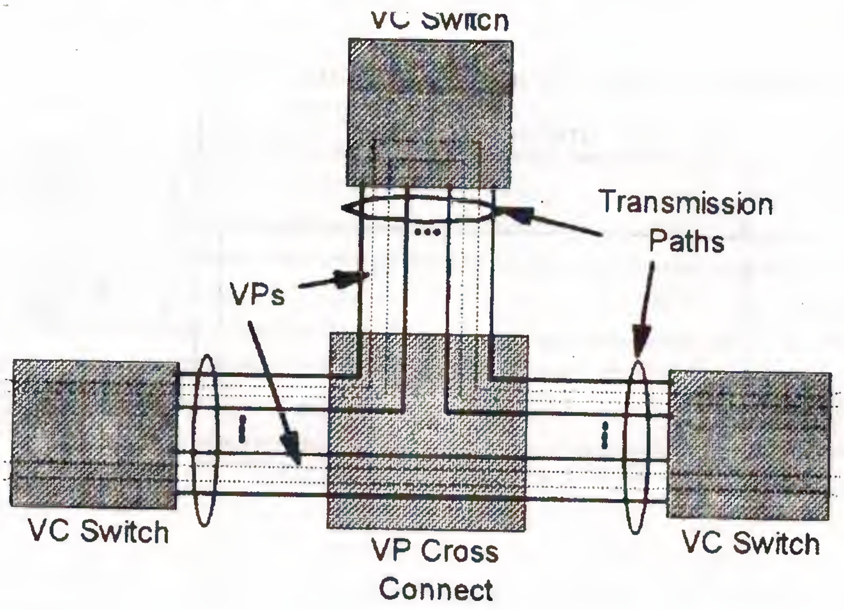

Figure 11. Illustration of VPI/VCI Usage on Link and End-to-End Basis

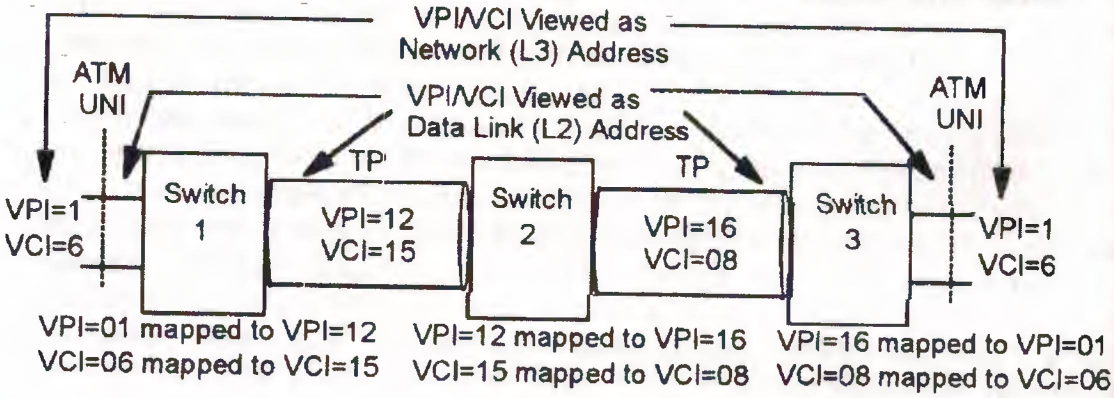

Both VPIs and VCIs are used to route cells through the network. Note that VPI and VCI values must be unique on a specific transmission path (TP). Thus, each TP between two network devices (such as ATM switches) uses VPIs and VCIs independently. This is demonstrated in Figure 11. Each switch maps an incoming VPI and VCI to an outgoing VPI and VCI. In this example, switch 1 and switch 2 have a single transmission path (TP) between them. Over this TP there are multiple virtual paths (VPs). At the ATM UNI, the input device to switch 1 provides a video channel over Virtual Path 1 (VPI 1) and Virtual Channel 6 (VCI 6): Switch 1 then assigns the VCI 6 to an outgoing VCI 15, and the incoming

VPI 1 to outgoing VPI 12. Thus, on VPI 12 switch 2 specifically operates on virtual channel (VC) number 15 (VCI 15). This channel is then routed from switch 2 to switch 3 over a different path and channel (VPI 16 and VCI 8). Thus, VPIs and VCIs are tied onto each individual link across the network. This is similar to frame relay, where Data Link Connection Identifiers (DLCIs) address a virtual circuit (VC) at each end of a link. Finally, switch 3 translates VPI 16 into VPI 1, and VCI 8 on VP 16 to VCI 6 on VP 1 at the destination UNI. The destination VPI and VCI need not be the same as at the origin. The sequence of VPI/VCI translation across the switches can be viewed as a network address in an extrapolation of the OSI layer 3 model.

ATM ARCHITECTURE, TECHNOLOGY, OR SERVICE?

ATM technology takes on many forms and means many different things to different people, from providing software and hardware multiplexing, switching, and cross-connect functions and platforms, to serving as an economical, integrated network access method, to becoming the core of a network infrastructure, to the much-touted ATM service. Let's now explore each.

As an Interface and Protocol

Asynchronous Transfer Mode (ATM) is defined as an interface and protocol designed to switch variable bit-rate and constant bit-rate traffic over a common transmission medium. The entire B-ISDN protocol stack is often referred to as ATM.

As a Technology

ATM is often referred to as a technology, comprised of hardware and software conforming to ATM protocol standards that can provide a multiplexing, cross-connect, and switching function in a network. ATM technology takes the form of a network interface card, multiplexer, cross-connect, or even a full switch.

As Economical, Integrated Access

Public ATM service providers offering ATM-based service are now appearing on the scene, enabling users t capitalise on a basic advantage of ATM - integrated acces to reduce cost. The development of circuit emulatio technology based upon ATM will make this benefit availabl to users who already have a large number of TDM acces lines today. The TDM access lines can be multiplexed ont an E3, DS3, or even SONET access line, leaving larg amounts of bandwidth available for ATM applications a little incremental cost.

As an Infrastructure

Where ATM technology can also have an advantage is as the core of a network infrastructure. ATM hardware an associated software together can provide the backbone technology for an advanced communications network. In fact, many experts view an ATM-based architecture as the future platform for data and eventually voice. ATM also provides a very scalable infrastructure, from the campus environment to the central office. Scalability occurs in the dimensions of interface speed, switch size, network size, and addressing.

As a Service

ATM is not a service, but services can be offered over ATM architecture. The Cell Relay Service (CRS) involves the direct delivery of ATM cells. Other services involve ATM Adaptation Layers, and include private line emulation service as defined using AAL1, variable-rate video as defined using AAL2, Switched Multimegabit Data Service as defined using AAL 3/4, and frame relay as one of the service-specific connection-oriented services defined for AAL 5.

PART 3 - Physical, ATM, and AAL Layers

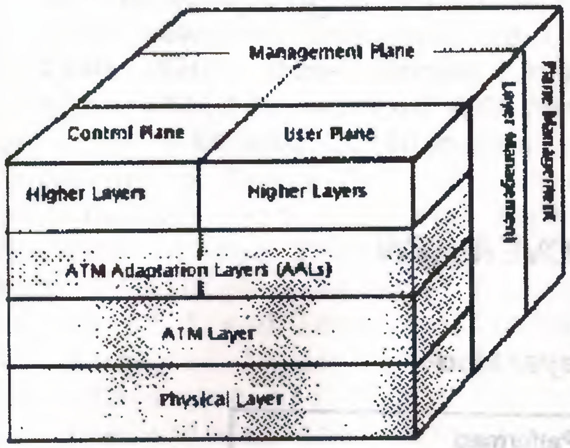

This part explores in detail the foundation of the entire ATM-based B-ISDN protocol stack. The three lowest protocol layers are introduced, first defining what functions they perform and then how they interface. It is logical to start at the bottom with the physical (PHY) layer, and then move to the ATM layer, which defines virtual paths and virtual channels, and finish with the ATM Adaptation Layer (AAL). Many of these concepts were introduced in the last part and are covered in greater detail in this part.

The primary layers of the B-ISDN protocol reference model are: the PHYsical layer, the ATM Layer where the cell structure occurs, and the ATM Adaptation Layer (AAL) that provides support for higher layer services such as circuit emulation, frame relay, and SMDS. The PHY layer corresponds to OSI Reference Model (OSIRM) layer 1, the ATM layer and part of the AAL correspond to OSIRM layer 2, and higher layers correspond to OSI layer 3 and above.

First the description covers the various physical interfaces and media currently defined and specified. A detailed discussion of definitions and concepts of the ATM layer, defining the cell structure for both the User-to-Network Interface (UNI) and the Network Node Interface (NNI), follows. At a lower level, a description of the meanings of the entire cell header fields, payload types, and generic functions that they support is provided. Lastly, the next higher layer in the protocol stack - the ATM Adaptation Layer (AAL) - is covered. An in-depth study of ATM Adaptation Layers (AALs) 1 through 5 relating them to the ITU-T, defined service classes A through D is a provided.

Throughout the remainder of this part will see figures li the one shown. They depict the B-ISDN protocol model fr I.321, with a portion of the B-ISDN/ATM protocol mod shaded out to represent the subject matter of th particular section. This figure shows what this pa covers: the physical layer, the ATM layer, and the AA that are common between the user and control planes.

THE PLANE-LAYER TRUTH - AN OVERVIEW

Figure 12. B-ISDN/ATM Layer and Sublayer Model

Layer Name		Functions Performed	
Higher Layers		Higher Layer Functions	Layer Management
AAL	Convergence Sublayer (CS)	Common Part (CP)	
		Service Specific (SS)	
	SAR Sublayer	Segmentation And Reassembly'	
ATM		Generic Flow Control Call Header Generation/Extraction Cell VCI/VPI Translation Cell Multiplexing/Demultiplexing	
Physical	Transmission Convergence (TC) Sublayer	Cell Rate Decoupling Cell Delineation Transmission Frame Adaptation Transmission Frame Generation/ Recovery	
	Physical Medium (PM)	Bit Timing Physical Medium	

If the front and right sides of the B-ISDN protocol cub were unfolded, they would yield a two-dimensional layere model like that shown in Figure 12.

Figure 12 lists the functions of the four B-ISDN/ATM layer along with the sublayer structure of the AAL and PHYsica (PHY) layer as defined in ITU-T Recommendation I.321 Starting from the bottom, the Physical layer has tw sublayers: Transmission Convergence (TC) and Physica

Medium (PM). The PM sublayer interfaces with the actual physical medium and passes the recovered bit stream to the TC sublayer. The TC sublayer extracts and inserts ATM cells within the Plesiochronous or Synchronous (PDH or SDH) Time Division Multiplexed (TDM) frame and passes these to and from the ATM layer, respectively. The ATM layer performs multiplexing, switching, and control actions based upon information in the ATM cell header and passes cells to, and accepts cells from, the ATM Adaptation Layer (AAL). The AAL has two sublayers: Segmentation And Reassembly (SAR) and Convergence Sublayer (CS). The CS is further broken down into Common Part (CP) and Service-Specific (SS) components. The AAL passes Protocol Data Units (PDUs) to and accepts PDUs from higher layers. PDUs may be of variable length, or may be of fixed length different from the ATM cell length. The Physical layer corresponds to layer 1 in the OSI model. The ATM layer and AAL correspond to parts of OSI layer 2, but the address field of the ATM cell header has a network-wide connotation that is like OSI layer 3. A precise alignment with the OSI layers is not necessary, however. The B-ISDN and ATM protocols and interfaces make extensive use of the OSI concepts of layering and sublayering, as we shall see. Figure 13 illustrates the mapping of the B-ISDN layers to the OSI layers and the sublayers of the PHY, ATM, and ATM adaptation layers that we describe in detail later.

Figure 13. B-ISDN Layers and Sublayers and OSI Layers

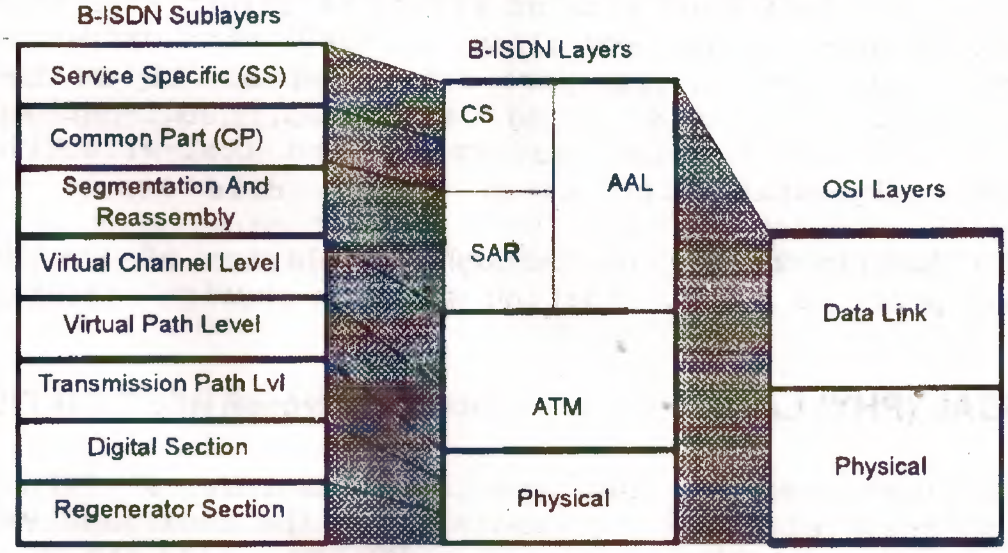

It is interesting to look at the number of instances of defined standardised protocols or interfaces that exist for each layer, and whether their target implementation is in hardware or software. Figure 14 depicts the number of

Figure 14. ATM Protocol Model Hardware to Software Progression

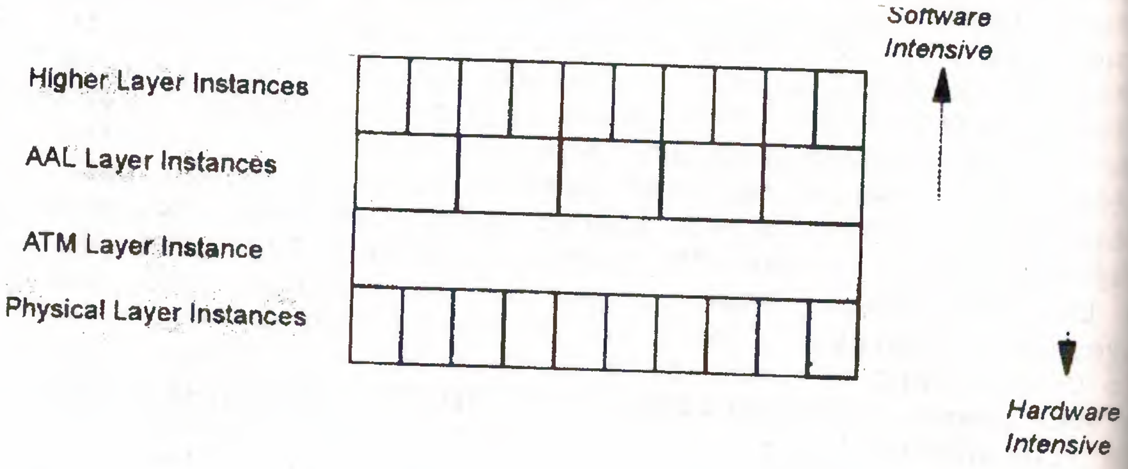

instances at each layer by boxes with the arrows on the right hand side showing how the layers are either more hardware- or software-intensive. The arrows illustrate the fact that ATM implementations move from being hardware-intensive at the lower layers (PHY and ATM layer) to software- intensive at the higher layers (AALs and higher layers). This shows how ATM is the pivotal protocol, for which there is only one instance, for a potentially large number of physical media, several AALs, and an ever-expanding set of higher layer functions. The inverted pyramid on the left-hand side of Figure 14 illustrates this concept. In other words, ATM allows machines with different physical interfaces to transport data independently of the higher layer protocols using a common, well-defined protocol amenable to high performance and cost-effective hardware implementation.

Now the journey begins up through the layers of the B-ISDN/ATM protocol model, starting with the physical layer.

PHYSICAL (PHY) LAYER

This section covers the key aspects of the PHYsical (PHY) Layer as they relate to the remainder of the book. The PHY Layer provides for transmission of ATM cells over a physical medium that connects two ATM devices. The PHY Layer is divided into two sublayers: the Physical Medium Dependent (PMD) sublayer and the Transmission Convergence sublayer. The TC sublayer transforms the flow of cells into a steady flow of bits and bytes for transmission over the physical medium. The PMD sublayer provides for the actual transmission of the bits in the ATM cells.

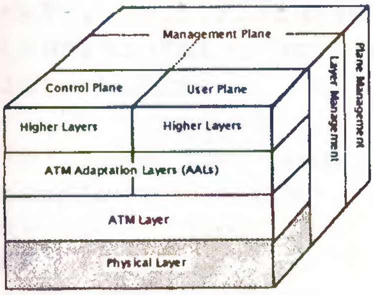

Physical Medium Dependent (PMD) Sublayer

The PMD sublayer provides for the actual clocking of bit transmission over the physical medium. There are three standards bodies that have defined the physical layer in support of ATM: ANSI, CCITT/ITU-T, and he ATM Forum. We summarise each of the standardised interfaces in terms of the interface clocking speed and physical medium below.

ANSI Standards

ANSI Standard T1.624 currently defines three single-mode optical ATM SONET- based interfaces for the ATM UNI:

- STS-1 at 51.84 Mbps
- STS-3c at 155.52 Mbps
- STS-12c at 622.08 Mbps

ANSI T1.624 also defines operation at the DS3 rate of 44.736 Mbps using the Physical Layer Convergence Protocol (PLCP) from the 802.6 Distributed Queue Dual Bus (DQDB) standard.

CCITT/ITU-T SDH Recommendations

CCITT/ITU-T recommendation I.432 defines two optical Synchronous Digital Hierarchy (SDH)-based physical interfaces for ATM which correspond to the ANSI rates mentioned in the last section. These are:

- STM-1 at 155.520 Mbps
- STM-4 at 622.08 Mbps

Since the transport rates (and the payload rates) of SDH STM- 1 and STM-4 correspond exactly to the SONET ST and STS-12c rates, interworking should be simplified. IT standardises additional electrical, physical interf rates of the following type and speeds:

- DS1 at 1.544 Mbps
- E1 at 2.048 Mbps
- DS2 at 6.312 Mbps
- E3 at 34.368 Mbps
- DS3 at 44.736 Mbps using PLCP
- E4 at 139.264 Mbps

ATM Forum Interfaces

The ATM Forum has defined four physical layer interfa rates. Two of these are interface rates intended for publ networks and are the DS3 and STS-3c standardised by AN and the ITU-T. The SONET STS-3c interfaces may be support on an OC-3, either single-mode or multimode fiber. T following three interface rates and media are for priva network application:

- FDDI-based at 100 Mbps
- Fibre Channel-based at 155.52 Mbps
- Shielded Twisted Pair (STP) at 155.52 Mbps

The FDDI-based PMD and fiber channel interfaces both us multimode fiber, while the STP interface uses type 1 and cable as specified by EIA/TIA 568. The ATM Forum i specifying ATM cell transmission over common buildin wiring, called Unshielded Twisted Pair (CTTP) types 3 an 5.

Transmission Convergence (TC) Sublayer

The TC sublayer converts between the bit stream clocked to the physical medium and ATM cells. On transmit, TC basically maps the cells into the Time Division Multiplexing (TDM) frame format. On reception, it must delineate the individual cells in the received bit stream, either from the TDM frame directly, or via the Header Error Check (HEC) in the ATM cell header. Generating the HEC on transmit and using it to correct and detect errors on receive are also important TC functions. Another important function that TC performs is cell rate decoupling by sending idle cells when the ATM layer has not provided a cell. This is a critical function that allows the ATM layer to operate with a wide range of different speed physical interfaces.

We cover two examples of TC mapping of ATM cells: direct mapping to a SONET payload; and the PLCP mapping to a DS3. We then cover the use of the Header Error Check (HEC) and why it is so important. We then complete our description of the TC sublayer with an illustration of cell rate decoupling using unassigned cells.

Examples of TC Mapping

In this section we give an example of direct and Physical Layer Convergence Protocol (PLCP) mapping by the Transmission Convergence (TC) sublayer of the physical layer.

Figure 15. B-ISDN UNI Physical Layer – STS-3c

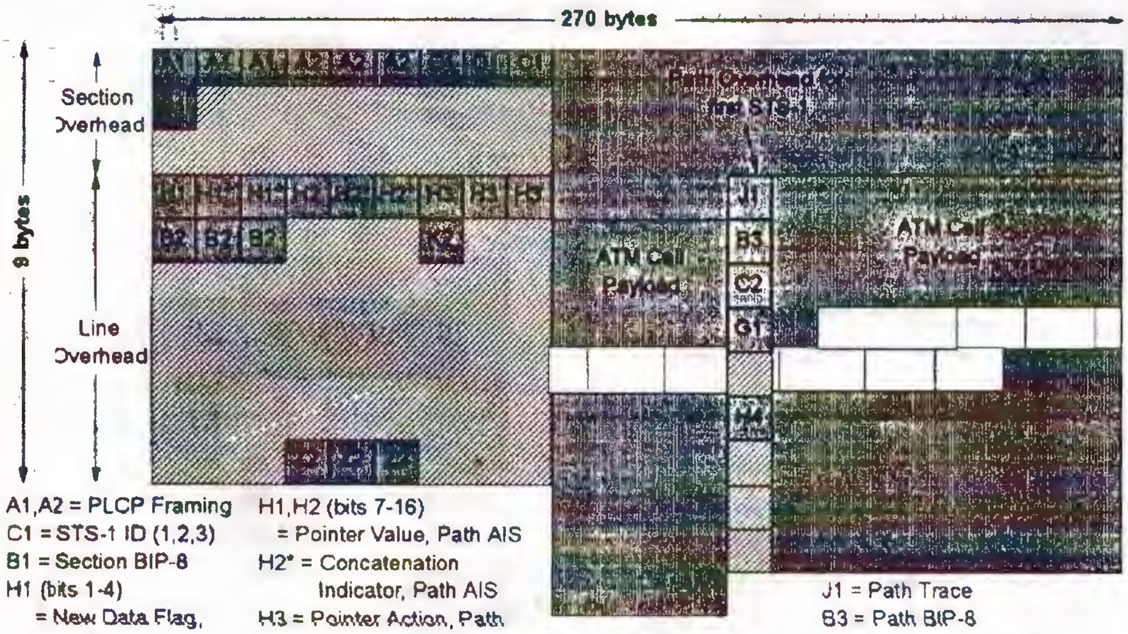

A1,A2 = PLCP Framing
C1 = STS-1 ID (1,2,3)
B1 = Section BIP-8
H1 (bits 1-4) = New Data Flag, Path AIS
H1* = Concatenation Indicator, Path AIS

H1,H2 (bits 7-16) = Pointer Value, Path AIS
H2* = Concatenation Indicator, Path AIS
H3 = Pointer Action, Path AIS
B2 = Line BIP-24
K2 = Line AIS, FERF
Z2 = Line FEBE

J1 = Path Trace
B3 = Path BIP-8
C2 = Path Signal Label
G1 = Path FEBE, RAI, FERF
H4 = Cell Offset Indicator

SONET STS-3c Direct Mapping

The SONET mapping is performed directly into the SONET STS-3c (155.52 Mbps) Synchronous Payload Envelope (SPE) as shown in Figure 15. ATM cells fill in the STS-3c payload continuously since an integer number of 53-octet cells do not fit in an STS-3c frame. This results in better efficiency than carriage of M13-mapped DS3s, or even VT1.5 multiplexing over SONET. The Data Communications Channel (DCC) overhead is not used on the User-Network Interface (UNI). The ATM layer uses the HEC field to delineate cells from within the SONET payload. The cell transfer rate is 149.760 Mbps. The mapping over STS-12c is very similar in nature. The difference between SONET and SDH is in the TDM overhead bytes.

DS3 PLCP Mapping

Figure 16. B-ISDN UNI Physical Layer

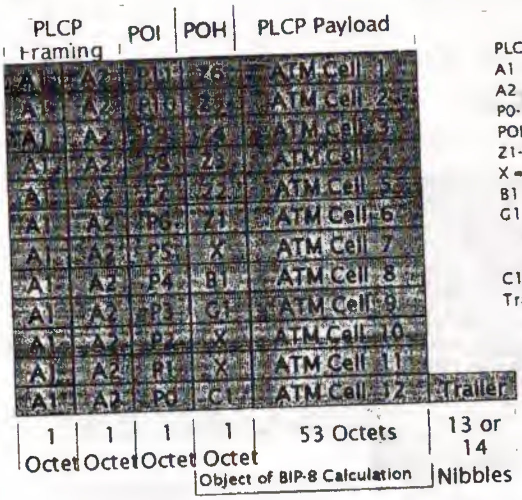

Definitions

PLCP = Physical Layer Convergence Protocol
A1 = 11110110
A2 = 00101000
P0-P11 = Path Overhead Identifier (POI)
POH = Path Overhead
Z1-Z6 = Growth Octets = 00000000
X = Unassigned
B1 = PCLP Bit Interleaved Parity -8 (BIP-8)
G1 = PLCP Path Status
  = AAAAXXXX = FEBE B1 Count
  = XXXXAXXX = RAI
C1 = Cycle Stuff Counter
Trailer Nibbles = 1100

Figure 16 illustrates the DS3 mapping using the Physical Layer Convergence Protocol (PLCP) defined in IEEE 802.6. The ATM cells are enclosed in a 125 μs frame defined by the PLCP, which is defined inside the standard DS3 M-frame. The PLCP mapping transfers 8 KHz timing across the DS3 interface which is somewhat inefficient in that the cell transfer rate is only 40.704 Mbps, which utilises only about 90% of the DS3's approximately 44.21-Mbps payload rate. Note that the BIP-8 indicator is computed over the POH and associated ATM cells of the previous PLCP frame.

TC Header Error Check (HEC) Functions

The Header Error Check (HEC) is a 1-byte code applied to the 5 byte ATM cell header. The HEC code is capable of correcting any single- bit error in the header. It is also capable of detecting many patterns of multiple-bit errors. The TC sublayer generates HEC on transmits and uses it to determine if the received header has any errors. If errors are detected in the header, then the received cell is discarded. Since the header tells the ATM layer what to do with the cell, it is very important that it not have errors; otherwise it might be delivered to the wrong user, or an undesired function in the ATM layer may be inadvertently invoked.

The TC also uses HEC to locate cells when they are directly mapped into a TDM payload. The HEC will not match random data in the cell payloads when the 5 bytes that are being checked are not part of the header. Thus, it can be used to find cells in a received bit stream. Once several cell headers have been located through the use of HEC, then TC knows to expect the next cell 53 bytes later. This process is called HEC-based cell delineation in standards.

TC Cell Rate Decoupling

Figure 17. Cell Rate Decoupling Using Unassigned Cells

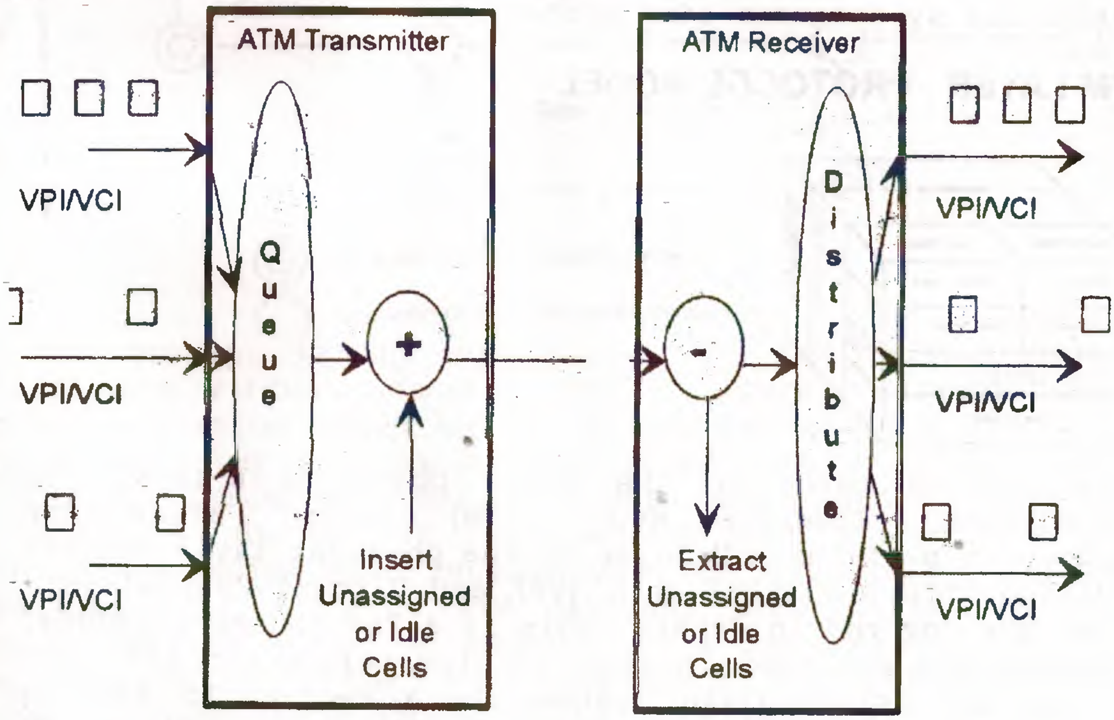

The TC sublayer performs a cell rate decoupling, or speed matching function, as well. Physical media that have synchronous cell time slots (eg, DS3, SONET, SDH, STP, and the Fiber Channel-based method) require this function, while asynchronous media such as the FDDI PMD do not. As we shall see in the next section, there are special coding of the ATM cell header that indicate that a cell is either unassigned or idle. All other cells are assigned which correspond to the cells generated by the ATM layer. Figure 16 illustrates this operation between a transmitting device and a receiving ATM device. The transmitter multiplexes multiple VPI/VCI cell streams, queuing them if an ATM slot is not immediately available. If the queue is empty when the time arrives to fill the next synchronous cell time slot, then the TC sublayer inserts an unassigned or idle cell. The receiver extracts unassigned or idle cells and distributes the other, assigned cells to the destinations.

ITU-T Recommendation I.321 places this function in the TC sublayer of the PHY layer and uses idle cells, while the ATM Forum places it in the ATM layer and uses unassigned cells. This presents a potential low-level incompatibility if different systems use different cell types for cell rate decoupling. Look for ATM systems that support both methods to ensure maximum interoperability. The ITU-T models views the ATM layer as independent of whether or not the physical medium has synchronous time slots.

ATM LAYER - PROTOCOL MODEL

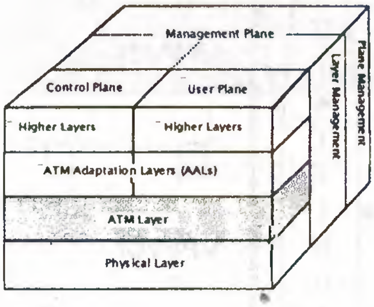

This section moves to the focal point of B-ISDN, the Asynchronous Transfer Mode (ATM) Layer. First the relationship of the ATM layer to the physical layer and its division into a Virtual Path (VP) and Virtual Channel (VC level are covered in detail. This is a key concept. Several examples are provided in this section portraying the role of end and intermediate systems in a real-world setting rather than just a formal model. This is accomplished by explaining how the ATM layer VP and VC functions are used in intermediate and end systems in terms of the layered protocol model. An example is then provided showing how intermediate systems perform ATM VP or VC switching or cross-connection, and how end systems pass cells to the ATM Adaptation Layer (AAL).

Physical Links and ATM Virtual Paths and Channels

Figure 18. Physical Layer, Virtual Paths, and Virtual Channels

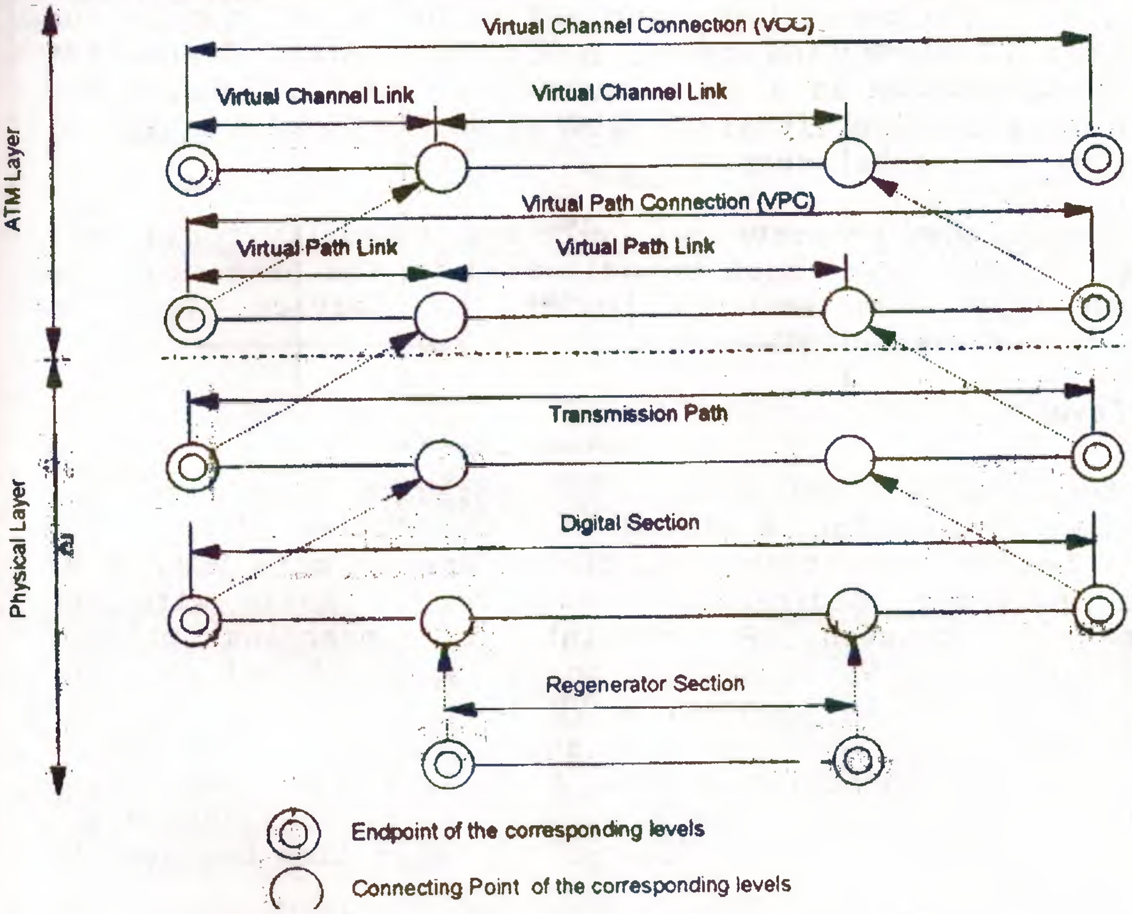

A key concept is the construction of ATM Virtual Paths (VPs) and Virtual Channels (VCs). Figure 18 illustrates this derivation based on ITU-T Recommendation I.311. The physical layer is composed of three levels: regenerator section, digital section, and transmission path as shown in the figure. At the ATM layer we are only concerned about the transmission path because this is essentially the TDM payload that connects ATM devices. Generically, an ATM device may be either an endpoint or a connecting point for a VP or VC. A Virtual Path Connection (VPC) or a Virtual Channel Connection (VCC) exists only between endpoint as shown in the figure. A VP link or a VC link can exist between an endpoint and a connecting point or between connecting points. A VPC or VCC is an ordered list of VP or VC links, respectively.

VC level

The Virtual Channel Identifier (VCI) in the cell header identifies a single VC on a particular Virtual Path (VP). Switching at a VC connecting point is done based upon the combination of virtual path and VCI. A VC Link is defined as a unidirectional flow of ATM cells with the same VCI between a VC connecting point and either a VC endpoint or another VC connecting point. A Virtual Channel Connection (VCC) is defined as a concatenated list of VC links. A VCC defines a unidirectional flow of ATM cells from one user to one or more other users.

A network must preserve cell sequence integrity for a VCC; that is, the cells must be delivered in the same order in which they were sent. A Quality of Service (QoS) is associated with a VPC.

VP Level

Virtual Paths (VPs) define an aggregate bundle of VCs between VP endpoint. A Virtual Path Identifier (VPI) in the cell header identifies a bundle of one or more VCs. A VP link provides unidirectional transfer of cells with the same VPI between VP endpoint or connecting points. Switching at a VP connecting point is done based upon the VPI - the VCI is ignored. A VP link is defined as a VP between a VP connecting point and either a VP endpoint or another VP connecting point. A Virtual Path Connection (VPC) is defined as a concatenated list of VP links. A VPC defines a unidirectional flow of ATM cells from one user to one or more other users.

Standards do not require a network to preserve cell sequence integrity for a VPC; however, the cell sequence integrity requirement of a VCC still applies. A Quality of Service (QoS) is associated with a VPC. If a VPC contains VCCs in different QoS classes, then the VPC assumes the QoS of the VCC with the highest quality.

Intermediate Systems (IS) and End Systems (ES)

As a more concrete illustration of connecting and endpoint functions, Figure 19 looks inside an Intermediate System (IS) and End System (ES) with reference to the ANSI T1.627-defined protocol model. The OSI terms Intermediate System (IS) and End System (ES) and this interpretation were

Figure 19. Intermediate System (IS)/End System (ES) VP/VC Functions

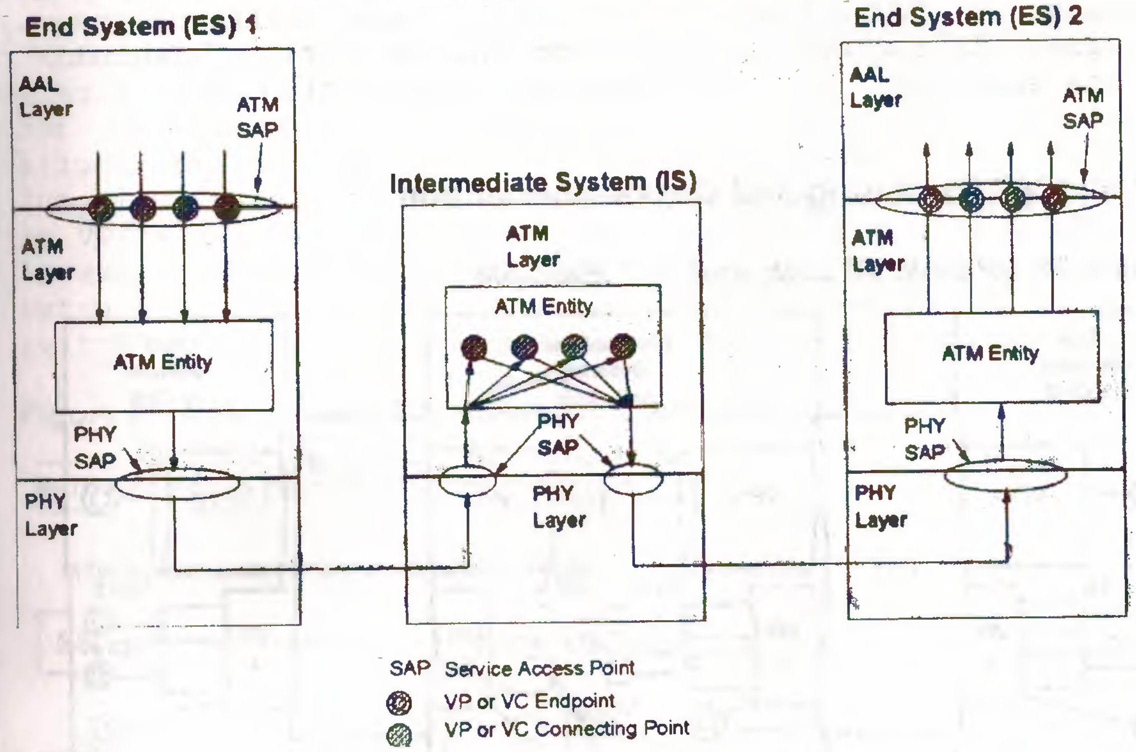

adapted from ANSI T1.627. This example depicts an intermediate network node (IS) connecting two ATM CPE devices (ESs). The ATM Adaptation Layer in End System 1 (ES1) generates cell payloads for a number of connection endpoint at the boundary between the AAL and the ATM layers, which is called the ATM Service Access Point (SAP). The ATM entity in ES1 multiplexes these virtual connections and passes ATM cells to the PHYsical (PHY) layer at the interface labelled PHY-SAP. The bit stream is passed between the PHY layer in ES1 and the Intermediate System (IS) over the physical interface. The IS demultiplexes the ATM connections and applies each to a connecting point in the IS ATM entity. This action is not specified in the protocol model, which is the reason we show it by dashed lines. The connecting point in the ATM entity translates the VPI and/or the VCI depending on whether it is a VP or VC connecting point, determines the outgoing physical interface, and performs other ATM layer functions that we will define later. The IS ATM entity then multiplexes these onto the outgoing PHY-SAP for transfer to the destination - End System 2. The PHY layer in ES2 delivers these to its ATM entity via the PHY-SAP. The ES2 ATM entity demultiplexes the cells and delivers these to the endpoint of the corresponding VP or VC via the ATM-SAP.

The ATM layer requires that cell sequence integrity be preserved. This means that cells are delivered to intermediate connecting points and the destination endpoint in the same order in which they were transmitted.

VP and VC Switching and Cross-Connection

Figure 20. VP Link, VC Link, and VCC Example

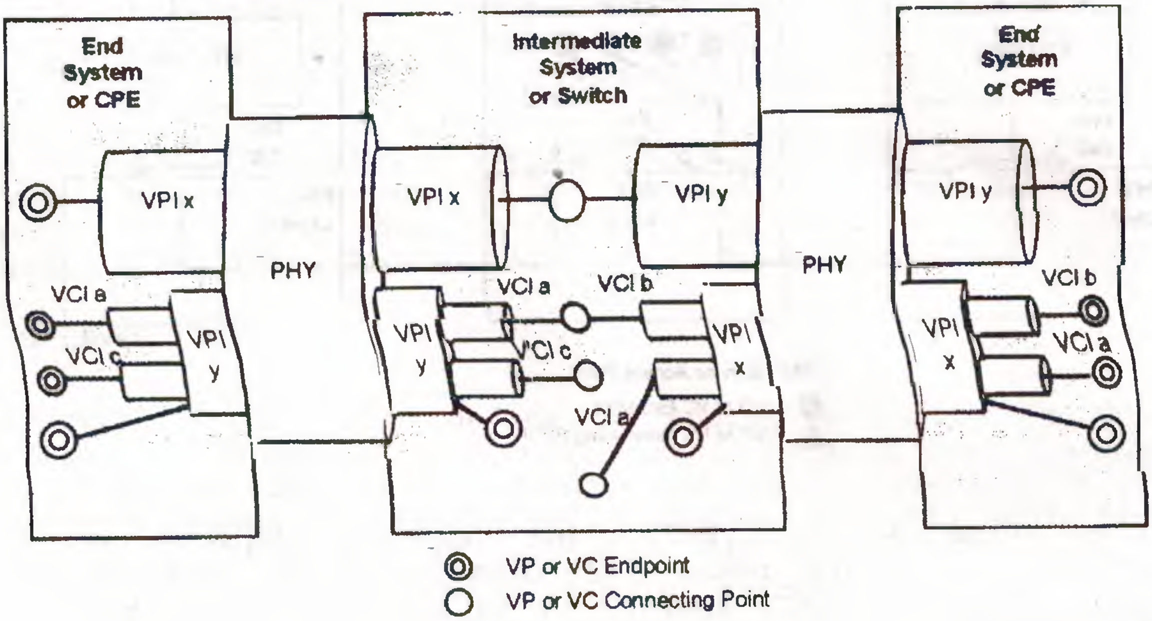

This section provides a specific example of VP and VC endpoint and connecting points in intermediate and end systems for VP links, VPCs, VC links, and VCCs. Figure 20 depicts two end systems (or CPE) and an intermediate system (or switch). The endpoint and connecting points are shown using the terminology from Figure 18. The transmission path, virtual path, and virtual channel are shown as a nested set of pipes using the convention from ITU-T Recommendation I.311. The transmission path PHY layer carries Virtual Paths (VPs) and Virtual Channels (VCs). These may be uni-directional, or bi-directional. Our example shows end systems (or CPE) that have both VP and VC endpoint. The left- hand-side end system, or CPE, originates a VP with VPI x and two VCs with VCI a and c.

The intermediate system contains VP and VC switching functions as shown in Figure 20. The intermediate system VP switching function translates the VPI from x to y since VPI x is already in use on the physical interface to the destination end system. All of the VCIs within VPI x are automatically connected to VPI y. This simultaneous switching of a large number of VCs within a VP is the principal reason for the standardisation of VPs. If only a single level of addressing were used this function would not be possible. VC switching operates within VPs as illustrated by the other VC connection The VC switching function translates the received VCI a to an outgoing VCI b on VPI x for delivery to the destination. VCI c from VPI y is switched to some other destination. Similarly VCI a is extracted from another physical interface and/or VPI on the switch and placed in VPI x for delivery to the end system.

Figure 21. ATM UNI and NNI Reference Configuration

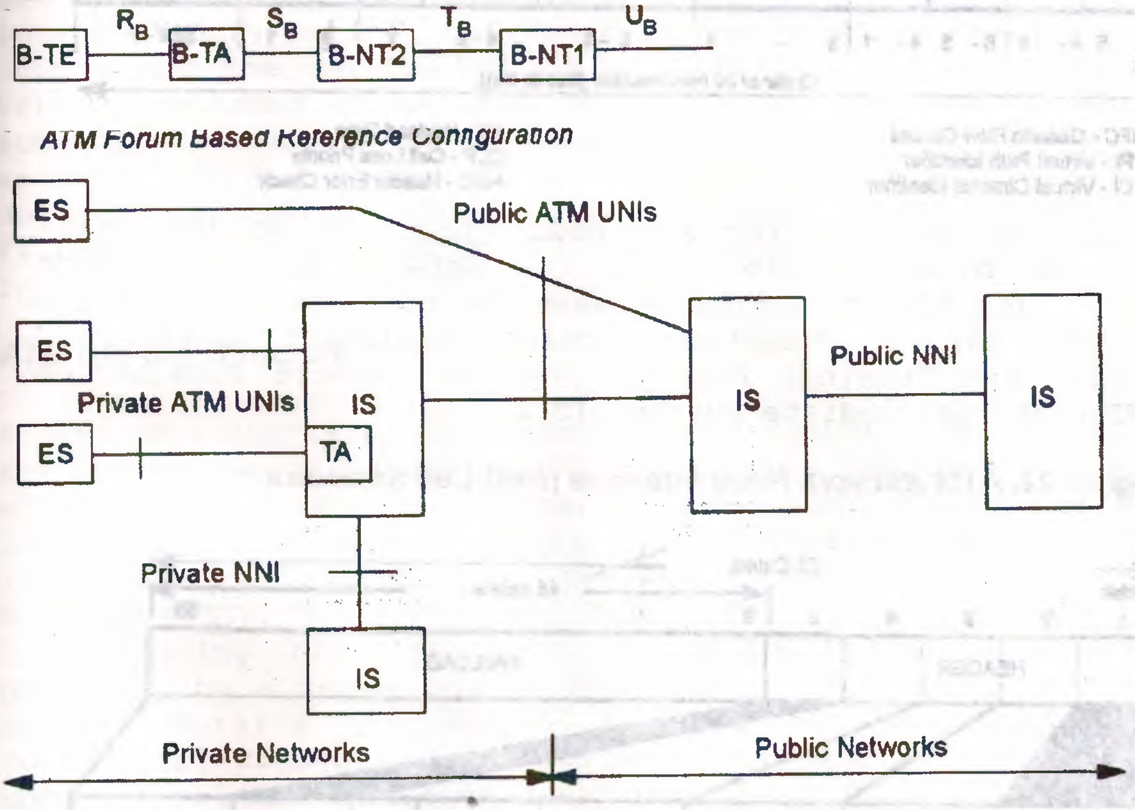

ATM LAYER AND CELL - DEFINITION

Now for a detailed look inside the ATM cell header and the meaning of each field. The User-Network and Network-Network Interfaces (CTNI and NNI) are defined first, followed by a summary of the ITU-T Recommendation I.361, ANSI, and ATM Forum definition of the cell structure at the ATM UNI and NNI. The basic functions of the ATM layer are then introduced, and each function is described in detail.

ATM UNI and NNI Defined

Figure 21. ATM User-Network Interface (UNI) Cell Structure

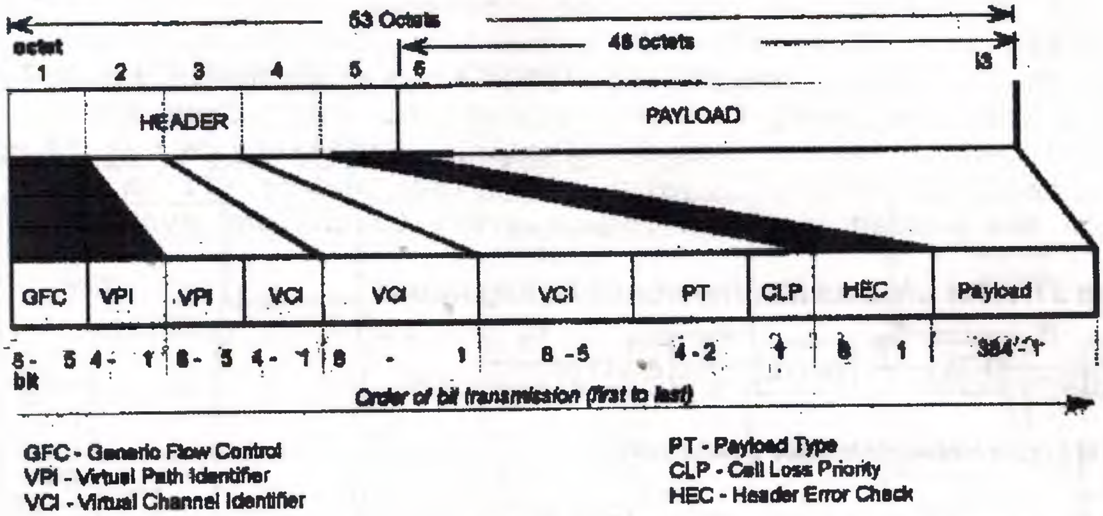

Figure 21 defines the ATM cell structure at the User-Network Interface (UNI) and the Network Node Interface (NNI). The ATM UNI occurs between the user equipment or End System (ES), or Broadband Terminal Equipment (B-TE), and either the Terminal Adaptor (TA) or Network Termination (NT), or Intermediate System (IS).

Figure 22. ATM Network Node Interface (NNI) Cell Structure

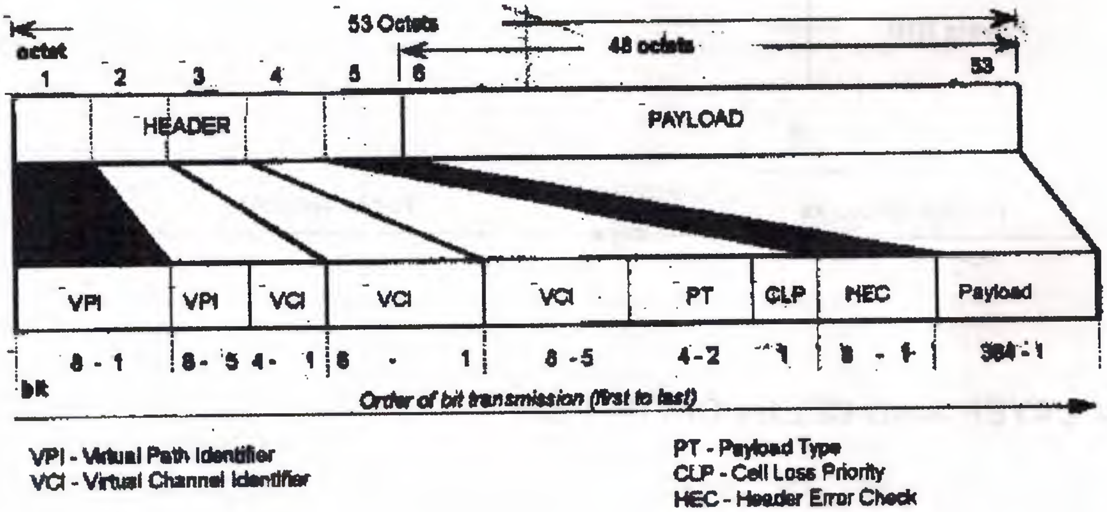

The ATM Forum terminology of private and public UNIs is mapped to the ITU-T reference point terminology in this figure. The ATM UNI may be a private ATM UNI, which would occur at the R or S reference points in ITU-T Recommendation I.413 and ANSI T1.624, or a public ATM UNI, which would occur at reference points T or U as shown in the figure. The Network Node Interface (NNI) defined in ITU-T Recommendation I.113 is normally thought of as the standard interface between networks, which will most likely also be the interface used between nodes within a network. The ATM Forum distinguishes between an NNI used for private networks and public networks as shown in the figure.

ATM Cell Structure at the UNI and NNI

Two standardised coding schemes exist for cell structure: the User-to-Network Interface (UNI) and the Network-Node, or Network-to-Network Interface (NNI). The UNI is the interface between the user [or customer premises equipment (CPE)] and the network switch. The NNI is the interface between switches or between networks. UNI and NNI coding schemes are introduced and each field is defined in this section. ITU-T Recommendation I.361 is the basis of these definitions, with further clarifications given in ANSI T1.627 and the ATM Forum UNI and Broadband Inter-Carrier Interface (B-ICI) specifications.

ATM UNI Cell Structure

Figure 22 illustrates the format of the 53-byte ATM cell at the User-Network Interface (UNI). The cell header contains a logical address in two parts: an 8 bit Virtual Path Identifier (VPI) and a 16-bit Virtual Channel Identifier (VCI). The cell header also contains a 4-bit Generic Flow Control (GFC), 3-bit Payload Type (PT), and a 1-bit Cell Loss Priority (CLP) indicator. The entire header is error-protected by a 1-byte Header Error Check (HEC) field. This section details the meaning of each header field. A fundamental concept of ATM is that switching occurs based upon the VPI/VCI fields of each cell. Switching done on the VPI only is called a Virtual Path Connection (VPC), while switching done on both the VPI/VCI values is called a Virtual Channel Connection (VCC). VPCs/VCCs may be either provisioned as Permanent Virtual Circuits (PVCs), or established via signalling protocols as Switched Virtual Circuits (SVCs). SVCs involve the control plane for the UNI.

ATM NNI Cell Structure

Figure 23 illustrates the format of the 53-byte ATM cell at the Network Node Interface (NNI). The format is identical to the UNI format with two exceptions. First, there is no

Generic Flow Control (GFC) field. Secondly, the NNI uses the 4 bits used for the GFC at the UNI to increase the VPI field to 12 bits at the NNI as compared to 8 bits at the UNI. SVCs involve the control plane for the NNI.

Definition of ATM Cell Header Fields

This section provides a description of each header field.

Generic Flow Control (GFC) is a 4-bit field intended to suppor simple implementations of multiplexing. In the early 1990s, GFC was being specified to implement a DQDB-like, multiple-access-type protocol. However, it appears unlikely tha this type of GFC will be standardised. The curren standards define the uncontrolled mode, where the 4-bit GF field is always coded as zeroes. If too many non-zero GF values are received, layer management should be notified.

Cell Loss Priority (CLP) is a 1-bit field that indicates the loss priority of an individual cell.

Payload Type (PT) is a 3-bit field that discriminates betwee a cell payload carrying user information or one carryin management information. A later section details the codin of the payload type field.

The Header Error Control (HEC) field provides error checking o the header for use by the Transmission Convergence (TC sublayer of the PHYsical layer as defined earlier in this part.

ATM Layer Functions

This section details the key functions of the ATM layer. The ATM Layer provides many functions, including:

- Cell Construction
- Cell Reception and Header Validation
- Cell Relaying, Forwarding, and Copying Using the VPI/VCI
- Cell Multiplexing and Demultiplexing Using the VPI/VCI
- Cell Payload Type Discrimination
- Interpretation of Pre-defined Reserved Header Values
- Cell Loss Priority Processing
- Support for Multiple QoS Classes
- Usage Parameter Control (UPC)
- Explicit Forward Congestion Indication (EFCI)

- Generic Flow Control
- Connection Assignment and Removal

Cell construction, reception and header validation, and several examples of relaying, forwarding, and copying were already covered in earlier examples. Subsequent sections cover descriptions of payload type discrimination, interpretation of predefined header values, and cell loss priority processing.

Relaying and Multiplexing Using the VPI/VCI

As shown through several earlier examples, the heart of ATM is in the use of the VPI and VCI for relaying or switching. ATM also effectively performs multiplexing and demultiplexing of multiple logical connections with different quality requirements using the fixed-length ATM cell.

The number of bits allocated in the ATM cell header limit each physical UNI to support of no more than $2^8 = 256$ virtual paths and each physical NNI to support of no more than $2^{12} = 4096$ virtual paths. Each virtual path can support no more than $2^{16} = 65,536$ virtual channels on the UNI or the NNI.

Although the UNI and NNI cell formats specify 8 and 12 bits for the VPI, respectively, and 16 bits for the VCI on both interfaces, an implementation need only support a smaller number of the lower order bits in the VPI and VCI. Thus, a real ATM application may differ markedly from the above maximums. This means that the number of virtual paths and virtual channels actually supported in a live ATM network may be far less than the maximum numbers defined above. This has important implications in interoperability if one ATM device expects the next ATM device to operate on VPI/VCI bits, but that device ignores these bits. One way to handle this is to allow each system to query the other about the number of bits that are supported. This function is supported in the ATM Forum Interim Local Management Interface (ILMI).

Meaning of Preassigned Reserved Header Values

A key function of the ATM Layer is the identification of preassigned, reserved header values. Figure 23 shows the preassigned (also called pre-defined) header field val for the UNI. The 4-bit GFC field can be used with all these values. The ITU-T has reserved the first 16 VCIs future assignment as preassigned, reserved header va functions. Other portions of the book cover the use these specific header values as indicated below.

Figure 23. Preassigned, Reserved Header Values

Usage	VPI*	VCI	PT	C
Unassigned Cell	00000000	00000000 00000000	XXX	
Idle Cell *	00000000	00000000 00000000	000	
Reserved for Physical layer *	00000000	00000000 00000000	PPP	
Meta-signalling (I.311)	XXXXXXXX	00000000 00000001	0A0	C
General broadcast signaling	XXXXXXXX	00000000 00000010	0AA	C
Point-point signaling	XXXXXXXX	00000000 00000101	0AA	C
Segment OAM F4 Flow Cell	YYYYYYYY	00000000 00000011	0A0	A
End-to-end OAM F4 Flow Cell	YYYYYYYY	00000000 00000100	0A0	A
Segment OAM F5 Flow Cell	YYYYYYYY	ZZZZZZZZ ZZZZZZZZ	100	A
End-to-End OAM F5 Flow Cell	YYYYYYYY	ZZZZZZZZ ZZZZZZZZ	101	A
Resource Management Cell	YYYYYYYY	ZZZZZZZZ ZZZZZZZZ	110	A

* Defined as invalid pattern by ATM Forum

X = "Don't Care"
Y = Any VPI value
Z = Any non-zero VCI

A = Use by appropriate function
C = Originator set CLP
P = Reserved for PHY Layer

The physical layer usage of ATM cells is still in t process of standardisation. The use of the unassigned a idle cell types was described earlier in this part. The N has an additional 4 bits in the VPI field. The N preassigned, reserved header fields have not be completely standardised. The current version of the A Forum B-ICI specification only requires that the F4 O flow, point-to-point signalling, invalid patterns, a unassigned cells be supported.

Meaning of the Payload Type (PT) Field

Figure 24 depicts Payload Type (PT) encoding. We see t first bit is an AAL indication bit (currently used by AA to identify the last cell), the second bit indicat upstream congestion, and the third bit discriminat between data and operations cells. Payload types carryi user information may also indicate whether congestion w experienced by Explicit Forward Congestion Indicati (EFCI) or whether the cell contains an indication to t AAL protocol. The management information payload ty indicates whether the cell is either a segment or end-to-end Operations Administration and Maintenance (OAM) cell for a VCC or a Fast Resource Management (FRM) cell. Use of the AAL indicates PT is covered in the next section of this part.

Figure 24. Payload Type (PT) Encoding

PT Coding	PT Coding
000	User Data Cell, EFCI = 0, AAL_indicate = 0
001	User Data Cell, EFCI = 0, AAL_indicate = 1
010	User Data Cell, EFCI = 1, AAL_indicate = 0
011	User Data Cell, EFCI = 1, AAL_indicate = 1
100	OAM F5 segment associated cell
101	OAM F5 end-to-end associated cell
110	Resource management cell
111	Reserved for future functions

EFCN = Explicit Forward Congestion Notification
AAL_indicate = ATM-layer-user-to-ATM-layer-user indication

Meaning of the Cell Loss Priority (CLP) Field

A value of 0 in the Call Loss Priority (CLP) field means that the cell is of the highest priority - or in other words, it is the least likely to be discarded. A value of 1 in the CLP field means that this cell has low priority - or in other words, it may be selectively discarded during congested intervals in order to maintain a low loss rate for the high-priority CLP=0 cells. The value of CLP may be set by the user or by the network as a result of a policing action.

ATM ADAPTATION LAYER (AAL) - PROTOCOL MODEL

CCITT Recommendations I.362 and I.363 define the next layer of the ATM B-ISDN protocol stack, the ATM Adaptation Layer (AAL). AAL service class attributes and example applications will be covered first followed by the generic AAL protocol model. The Common Part (CP) AALs format and protocol are described in detail with an example of each.

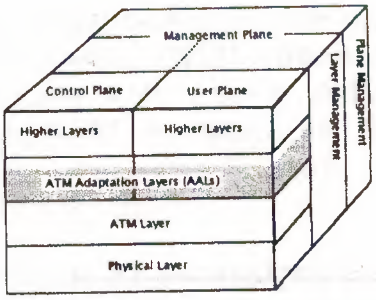

The AAL Protocol Structure Defined

Figure 25. Generic AAL Protocol Sublayer

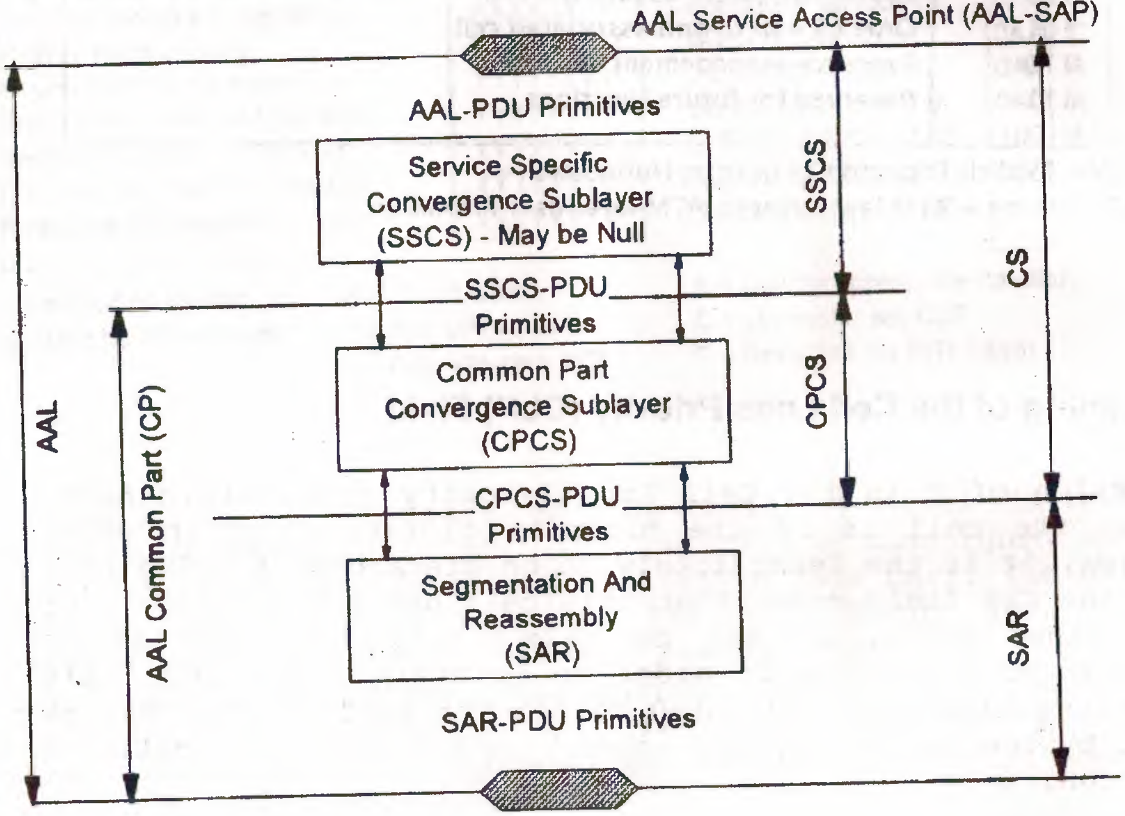

The B-ISDN protocol model adapts the services provided by the ATM Layer to those required by the higher layers through the ATM Adaptation Layer (AAL). Figure 25 depicts the structure and logical interfaces of the AAL. Services provided to higher layers by an AAL Service Access Point (SAP) is shown at the top of the figure, across which primitives regarding the AAL Protocol Data Units (AAL-PDUs) are passed. The AAL is divided into the Convergence Sublayer (CS) and the Segmentation And Reassembly (SAR) sublayer. The CS sublayer is further subdivided into

Service Specific (SS) and Common Part (CP) components. The SSCS sublayer may be null, which means that it need not be implemented. The CPCS must always be implemented along with the Segmentation And Reassembly (SAR) sublayer. These layers pass primitives regarding their respective PDUs between them as labelled in the figure, resulting in the passing of SAR-PDU primitives (which is the ATM cell payload) to and from the ATM layer via the ATM- SAP.

This protocol model may seem somewhat abstract now; however, specific examples clarifying these concepts will soon be provided. This part provides an explanation of the Common Part CS (CPCS) and SAR model, leaving the next part to give specific examples of the SSCS for the control plane that map this model to specific message formats. It is nearly impossible to generalise the CS and SAR functions since, as we shall see, there are significant differences between every AAL.

The protocol primitives will not be covered in detail in this project. Instead, their actions resulting in the transfer of PDUs either between sublayers or across a SAP will be viewed. Standards use the primitives: request, indicate response, and conform.

AAL Service Attributes Classified

ITU-T Recommendation I.362 defines the basic principles and classification of AAL functions. The attributes of the service class are the timing relationship required between the source and destination, whether the bit rate is constant or variable, and the connection mode is connection-oriented or connectionless. Figure 26 depicts the four currently defined AAL service classes, labelled A through D; summarised as follows:

- Class A - constant bit-rate (CBR) service with end-to-end timing, connection-oriented
- Class B - variable bit-rate (VBR) service with end-to-end timing, connection-oriented
- Class C - variable bit-rate (VBR) service with no timing required, connection-oriented
- Class D - variable bit-rate (VBR) service with no timing required, connectionless

The mapping of service classes to AALs is only partially complete in the standards. The next section indicates the AAL(s) that can support the attributes of the defined AAL service class and also gives several application examples for each service class and AAL.

Figure 26. ATM ITU ATM/B-ISDN Service Classes

	Service Class			
Attribute	Class A	Class B	Class C	Class D
Timing relation between source and destination	Required		Not Required	
Bit Rate	Constant	Variable		
Connection Mode	Connection-Oriented			Connection-less
AAL(s)	AAL1	AAL2	AAL3/4 or AAL5	AAL3/4 or AAL5
Example(s)	DS1, E1, nx64 kbps emulation	Packet Video, Audio	Frame Relay, X.25	IP, SMDS

ATM ADAPTATION LAYER (AAL) - DEFINITION

AAL1 through AAL4 were initially defined by the CCITT to directly map to the AAL service classes A through D. ITU-T Recommendation I.363 states the standards for the AALs. AAL1 has been defined by the ITU-T and further clarified in the ANSI T1.630 standard for Continuous Bit Rate (CBR) applications. The history of AAL development for Variable Bit Rate (VBR) services is interesting. Initially, AAL3 was being developed for connection-oriented services and AAL4 for connectionless services. As the details were being defined, it was realised that AAL3 and AAL4 were common enough in structure and function that they were combined into a single class called AAL3/4. A newcomer, AAL5, was conceived by the computer industry in response to perceived complexity and implementation difficulties in the AAL3/4 AAL, which had become aligned with the IEEE 802.6 Layer 2

PDU (L2 PDU). Initially, AAL5 was named the Simple Efficient Adaptation Layer (SEAL) for this reason. AAL5 was adopted by the ATM Forum, ANSI, and the CCITT in a relatively short time frame compared to the usual standards process and has become the predominant AAL of choice in a great deal of data communications equipment. AAL5 is currently standardised for the transport of signalling messages and frame relay. AAL3/4 will likely be chosen for the support of SMDS since it is essentially identical to the IEEE 802.6 L2_PDU.

Figure 27. AAL1 Common Part Convergence Sublayer (CPCS)

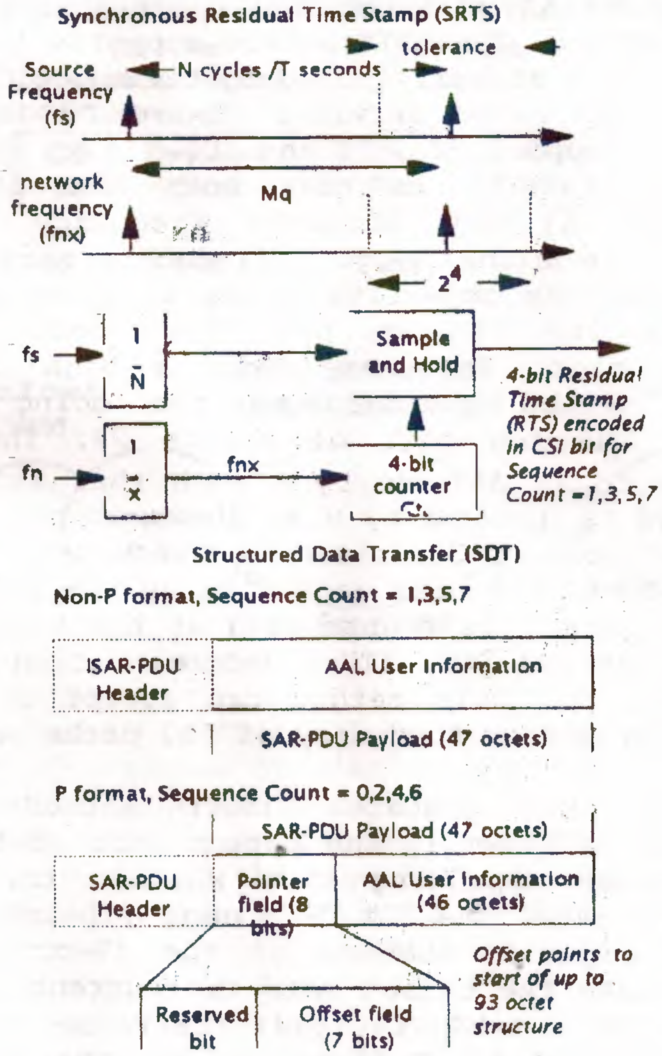

We describe the Common Part Convergence Sublayer (CPCS) and Segmentation And Reassembly (SAR) sublayer for each of the currently standardised Common Part (CP) AALs:

- AAL1 - Constant Bit-Rate (CBR) traffic
- AAL3/4 - Variable Bit-Rate (VBR) traffic
- AAL5 - Lightweight Variable Bit-Rate (VBR) traffic

AAL1

AAL1 specifies how TDM-type circuits can be emulated over an ATM network. Circuit emulation is specified in detail for DS1, DS3, and nxDSO support in ANSI T1.330. AALI supports circuit emulation in one of two modes: the Synchronous Residual Time Stamp (SRTS) or Structured Data Transfer (SDT) method. The SRTS method supports transfer of a DS1 or DS3 digital stream, including timing. SDT supports an octet-structured nxDSO service. Figure 27 depicts the CPCS for AALI in support of SRTS and SDT.

A key concept in SRTS is that both the origin and destination have a very accurate frequency clock of frequency *fn*. The signal (eg. DS1) has a service clock frequency *fs* with the objective being to pass sufficient information via the AAL so that the destination can reproduce this clock frequency with a high degree of accuracy. The method standardised for doing this is illustrated in the top part of Figure 27. The network reference clock *fn* is divided by x such that $1 \leq fnx/fs \leq 2$. The source clock is divided by N as shown in the figure to sample the 4-bit counter Ct driven by the network clock fnx once every N=3008=47*8*8 bits generated by the source. This sampled counter output is transmitted as the Residual Time Stamp (RTS) in the SAR PDU. ITU-T Recommendation I.363 and ANSI T1.630 show how this method can accept a frequency tolerance for the source frequency of 200 parts per million (ppm).

The Structured Data Transfer (SDT) method is more straightforward, as shown in the bottom part of Figure 27. SDT has two modes depending upon whether the sequence number is odd or even. The SDT CPCS uses a 1-octet pointer field in even sequence numbers of the 47-octet SAR-PDU payload to indicate the offset into the current payload of the first octet of a nxDSO payload. The value of n may be as large as 92 in the P-format since the pointer is repeated every other cell when supporting AAL1.

Figure 28 depicts the SAR for AAL1. The 1 octet of overhead is broken down into four fields as shown in the figure. Since the AAL1 SAR uses 1 octet, this leaves 47 octets for user data. There are two major fields: the Sequence Number (SN) and the Sequence Number Protection (SNP) field. The 3-

bit sequence count is incremented sequentially by the origin. The receiver checks for missing or out-of sequence SAR-PDUs and generates a signal alarm when this occurs. The Convergence Sublayer Indication (CSI) bit in the SN field is used differently in the SRTS and SDT modes. In SRTS mode, the 4-bit RTS is sent in odd-sequence-numbered PDUs. In SDT mode, the CSI bit is used to indicate if the pointer field is present in even-sequence-numbered SAR-PDUs. The 3-bit CRC field computes a checksum across the SN field. The parity bit represents even parity across the first 7 bits in the 1-octet SAR-PDU overhead.

Figure 28. AAL1 Segmentation and Reassembly (SAR) Sublayer

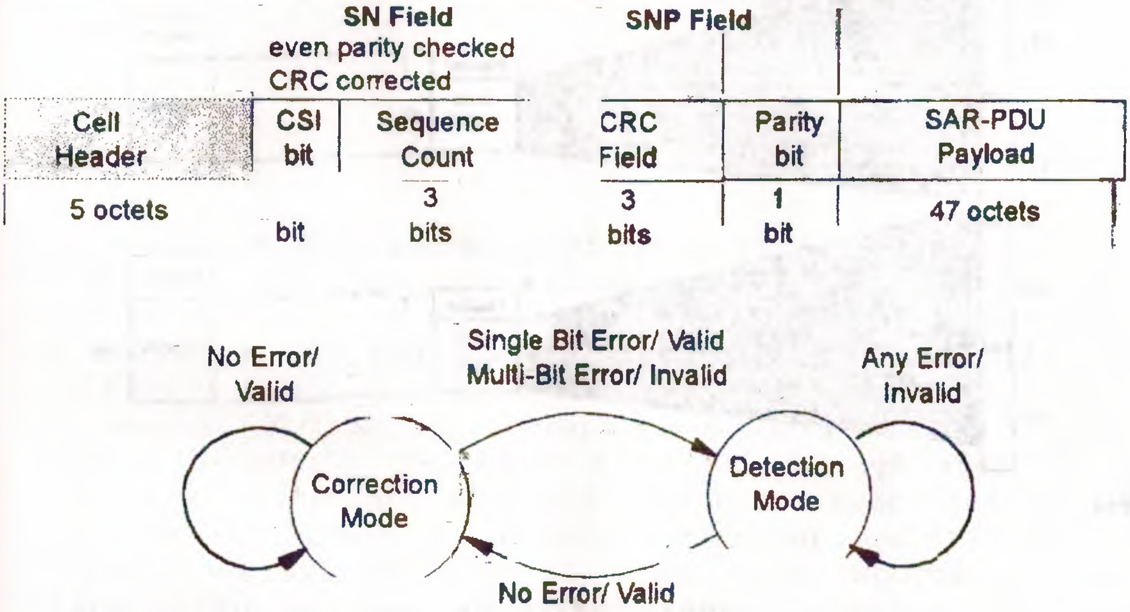

The sequence number is critical to proper operation of AALI since an out-of-sequence or missing SAR-PDU will disrupt at least 47 octets of the emulated circuit bit stream. A well-defined procedure is standardised to correct many problems due to bit errors in the sequence number field, or to accurately detect errors that are not corrected. The operation at the receiver is illustrated in the state machine at the bottom of Figure 28. While in the correction mode the receiver can correct single-bit errors using the CRC, but, if after CRC correction the parity check fails, then either a single, or multiple, bit error has been detected and the receiver switches to detection mode. The receiver stays in detection mode until no error is detected and the sequence number is sequential (ie. valid).

Example of DS1 Circuit Emulation Using AAL1

Figure 29. DS1 Circuit Emulation Using AAL1

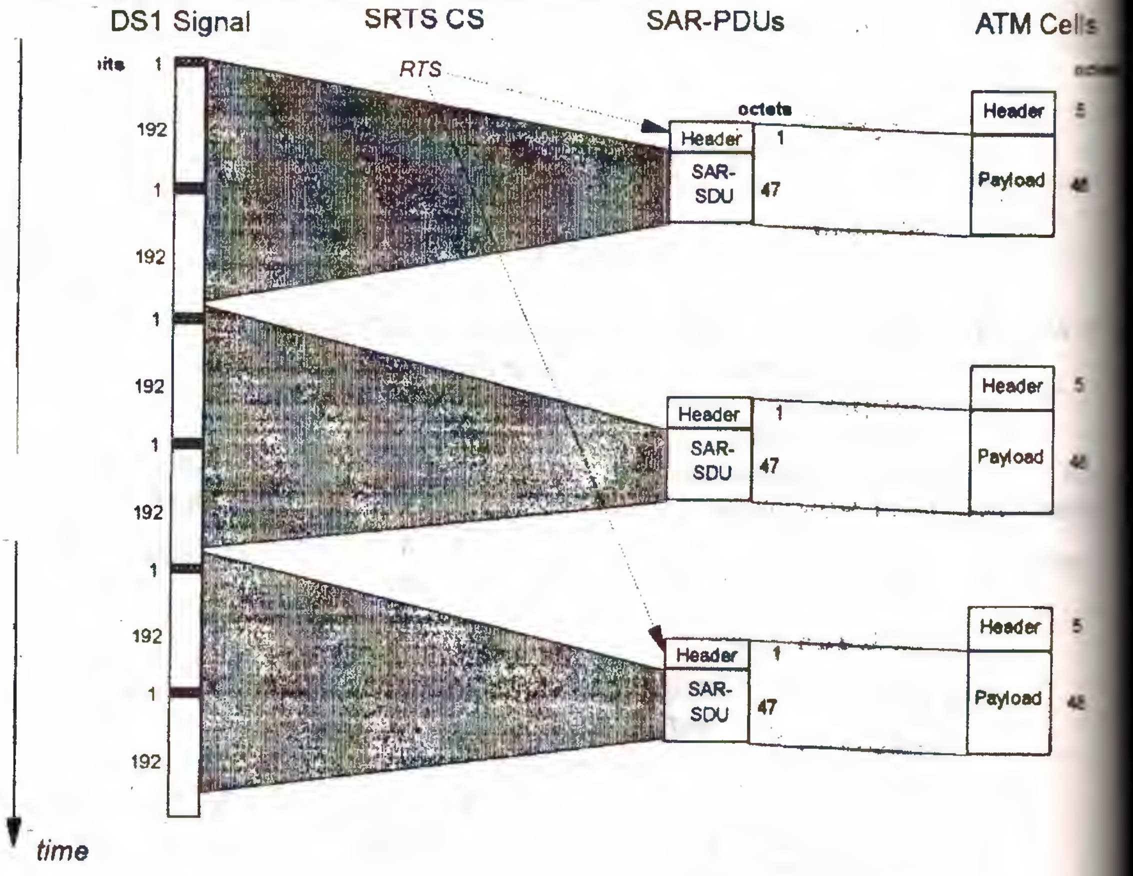

Figure 29 shows an example of a transmitter using AA operating in SRTS mode to emulate a DS1 digital bit stre created by a video codec. Recall from Chapter 4 that a DS frame has a 193-bit frame structure that repeats 8000 ti per second. The DS1 frame has 1 framing bit and 192 us data bits as shown coming out of the codec on the left-ha side of the figure with time running from top to botto The Convergence Sublayer (CS) computes the Residual Ti Stamp (RTS) once every 8 cell times and provides this t the SAR sublayer for insertion in the SAR header. The 19 bit frames are packed into the 47-octet (376 bits) SA Protocol Data Units (PDUs) by the Segmentation A Reassembly (SAR) sublayer. The SAR sublayer then adds th sequence number, inserts the data from the CS and compute the CRC and parity over the SAR header, and passes the 48-octet SAR-PDU to the ATM layer. The ATM layer adds the 5-byte ATM header and outputs the sequence 53-byte cell shown in the figure. The process at the receiver i analogous to that shown here, except the steps ar reversed.

Figure 30. AAL3/4 CPCS Sublayer

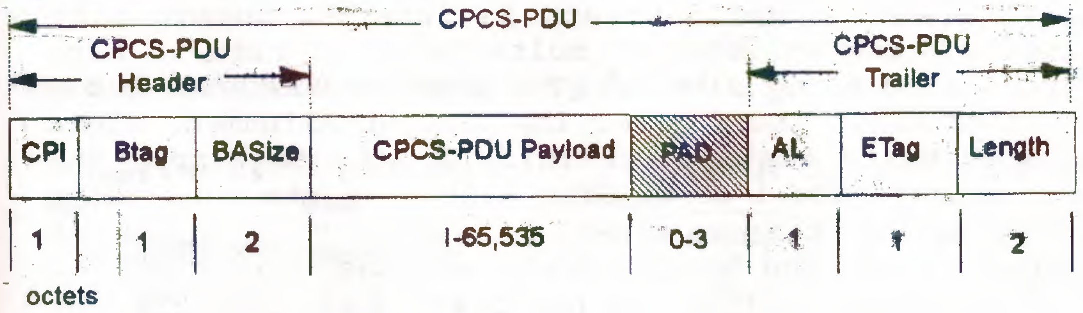

CPI = Common Part Indicator
Btag = Beginning Tag
BASize = Buffer Allocation Size

AL = 32 bit ALignment
ETag = Ending Tag
Length = CPCS-PDU Length

AAL2

AAL2 specifies ATM transport of connection-oriented circuit and VBR high bit-rate packeted audio and video. The current standard at time of publication was not well defined. AAL2 may become a key protocol in future ATM implementations requiring support for variable bit-rate audio and video. The second Motion Photographic Experts (MPEG) video encoding standard, called MPEG2, can be operated at a variable bit rate. The standardisation and specification of interoperable video and audio encoding using ATM is currently an active area of work. Some approaches are investigating the use of either AAL5 or AALI to provide this function.

AAL3/4

AAL3 and AAL4 are combined into a single Common Part (CP) AAL3/4 in support of Variable Bit Rate (VBR) traffic, both connection-oriented or connectionless. Support for connectionless service is provided at the Service Specific Convergence Sublayer (SSCS) level. Figure 30 depicts the CPCS-PDU for AAL3/4. The header has three components as indicated in the figure. The 1-octet Common Part Indicator (CPI) indicates the number of counting units (bits or octets) for the Buffer Allocation Size (BASize) field. The sender inserts the same value for the 2-octet Beginning Tag (BTag) and the Ending Tag (ETag) so that the receiver can match them as an additional error check. The 2-octet BASize indicates to the receiver how much buffer space should be reserved to reassemble the CPCS-PDU. A variable-length PAD field of between 0 and 3 octets is inserted in order to make the CPCS-PDU an integral multiple of 32 bits to make end system processing simpler. The trailer also has three fields as shown in the figure. The 1-octet ALignment field (AL) simply makes the trailer full 32 bits to simply the receiver design. The 1-octet ETag must have the same value as the BTag at the receiver for the CPCS-PDU to be considered valid. The length field encodes the length of the CPCS-PDU field so that the pad portion may be taken out before delivering the payload to the CPCS user.

Figure 31. AAL3/4 Segmentation And Reassembly (SAR) Sublayer

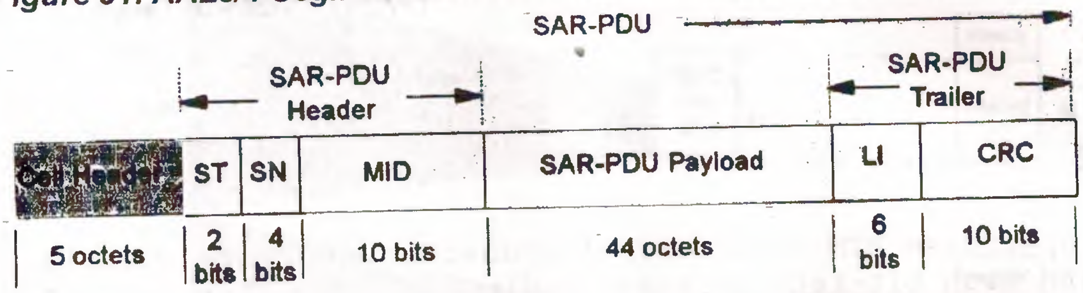

ST = Segment Type (BOM,COM, EOM, SSM)
SN = Sequence Number
MID = Multiplex IDentification

LI = Length Indicator
CRC = Cyclic Redundancy Check

Figure 31 depicts the SAR for AAL3/4. The SAR-PDU encoding and protocol function and format are nearly identical to the L2 PDU from IEEE 802.6. The SAR-PDU has a 2-octet header and trailer. The header contains three fields as shown in the figure. The 2-bit Segment Type (ST) field indicates whether the SAR-PDU is a Beginning Of Message (BOM), a Continuation Of Message (COM), an End Of Message (EOM), or a Single Segment Message (SSM). The 2- bit Sequence Number (SN) is incremented by the sender and checked by the receiver. The numbering and checking begins when a ST of BOM is received. The 10-bit Multiplex IDentification (MID) field allows up to 1024 different CPCS-PDUs to be multiplexed over a single ATM VCC. This is a key function of AAL3/4 since it allows multiple logical connections to be multiplexed over a single VCC. This function is essentially the same one used in the 802.6 L2 protocol where there was effectively no addressing in the cell header. The MID is assigned for a BOM or SSM segment type. The trailer has two fields. The 6-bit Length Indicator (LI) specifies how many of the octets in the SAR-PDU contain CPCS-PDU data. LI has a value of 44 in BOM and

COM segments, and may take on a value less than this in EOM and SSM segments.

Figure 32. Multiplexing Example Using AAL3/4

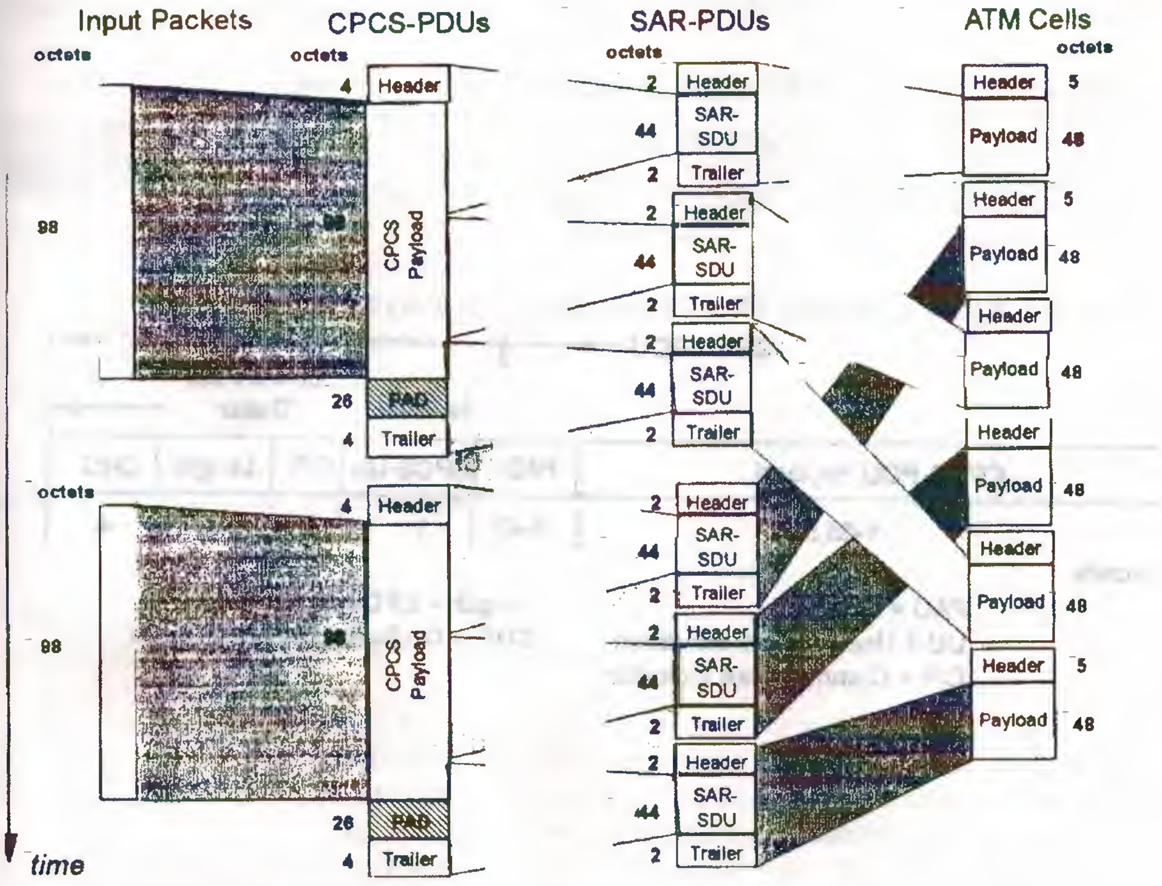

AAL3/4 Multiplexing Example

Figure 32 depicts a data communications terminal that has two inputs with two 98-byte (or octet) packets arriving simultaneously destined for a single ATM output port using the AAL3/4 protocol. On the left-hand side of the figure two 98-byte packets are shown arriving simultaneously. Two parallel instances of the CPCS sublayer encapsulate the packets with a header and trailer. These are then passed to two parallel Segmentation And Reassembly (SAR) processes that segment the CPCS-PDU on two different MIDs, resulting in a BOM, COM, and EOM segment for each input packet. Because all of this occurred in parallel, the ATM cells resulting from this process are interleaved on output. This is the major additional function of AAL3/4 over AAL5, as we shall see by comparison with the AAL5 example late in this part. This also introduces complexity in AAL3/4 that is not present in AAL5.

AAL5

The Common Part (CP) AAL5 supports Variable Bit Rate (VBR traffic, both connection-oriented or connectionless Support for connectionless or connection-oriented service is provided at the Service Specific Convergence Sublayer (SSCS) level.

Figure 33. AAL5 Common Part Convergence Sublayer (CPCS)

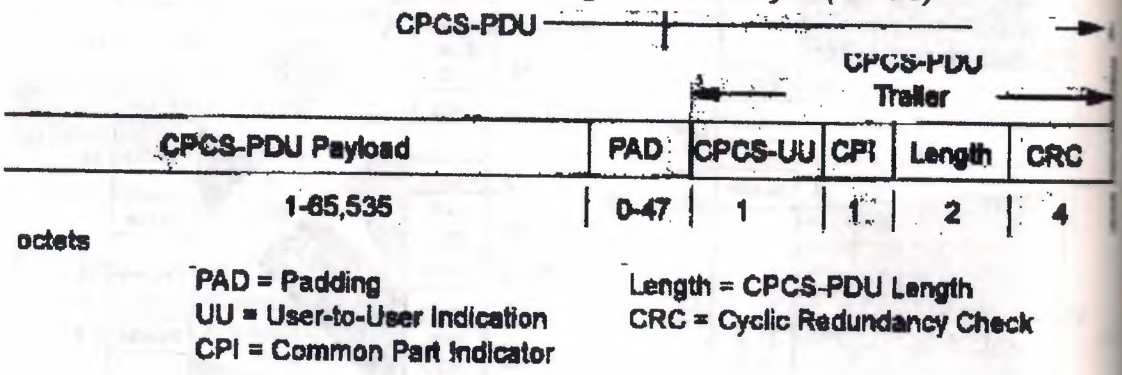

Figure 33 depicts the CPCS for AAL5. The payload may be an integer number of octets in the range of 1 to 2^{16}-1 (65,535). The Padding field is of a variable length choser such that the entire CPCS-PDU is an exact multiple of 48 so that it can be directly segmented into cell payloads. The User-to-User (UU) information is conveyed between AAL users transparently. The only current function of the Common Part Indicator (CPI) is to align the trailer to a 64-bit boundary, with other functions for further study. The length field identifies the length of the CPCS-PDU payloa so that the PAD can be removed. Since 16 bits are allocate to the length field, the maximum payload length is 2^{16}-1 = 65,535 octets. The CRC-32 detects errors in the CPCS-PDU The CRC-32 is the same one used in IEEE 802.3, IEEE 802.5, FDDI, and Fiber Channel.

Figure 34 depicts the SAR for AAL5. The SAR-PDU is simply 48 octets from the CPCS-PDU. The only overhead the SAR sublayer makes use of is the Payload Type code points for AAL indicate. AAL_indicate is zero for all but the last cell in a PDU. A nonzero value of AAL_indicate identifies the last cell of the sequence of cells indicating that reassembly should begin. This was intended to make the reassembly design simpler and make more efficient use of ATM bandwidth, which was the root of the name for the original AAL5 proposal [1], called the Simple Efficient Adaptation Layer (SEAL).

Figure 34. AAL5 Segmentation and Reassembly (SAR) Sublayer

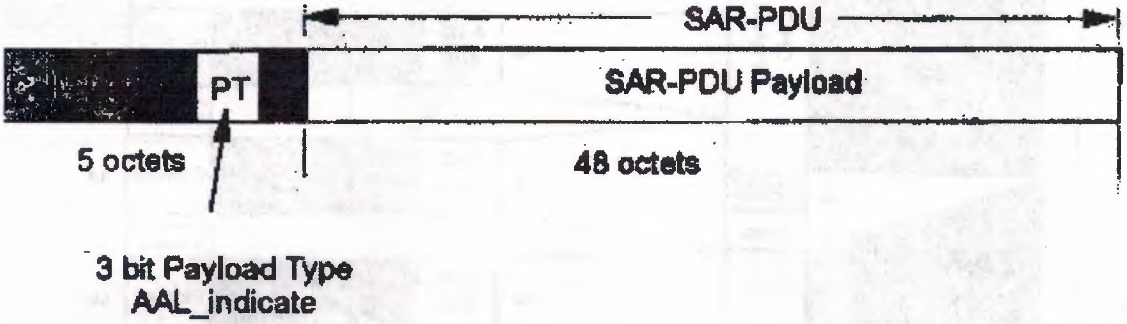

AAL5 Multiplexing Example

Figure 35 depicts the same example previously used for the AAL3/4 example to illustrate a major difference in AAL5. The figure depicts a data communications terminal that has two 98-byte packets arriving simultaneously destined for a single ATM output port, this time using the AAL5 protocol. On the left-hand side of the figure the two 98-byte packets are shown arriving simultaneously. Two parallel instances of the CPCS sublayer add a trailer to each packet. Note that the entire packet does not have to be received before it can begin the SAR function as would be required in AAL3/4 to insert the correct Buffer Allocation Size (BASize) field. The packets are acted on by two parallel Segmentation And Reassembly (SAR) processes which segment the CPCS-PDU into ATM cells. In our example these are destined for the same VPI/VCI, and hence only one can be sent at a time. This implementation is simpler than AAL3/4, but is unable to keep the link as fully occupied as the additional multiplexing of AAL3/4 could if the packets arrive much faster than the rate at which SAR, and ATM cell transmission occur.

Figure 35. Multiplexing Example Using AAL5

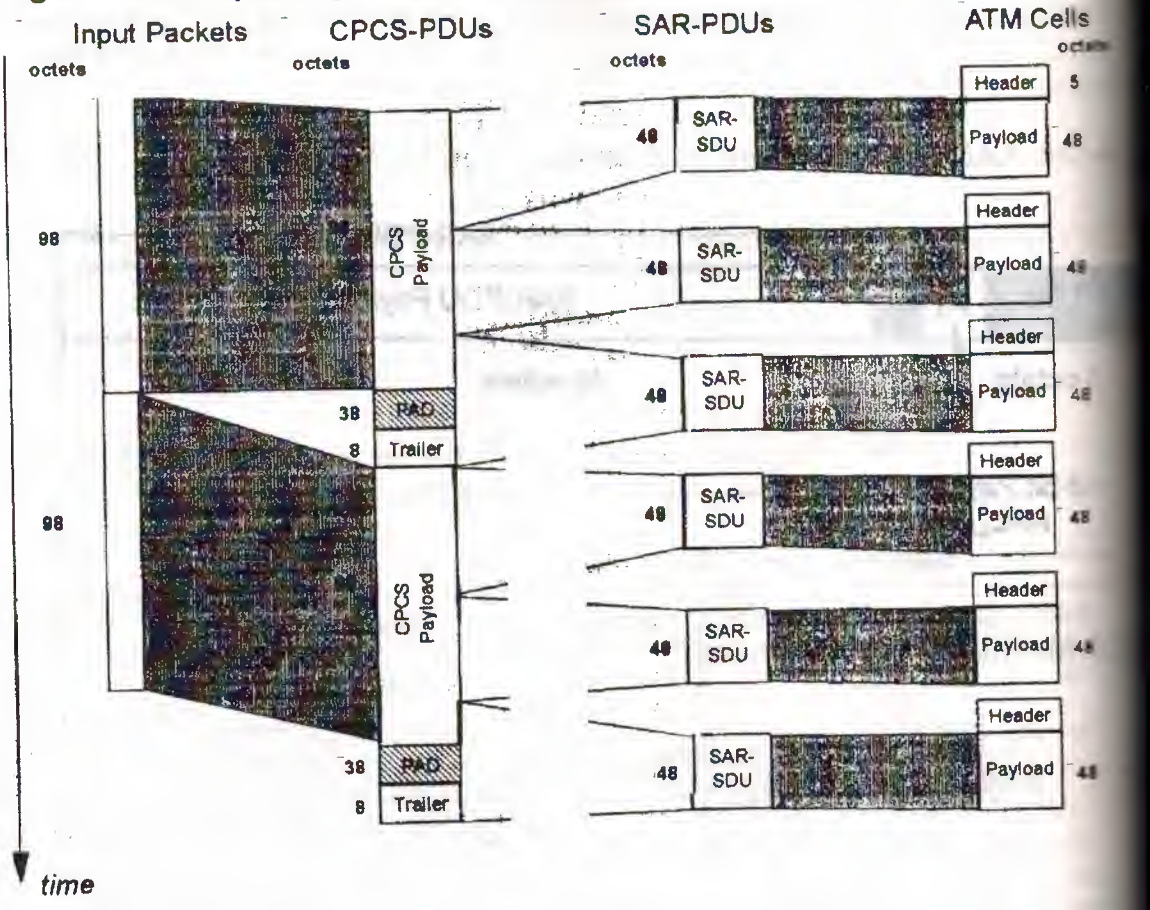

User, Control, and Management Planes

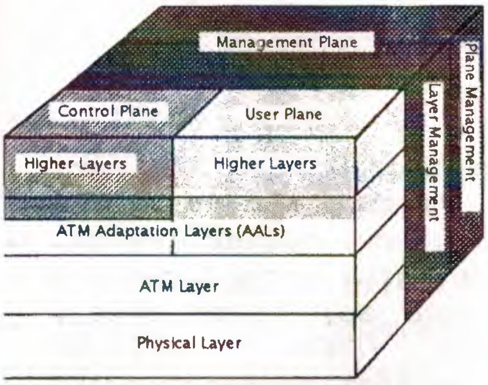

This part provides an overview of the higher layers and the Service Specific (SSCS) portion of the ATM Adaptation Layer (AAL). It includes the user and control planes, along with the management plane, as illustrated in the B-ISDN cube. Of course; a key point is that the principal purpose of the control and management planes is to support the services provided by the user plane. Next, the part on to the very important control plane, which is central in performing the functions needed in a Switched Virtual Connection (SVC) service. The B-ISDN signalling protocol, Service Specific Coordination Function (SSCF), and Service Specific Connection-Oriented Protocol (SSCOP) SSCS protocols are then covered in detail. Finally, the management plane is covered as an introduction. The management plane is composed of overall plane management and management of each of the user and control plane layer components.

USER PLANE OVERVIEW

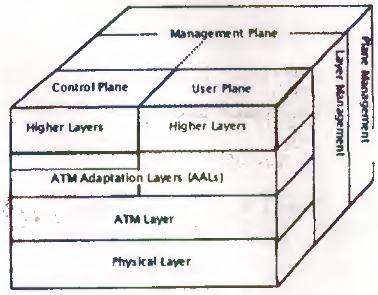

As shown in the shaded portion of the figure first covers the general purpose and function of the user plane from a high level. The state of standardisation in the Service Specific Convergence Sublayer (SSCS) and higher layers of the user plane are summarised as an introduction In fact, a great deal of standardisation work still remains to be done in the area of higher-layer user plane functions.

User Plane - Purpose and Function

The protocol cube figure at the introduction of this part shows that the user plane spans the PHYsical Layer (PHY), ATM Layer, ATM Adaptation Layer (AAL), and higher layers. The AAL and higher layers provide meaningful interfaces and services to end user applications such as; frame relay, Switched Multimegabit Data Service (SMDS), Internet Protocol (IP), other protocols, and Application Programming Interfaces (APIs).

It is also important to note that the control and management planes exist in support of the user plane, in a manner similar to that developed for ISDN. The control plane provides the means to support the following types of connections on behalf of the user plane:

- Switched Virtual Connections (SVCs)
- Permanent Virtual Connections (PVCs)

SVCs and PVCs can be either point-to-point, point-to-multipoint, multipoint-to-point, or multipoint-to-multipoint Virtual Path Connections (VPCs) or Virtual Circuit Connections (VCCs). A VPC or VCC provides a specified Quality of Service (QoS) with a certain bandwidth defined by traffic parameters in an ATM layer traffic contract.

The entire B-ISDN architecture must support the user's application needs to be successful.

User Plane - SSCS Protocols

To date, two Service Specific Convergence Sublayer (SSCS) protocols have been developed specifically for the user plane:

- Frame Relay SSCS
- SMDS SSCS

There is no SSCS required for support of IP or circuit emulation over ATM since the common part AAL directly supports them. There is some discussion and the likely possibility that the Service Specific Connection-Oriented

Protocol (SSCOP) defined for signalling could be used to provide an assured data transfer service in the user plane. It is anticipated that SSCS protocols will be developed for the following user-driven applications:

- Desktop-quality video
- Entertainment-quality video
- Multicast LAN support
- LAN emulation
- Reliable data delivery (like the X.25 capability)
- Interactive, cooperative computing support
- Database Concurrency, Commitment, and Recovery (CCR) function support

A large amount of additional standardisation work is required to support the applications listed above and put a real smile on the user's face. It is likely that many more SSCS protocols will be required to support these applications before the B-ISDN protocol suite is mature.

User Plane - Higher Layers

The area of higher layer protocol support for ATM has only one standard to date, which underscores the need for further standards work:

- IETF RFC 1483 - Multiprotocol Encapsulation Over ATM

That the higher-layer standardisation is not as far along as the lower layers is to be expected. The foundation must first be built before adding the walls and finally the roof. Of course; work continues on the foundational layers by defining new physical layers, new ATM capabilities, and possibly even new AALs in the ever expanding B- ISDN/ATM mansion.

- ATM Forum LAN Emulation working group
- ATM Forum System Aspects and Applications (SAA) working group
- ATM Forum Private Network-Network Interface (P-NNI) working group
- IETF work group supporting IP over ATM
- IETF work group for routing over ATM networks

It is suggested that users follow the progress of these activities and provide inputs to these groups on requirements. Users need these functions, but they are no defined yet; that is why they don't have smiles on thei faces in the above list. You have a voice in the standard process - use it!

CONTROL PLANE AAL OVERVIEW

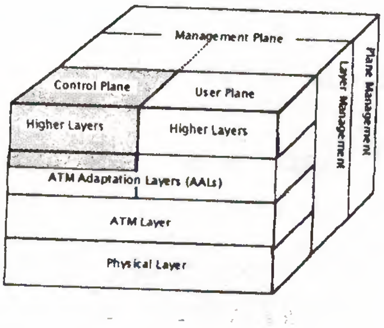

The control plane handles all virtual connection-related functions, most importantly the Switched Virtual Circuit (SVC) capability. The control plane also performs the critical functions of addressing and routing. The higher-layer and service-specific AAL portions of the signalling protocol have recently reached an initial level of standardisation. This chapter covers these functions as indicated by the shaded portions of the B-ISDN cube in the figure.

The signalling protocol architecture of the control plane is very similar to that of Narrowband Integrated Services Digital Network (N-ISDN) as depicted in Figure 36.

The specifications for the Signalling AAL (SAAL) are being developed in the ITU-T and are being adopted by ANSI and the ATM Forum. ITU-T Recommendation Q.2931 (previously called Q.93B) specifies the B-ISDN signalling on the ATM UNI. Q.2931 was derived from both the Q.931 UNI signalling protocol specified for N- ISDN, and the Q.933 UNI signalling protocol for frame relay. The formal name for the ATM UNI signalling protocol is the Digital subscriber Signalling System 2 (DSS2), while the name for ISDN UNI signalling was DSSS 1. ITU-T Recommendation Q.2130 (previously called Q.SAAL.2) specifies the Service Specific Coordination Function (SSCF) for the UNI. ITU-T Recommendation Q.2110 (previously called Q.SAAL.1) specifies SErvice Specific Connection-Oriented Protocol (SSCOP). The ISDN User Part (ISUP) is being adapted in a similar way to broadband as the UNI protocol was in defining the broadband NNI signalling which is called B-

ISUP. The B-ISUP protocol operates over the Message Transfer Protocol 3 (MTP3), identical to that used in Signalling System 7 (SS7) for out-of band N- ISDN and voice signalling. This will allow B-ISDN network signalling the flexibility to operate over existing signalling networks or directly over new ATM networks. The series of ITU-T Recommendations Q.2761 through Q.2764 specify the B-ISUP protocol. ITU-T Recommendation Q.2140 specifies the SSCF at the NNI. The NNI signalling uses the same SSCOP protocol as the UNI.

Figure 36. Overview of Control Plane Architecture

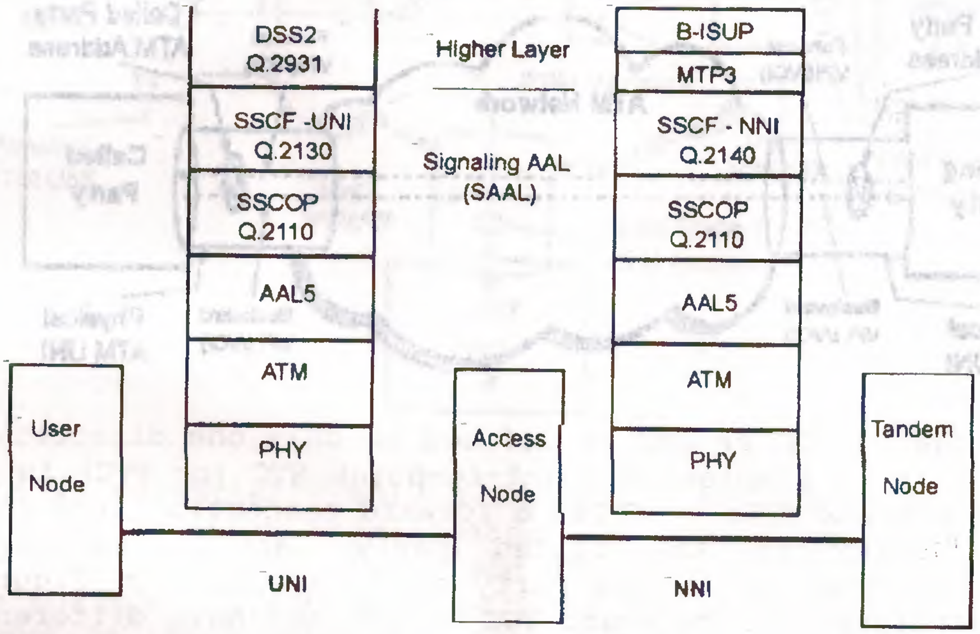

CONTROL PLANE ADDRESSING AND ROUTING

There are two capabilities that are critical to a switched network: addressing and routing. Addressing occurs at the ATM VPI/VCI level and at the logical network level. Since the VPI/VCI is unique only to a physical transmission path, there is a need to have a higher level address that is unique across at least each network. Ideally, the address should be unique across all networks in order to provide universal connectivity. Once each entity involved in switching virtual connections has a unique address, there is another even more onerous problem of finding a route from the calling party to the called party. Using routing solves this problem.

ATM Layer VPI/VCI Level Addressing

The signalling protocol automatically assigns the VPI/VCI values to ATM addresses and physical ATM UNI ports based upon the type of SVC requested according to the following set of rules: either point-to-point or point-to-multipoint. A physical ATM UNI port must have at least one unique ATM address. An ATM UNI port may also have more than one ATM address.

Figure 37. Point-to-Point Switched Virtual Connection

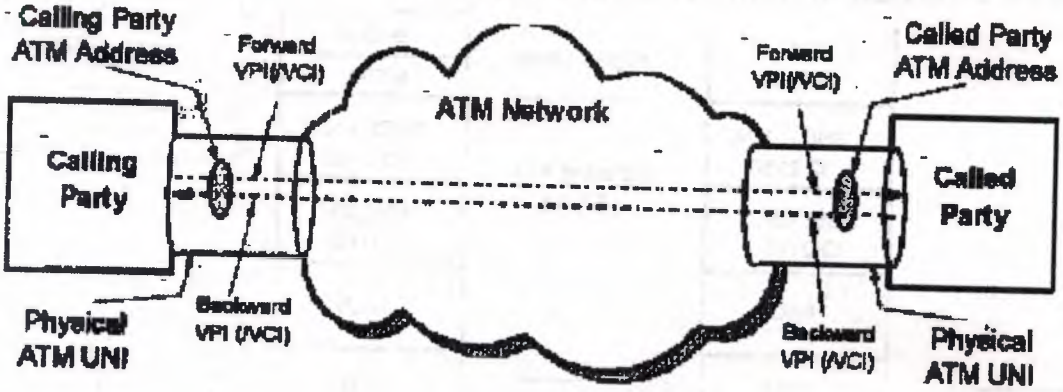

Recall that a VCC or VPC is defined in only one direction; that is, it is simplex. A point-to-point SVC (or PVC) is a pair of simplex VCCs or VPCs: a forward connection from the calling party to the called party, and a backward connection from the called party as illustrated in Figure 37. The forward and backward VCC or VPC can have different traffic parameters. A point-to-point SVC is defined by the forward and backward VPI (and VCI for a VCC) as well as the ATM address associated with the physical ATM UNI ports on each end of the connection. The VPI (/VCI) assignment can be different for the forward and backward directions of a VPC or VCC at the same end of the connection as well as being different from the other end of the connection. A convention where the VPI (and VCI for a VCC) is identical at the same end of a connection may be used.

A point-to-multipoint SVC (or PVC) is defined by the VPI and the ATM address associated with the physical ATM UNI port of the root node, and the ATM address and VPI (/VCI) for each leaf node of the connection, as shown in Figure 38.

Figure 38. Point-to-Multipoint Switched Virtual Connection

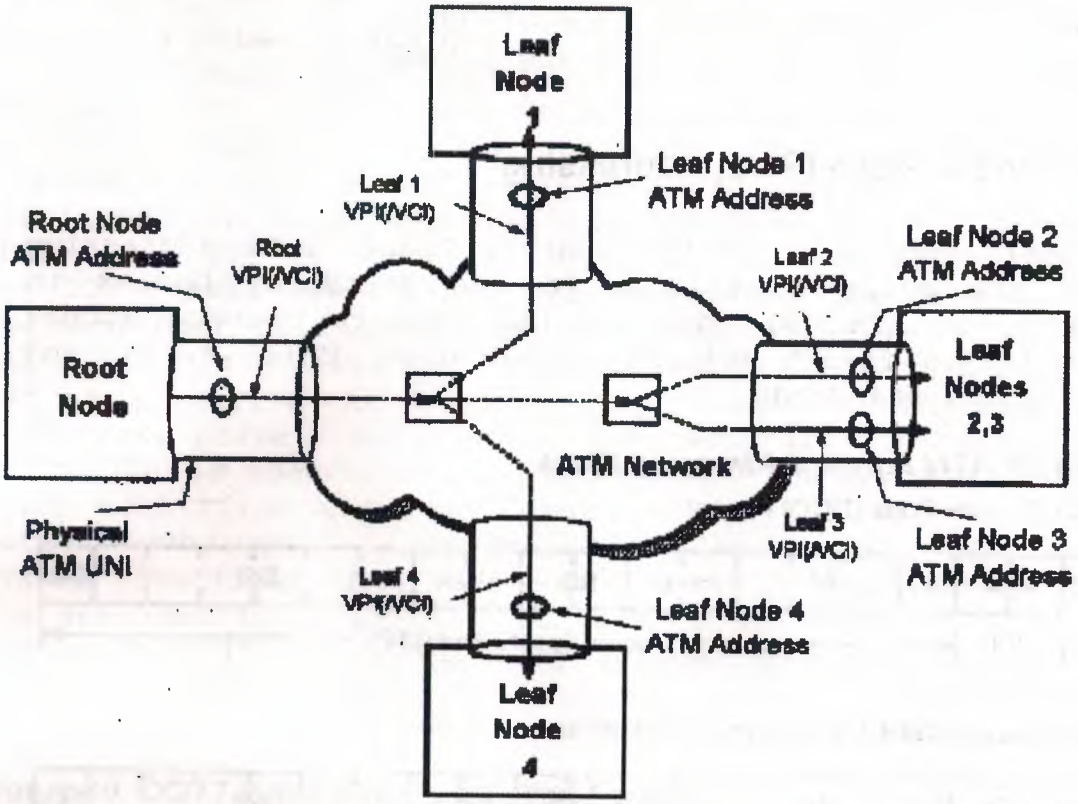

There is essentially only a forward direction because the backward direction is allocated zero bandwidth. Note that more than one VPI/VCI value and ATM address can be assigned to a physical interface as part of the point-to-multipoint connection. This means that the number of physical ATM UNI ports is always less than or equal to the number of logical leaf endpoints of the point-to-multipoint connection. The implementation of a point-to-multipoint connection should efficiently replicate cells within the network. A minimum spanning tree is an efficient method of constructing a point-to-multipoint connection, as illustrated in Figure 38.

Desirable Attributes of ATM Layer Addressing

A set of desirable attributes should be followed when designing an ATM Layer addressing scheme. Of course; each address must be unique. The desirable attributes of the addressing scheme include at least the following:

- Simplicity
- Automatic assignment

- Efficient usage of the address space
- Ease of managing changes in addresses
- Extensibility of the addressing scheme

ATM Control Plane (SVC) Addressing

Currently two types of ATM Control Plane (SVC) addressing plans are being considered in the standards bodies to identify an ATM UNI address: the Network Service Access Point (NSAP) format defined in ISO 8348, CCITT X.213; and CCITT E.164 standards.

Figure 39. ATM Forum Addressing Plans

a) Data Country Code (DCC) Format

b) International Code Designator (ICD) Format

c) E.164 ATM Address Format

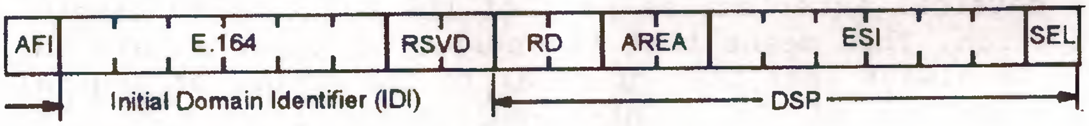

AFI	Authority and Format Identifier	RSVD	Reserved
DCC	Data Country Code	RD	Routing Domain
DFI	DSP Format Identifier	ESI	End System Identifier
AA	Administrative Authority	SEL	SELector

Figure 39 summarises the current version of the NSAP addressing plans from the ATM Forum UNI version 3.0 specification. The Authority and Format Identifier (AFI) identifies which of the formats is being used and designates the authority that allocates the Data Country Code (DCC). Each address is composed of an Initial Domain Identifier (IDI) and a Domain Specific Part (DSP). These may be subject to change as a result of further study due to the strong relationship between addressing and routing.

The ATM Forum currently specifies three IDI formats: DCC, ICD, and E.164. The Data Country Code (DCC) specifies the country in which an address is registered as defined in ISO 3166. The International Code Designator (ICD) identifies an international organisation as administered by the British Standards Institute. E.164 specifies ISDN and telephone numbers that will be defined later in this part.
Within each of these domains lies the Domain Specific Part (DSP). The DSP is identical for the DCC and ICD formats. The DSP Format Identifier (DFI) specifies the meaning of the remainder of the address. The Administrative Authority (AA) field identifies the organisation that is responsible for administering the remainder of the address, for example a carrier, private network, or manufacturer. The remainder of the DSP is identical for all domains. The Routing Domain (R.D) identifies a unique domain within the IDI format - further subdivided using the 2-byte AREA field. The End System Identifier (ESI) and SELector (SEL) portions of the DSP are identical for all IDI formats as specified in ISO 10589. The ESI can be globally unique, for example, a 48-bit IEEE MAC address. Routing but may be used by End Systems (ES) does not use the SELector (SEL) field.

Figure 40. CCITT E.164 Numbering Plan Format

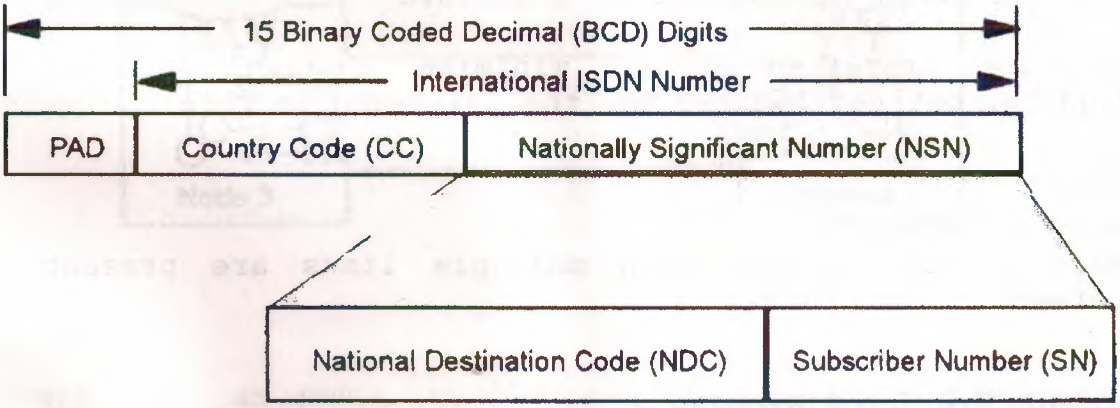

Figure 40 summarises the CCITT E.164 numbering plan format. The international number is coded in Binary Coded Decimal (BCD) and is padded with zeroes on the left-hand side to result in a constant length of 15 digits. There is a Country Code (CC) of one to three digits as standardised in CCITT Recommendation E.163. The remainder of the address is a Nationally Significant Number (NSN). The NSN can be further broken down as a National Destination Code (NDC) and Subscriber Number (SN). The North American Numbering

Plan (NANP) is a subset of E.164. The NDC currently corresponds to an area code and switch NXX identifier for voice applications. Further standardisation of the E.164 number for data and other applications, such as SMDS, is in progress.

Basic Routing Requirements

Cells from the same VPC or VCC must follow the same route, defined as the ordered sequence of physical switch ports that the cells traverse from source to destination. A route is established in response to the following events:

- PVC is newly provisioned
- A SVC connection request is made
- Failed PVC is being automatically reestablished

A route is cleared in response to the following events:

- A PVC disconnect order is processed
- A failure is detected on restorable PVC
- A SVC disconnection request is made
- Call clearing in response to a failure

The route traversed should minimise a cost function including, but not limited to, the following factors:

- Delay
- Economic expense
- Balance utilisation (when multiple links are present between a node-pair)

Desirable Routing Attributes

There are desirable attri6utes to follow when designing an ATM layer routing scheme. Attributes of the routing scheme include at least the following:

- Simplicity
- Automatic determination of least-cost route(s)
- Ease of managing changes in the network in terms of new links and nodes
- Extensibility of the routing scheme to a large network

A Simple ATM Layer VCC Routing Design

A simple routing design for VCCs utilises routing based upon the VPI value only. Each physical node is assigned a VPI value, which means that it is a VPC endpoint as illustrated in Figure 41.

Figure 41. Illustration of simple VPC-Based Routing

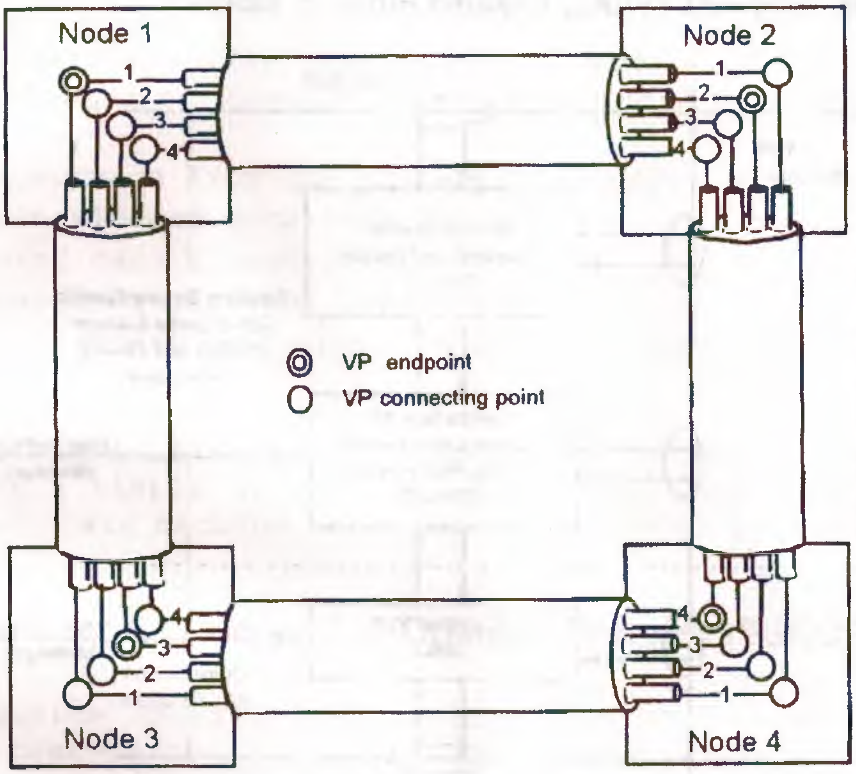

Every node can route traffic to a destination node using a VPC connecting point with the VPI corresponding to the destination node number. Each node - knowing that the tandem nodes will connect this VPC through to the destination node accomplishes this routing.
The principal advantage of this method is that it is very simple - no VPI or VCI translation is required. This method has several disadvantages; it is inefficient since VPIs are allocated on routes that are not used, and it limits the number of VPCs that can be assigned to user applications.

CONTROL PLANE - PROTOCOL MODEL

This section introduces the protocol model for the signalling Service Specific Convergence Sublayer (SSCS), and also summarises the B-ISDN UNI protocol for signalling.

Layered SSCS Model

Figure 42. Signalling AAL (SAAL) Layered Protocol Model

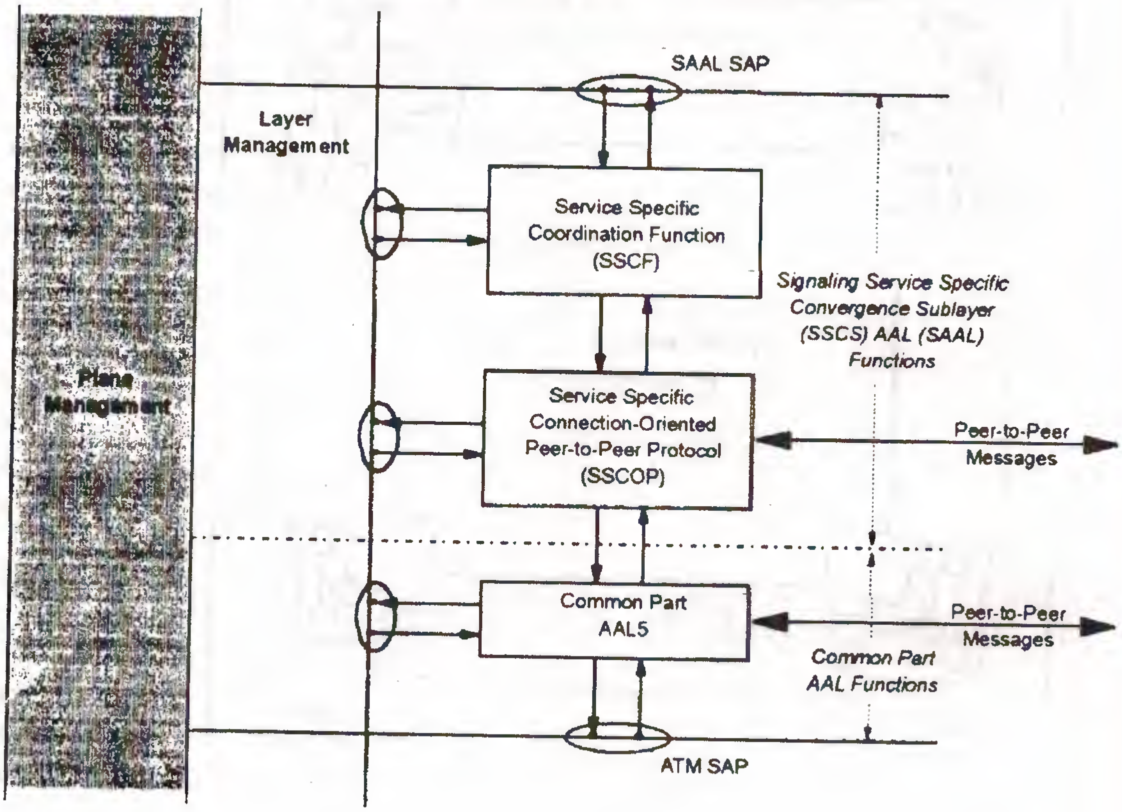

Figure 42 illustrates the protocol model for the Signalling AAL (SAAL). The Common Part AAL (CP-AAL) is AAL5 as defined in previous part. The SSCS portion of the SAAL is composed of the following two protocols:

- Service Specific Coordination Function (SSCF)
- Service Specific Connection-Oriented Protocol (SSCOP)

The SAAL primitives are provided at the SAAL Service Access Point (SAP). The CP AAL5 interfaces with the ATM layer at the ATM SAP. There is a one-to-one correspondence between an SAAL SAP and an ATM SAP. The signalling SSCF and SSCOP protocols and the CP-AAL are all managed as separate layers, by corresponding layer management functions as indicated on the left-hand side of Figure 42. Layer management is responsible for setting parameters in the individual layer protocols and monitoring their state and performance. Plane management coordinates across the layer management functions so that the overall end-to-end signalling capability is provided.

Service Specific Coordination Function (SSCF)

The Service Specific Coordination Function (SSCF) provides the following services to the Signalling AAL (SAAL) user:

- Independence from the underlying layers
- Unacknowledged data transfer mode
- Assured data transfer mode
- Transparent relay of information
- Establishment of connections for assured data transfer mode

The SSCF provides these capabilities primarily by mapping between a simple state machine for the user and the more complex state machine employed by the SSCOP protocol.

Service Specific Connection-Oriented Protocol (SSCOP)

The Service Specific Connection-Oriented Protocol (SSCOP) is a peer-to-peer protocol that performs the following functions:

- Guaranteed sequence integrity, or ordered delivery
- Error correction via error detection and retransmissions Receiver-based flow control of the transmitter
- Error reporting to layer management
- Keep alive messaging when other data is not being transferred
- Local retrieval of unacknowledged or enqueued messages Capability to establish, disconnect, and synchronise an SSCOP connection
- Transfer of user data in either unassured or assured mode
- Protocol level error detection
- Status reporting between peer entities

SSCOP is a fairly complicated protocol, but is specified in the same level of detail as a successful protocol like HDLC. The unassured mode is a simple unacknowledged datagram protocol, similar to the User Datagram Protocol (UDP).

Figure 43. Example of SSCOP Retransmissions Strategy

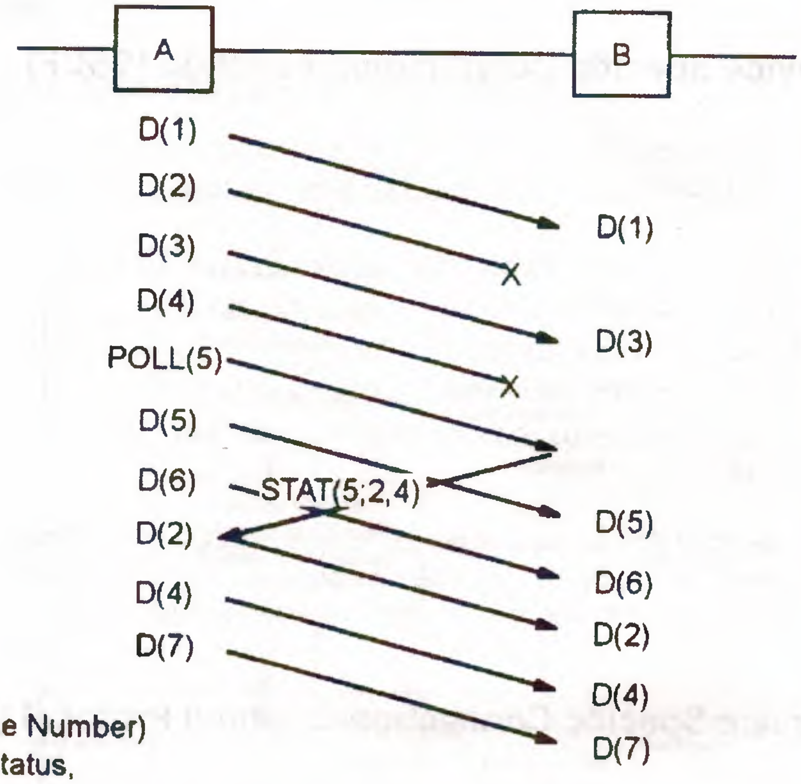

D(SN) = Data(Sequence Number)
POLL(Next) = Poll for status, identifying next SN
STAT(Next;Missing) = Status acknowledging next and list of missing data

Much of the richness of SSCOP is provided in the assured data transfer mode. A connection must be established before any data can be transferred. Figure 43 illustrates an example of the SSCOP retransmission strategy. The error detection capability of AAL5 determines whether a frame is successfully received. SSCOP requires that the transmitter periodically poll the receiver as a keep-alive action and to determine if there is a gap in the sequence of successfully received frames. The receiver must respond to the poll, and if more than a few poll responses are missed, the transmitter will take down the connection. A key feature is where the receiver identifies that one or more frames are missing in its sequence, as illustrated in the figure. The transmitter then only resends the missing frames. This selective reject type of retransmission protocol significantly reduces unproductive retransmissions for very high-speed links, such as those in ATM networks, as compared to a "Go-Back N" retransmission strategy employed in X.25. SSCOP PDUs also employ a 24-bit sequence number that allows very high sustained rates to be achieved in a window flow controlled protocol for example, the Transmission Control Protocol (TCP).

CONTROL PLANE - SIGNALLING FUNCTIONS

The control plane signalling functions currently defined by the ATM Forum and the ITU-T are first described, identifying the common functions and differences. Next, the plans for the next phase of signalling functions identified by the ITU-T are summarised.

Signalling Functions - Current

The current version of the ATM Forum UNI version 3.0 signalling specification and the ITU-T Q.2931 standard are closely aligned. First the functions of Q.2931 are summarised, followed by a description of how the ATM Forum UNI 3.0 both subsets Q.2931, and defines additional functions.
The major functions defined in ITU-T Recommendation Q.2931 are:

- Point-to-point connection setup and release
- VPI/VCI selection and assignment
- Quality of Service (QoS) class request
- Identification of calling party
- Basic error handling
- Communication of specific information in setup request
- Subaddress support
- Specification of Peak Cell Rate (PCR) traffic parameters
- Transit network selection

The ATM Forum UNI version 3.0 specification does not require the following capabilities from Q.2931:

- No alerting message sent to called party

- No VPI/VCI selection or negotiation
- No overlap sending
- No interworking with N-ISDN
- No subaddress support
- Only a single transit network may be selected

The ATM Forum UNI version 3.0 specification defines the following capabilities in addition to Q.2931:

- Support for a call originator setup of a point-to-multipoint call
- Extensions to support symmetric operation
- Addition of sustainable cell rate and maximum burst size traffic parameters
- Additional information elements for point-to-multipoint endpoints
- Additional NSAP address structures

There is a general intention stated by both the ITU-T and the ATM Forum to align the specifications of future releases.

Signalling Functions - Next Phase

The ITU-T has specified a capability set 2 that defines a set of capabilities that will be standardised next. These capabilities must first be described in more detail before protocol standardisation can begin. The following general functions are part of capability set 2:

- Specification of a call model where each call may have multiple connections, for example, in multimedia
- Support for a distributed point-to-multipoint call setup protocol
- Renegotiation of traffic parameters during the course of a connection
- Support for multipoint and multipoint-to-point calls
- Specification of metasignaling which establishes additional connections for signalling

CONTROL PLANE - SIGNALING PROTOCOL

This section provides an overview of the signalling messages and their key parameters. The basics of the signalling protocol are then presented. Finally, an example of a point-to-point call setup and release and the establishment of a point-to-multipoint call illustrate the signalling messages and protocol in action.

The Signalling Messages

The following sections summarise key parameters for the following major signalling messages from the Q.2931 protocol as defined by the ATM Forum UNI specification version 3:0:

Point-to-Point Connection Control:

- Call Establishment Messages
  CALL PROCEEDING
  CONNECT ACKNOWLEDGE SETUP
- Call Clearing Messages
  RELEASE
  RELEASE COMPLETE
- Status Messages
  STATUS ENQUIRY
  STATUS (Response)
- Global Call Reference Related Messages:
  RESTART (All)
  RESTART ACKNOWLEDGE
  STATUS
- Point-to-Multipoint Connection Control:
  ADD PARTY
  ADD PARTY ACKNOWLEDGE
  ADD PARTY REJECT
  DROP PARTY
  DROP PARTY ACKNOWLEDGE

Each signalling message has a number of Information Elements (IE), some of which are Mandatory (M) and others of which are Optional (O). The key mandatory information elements used in the protocol are:

- ATM user cell rate requested
- Called party number
- Connection identifier (assigned VPI/VCI value)
- QoS Class requested

The messages related to a particular call attempt each contain a common mandatory information element, the call reference, that is unique on a signalling interface. All messages must also contain an information element for their type, length, and protocol discriminator (ie. the set from which these messages are taken). There are an even larger number of optional parameters, such as:

- Broadband bearer capability requested
- Broadband lower and higher-layer information
- AAL parameters
- Called party subaddress
- Calling party number and subaddress
- Transit network selection
- Cause Code
- Endpoint reference identifier and endpoint state number

The Signalling Protocol

The signalling protocol specifies the sequence of messages that must be exchanged, the rules for verifying consistency of the parameters, and the actions to be taken in order to establish and release ATM layer connections. A significant portion of the specification is involved with handling error cases, invalid messages, inconsistent parameters, and a number of other unlikely situations. These are all-important functions since the signalling protocol must be highly reliable in order for users to accept it.
Signalling protocols may be specified in several ways: via narrative text, via state machines, or via a semigraphical Specification Definition Language (SDL). The ATM Forum UNI specification version 3.0 uses the narrative method. For complicated protocols, such as Q.2931, a very large sheet of paper would be needed to draw the resulting state machine in a manner such that a magnifying glass is not required to read it. The SDL allows a complicated state machine to be formally documented on multiple sheets of paper in a tractable manner. For example, Q.921 and SSCOP are specified using SDL.

The Q.2931 protocol is based upon the ISDN Q.931 and Frame Relay Q.933 protocols. Since a large amount of expertise has been built up over the years on these subjects, the prospects for implementing the rather complex Q.2931 protocol in an interoperable manner are encouraging.

Point-to-Point Call Setup and Release Examples

Figure 44. Point-to-Point Call setup Example

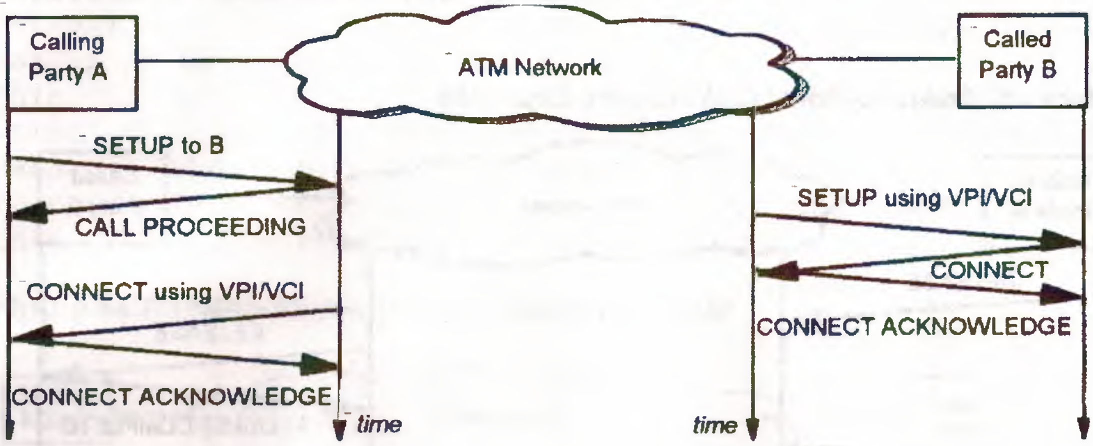

Two simple, but relevant, examples for a point-to-point call, call setup and call release, illustrate the basic aspects of the ATM Forum- and Q.2931-based signalling protocol for the ATM UNI. Figure 44 illustrates the point-to-point call setup example. The examples employ: a calling party with ATM address A on the left, a network shown as a cloud in the middle, and the called party with ATM address B on the right. Time runs from top to bottom in all of the examples. The calling party initiates the call attempt using a SETUP message indicating B as the called party number. The network routes the call to the physical interface on which B is connected and outputs a SETUP message indicating that the VPI/VCI that should be used if the call is accepted. Optionally, the SETUP message may communicate the identity of the calling party A. The called party accepts the call attempt by returning the CONNECT message, which is propagated back to the originator by the network as rapidly as possible in order to keep the call setup time low. The CONNECT ACKNOWLEDGE message is used from the network to the called party and from the calling party to the network as the final stage of the three-way handshake to ensure that the connection is indeed active.

Figure 45 illustrates the point-to-point call release example. The reference configuration and conventions are the same as in the call setup example above. Either party may initiate the release process. This example illustrates the calling party as the one that initiates the disconnect process by sending the RELEASE message. The network then propagates the RELEASE message across the network to the other party B. The other party acknowledges the RELEASE request by returning a RELEASE COMPLETE message, which is then propagated back across the network to the calling-party RELEASE originator. This two-way handshake completes the call release process.

Figure 45. Point-to-Point Call release Example

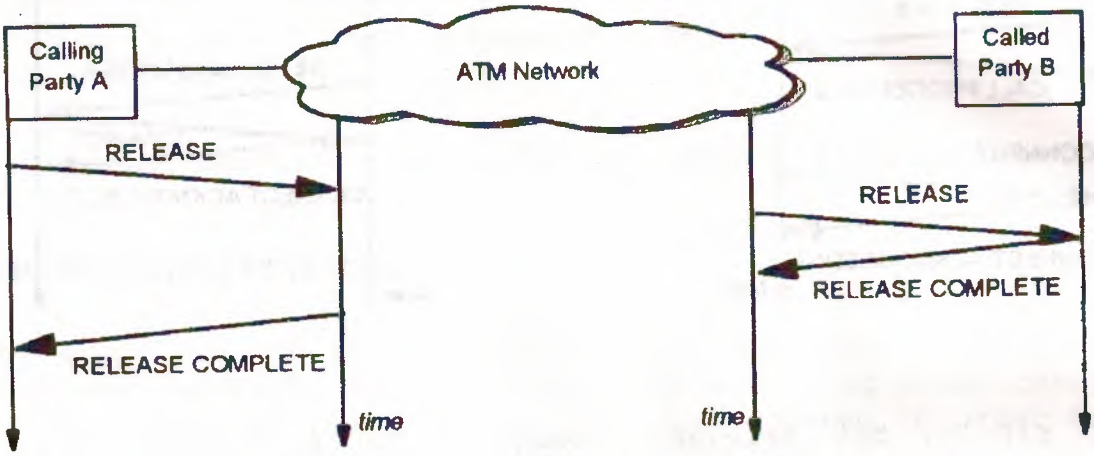

Point-to-Multipoint Call Setup Example

Figure 46 now illustrates an example of setting up a point-to-multipoint call from an originator (root) node A to two leaves, specifically two leaf nodes B and C connected to a local ATM switch on a single ATM UNI, and a third leaf node D connected to a separate ATM UNI. Node A begins the point-to-multipoint call by sending a SETUP message to the network requesting setup of a point-to-multipoint call identifying leaf node B's ATM address. In the example, node A requests a call SETUP to node B, and the network responds with a CALL PROCEEDING message in much the same way as a point-to-point call. The network switches the call attempt to the intended destination and issues a SETUP message to node B identifying the assigned VPI/VCI. The first leaf node then indicates its intention to join the call by returning a CONNECT message that the network in turn acknowledges with a CONNECT ACKNOWLEDGE message. The network informs the calling root node A of a successful addition of party B through a CONNECT and CONNECT ACKNOWLEDGE handshake as shown in the figure.

Continuing on with the same example, the root node requests that party C be added through the ADD PARTY message, which the network relays to the same ATM UNI as party B through the ADD PARTY message to inform the local switch of the requested addition. Party C responds with an ADD PARTY ACKNOWLEDGE message that is propagated by the network back to the root node A. The root node A requests that the final leaf party D be added through an ADD PARTY message. The network routes this to the UNI connected to party D, and issues a SETUP message since this is the first party on this ATM UNI. Node D responds with a CONNECT message to which the network responds with a CONNECT ACKNOWLEDGE message. The fact that leaf party D has joined the point-to-multipoint call is communicated to the root node A through the ADD PARTY ACKNOWLEDGE message.

Figure 46.Point-to-Multipoint Call Setup Example

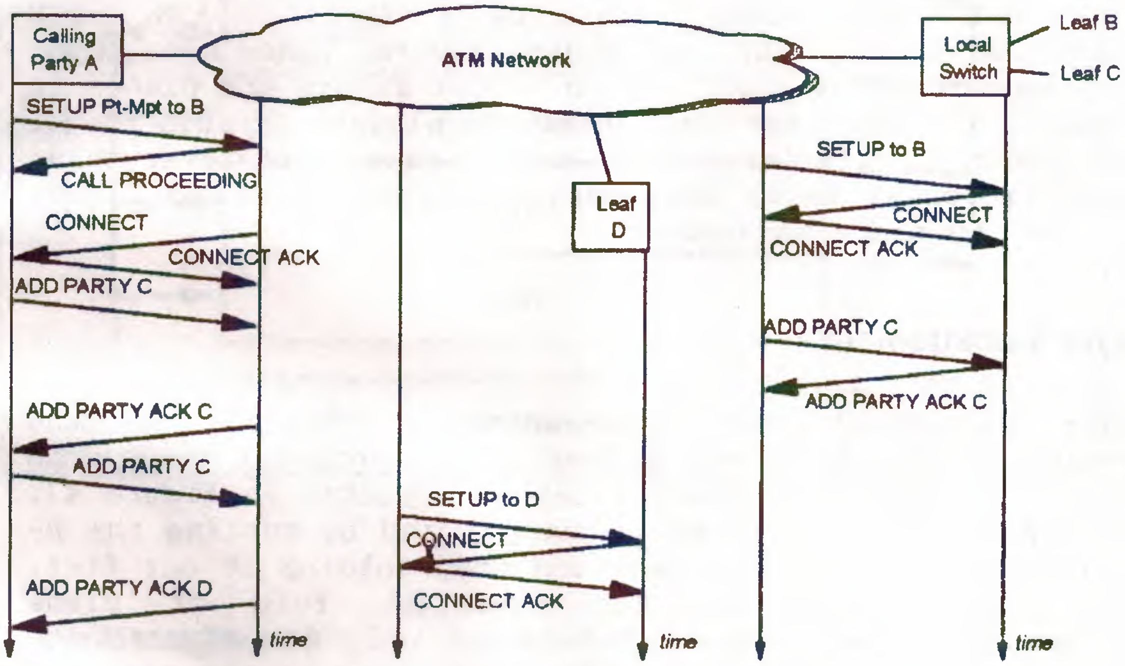

The leaves of the point-to-multipoint call may be removed from the call by the DROP PARTY message if one or more parties would remain on the call on the same UNI, or by the RELEASE message if the party were the last leaf present on the same UNI. The root node should drop each leaf in turn and then release the entire connection.

MANAGEMENT PLANE

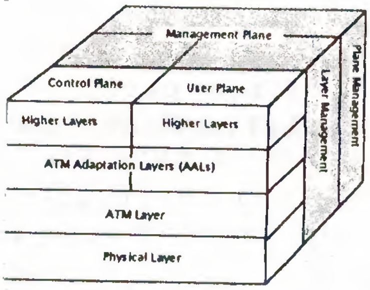

The management plane covers the layer management and plane management functions as shown in the B-ISDN cube layer management interfaces with the PHYsical, ATM, ATM Adaptation Layer (AAL), and higher layers. Plane management is responsible for coordination across layers and planes in support of the user and control planes through layer management facilities. This makes sure that everything works properly. Layer management will be discussed first, followed by plane management.

Layer Management

Layer management has a management interface to the PHYsical, ATM, AAL, and higher layer protocol entity in both the control and user planes as depicted in Figure 47. This two-dimensional view is constructed by cutting the B-ISDN cube open from the back and then folding it out flat. This view illustrates the oversight role of plane management as well. Plane management only interfaces with layer management, which provides interfaces to the user and control plane layers. Standards for these management interfaces are being defined by the ITU-T and ANSI for telecommunications equipment using the Common Management Information Protocol (CMIP), and by the IETF for data communications equipment using the Simple Network Management Protocol (SNMP).
Layer management has the responsibility for monitoring the user and control plane for faults, generating alarms, and taking corrective actions, as well as monitoring for compliance to the performance stated in the traffic contract. Layer management handles the operation and maintenance information functions found within specific layers. These functions include fault management, performance management, and configuration management. The

standards for PHY layer management are very mature. The standards for ATM layer fault and performance management are nearing the first stage of useability. Standardisation for management for the AAL and higher layers is just beginning.

Figure 47. Layer Management Relation to User and Control Planes

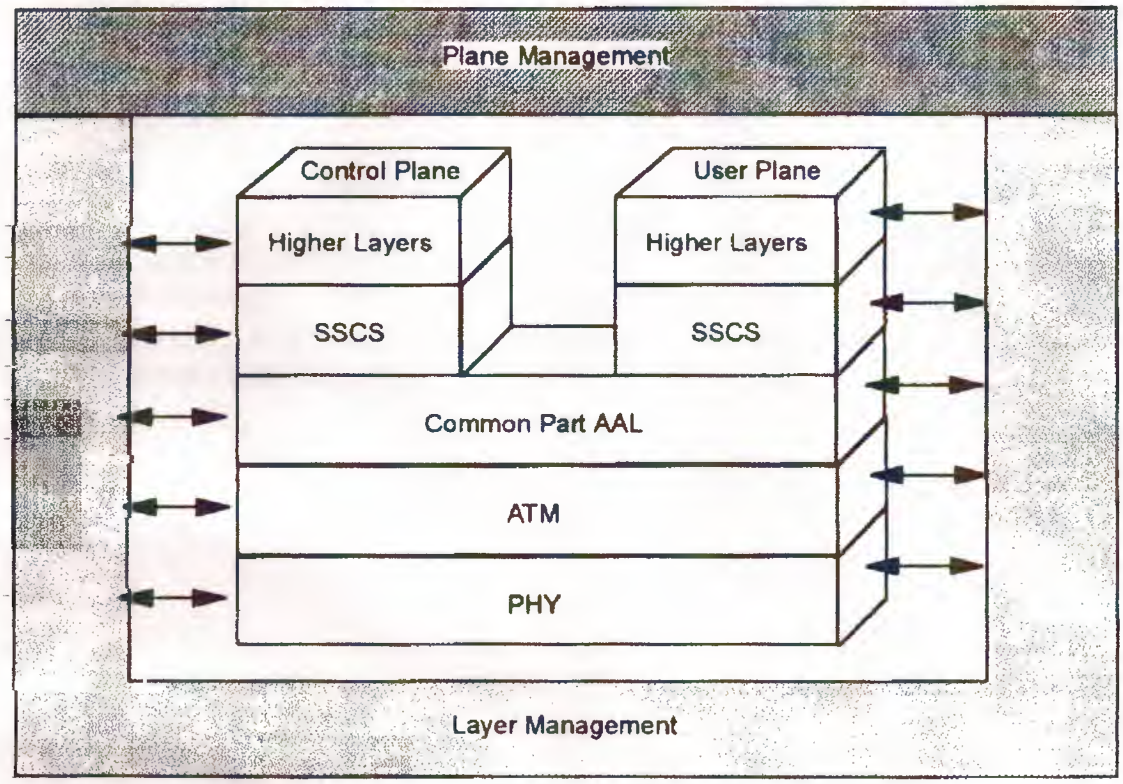

Plane Management

Plane management has no defined structure, but instead performs management functions and coordination across all layers and planes in the entire system. The Telecommunication Management Network (TMN) architecture developed by the ITU-T for managing all types of telecommunications networks is being extended to perform the B-ISDN plane management role.

References

- McDysan, David E.,"ATM:Theory and Application", McGrawHill, International Editions 1995
- E. Biagioni, E. Cooper, R.Sansom, "Designing a Practical ATM LAN", IEEE Network, March 1993.
- E.Ratheb, "Modeling and Performance Comparison of Policing Mechanisms for ATM Networks", IEEE Network, July 1991
- P.Boyer, F.Guillemin, M.Servel, J-P.Coudreuse, "Spacing Cells Protects and Enhances Utilization of ATM Network Links", IEEE Network, September 1992
- G.Woodruff, R.Kositpaiboon, "Multimedia Traffic Management Principles for Guaranteed ATM Network Performance", IEEE JSAC, April 1990
- G.Smith, J.P. van Steerteghem, "Seamless ATM Networks", Telecommunications, September 1993.