NEAR EAST UNIVERSITY

Faculty of Economics & Administrative Sciences

Department of Computer Informations Systems

ATM NETWORK AND B-ISDN

Graduation Project
COM-400

Student: S.Hüseyin Özgün

Supervisor: Assoc. Prof. Dr. Rahib Abiyev

Nicosia-2004

ADKNOWLEDGEMENTS

"First, I would like to thank my supervisor Prof. Dr. Fakhreddin Mamedow for his invaluable advice and belief in my work and myself over the course of this Graduation Project...

Second, I would like to express my gratitute to Near East University for the scholarship that made the work possible.

Third, I thank my family for their constant encouragement and support during the preparation of this project.

Finally, I would also like to thank all my friends for their advice and support."

ABSTRACT

Modern applications in today's Computer Networks include the use of multimedia applications such as Asynchronous Transfer Mode (ATM). Asynchronous Transfer Mode (ATM) has been recommended and been accepted by industry as the transfer mode for Broadband Integrated Services Digital Networks (B-ISDN). ATM has been designed to be able to handle different types of services and applications such as voice, data, image,text and video and mixture of all these.

ATM provides a good bandwidth flexibility and can be used efficiently from local area networks (LANs) and wide area networks (WANs). ATM is a connection-oriented packet switching technique in which all packets are of fixed length of 53 octets where 5 octets for header and 48 octets for information payload field. No processing like error control is done on the information field of ATM cells inside the network and it is carried transparently in the network.

LIST OF ABBREVIATIONS

ACM Address Complete Message: A BISUP call control message from the receiving exchange to sending exchange indicating the completion of address information.

ACR Attenuation to Crosstalk Ratio: One of the factors that limits the distance a signal may be sent through a given media. ACR is the ratio of the power of the received signal, attenuated by the media, over the power of the NEXT crosstalk from the local transmitter, usually expressed in decibels (db). To achieve a desired bit error rate, the received signal power must usually be several times larger than the NEXT power or plus several db. Increasing a marginal ACR may decrease the bit error rate.

AMI Alternate Mark Inversion: A line coding format used on T1 facilities that transmits ones by alternate positive and negative pulses.

ANI Automatic Number Identification: A charge number parameter that is normally included in the Initial Address Message to the succeeding carrier for billing purposes.

ANM Answer Message: A BISUP call control message from the receiving exchange to the sending exchange indicating answer and that a through connection should be completed in both directions.

ANSI American National Standards Institute: A U.S. standards body.

API Application Program Interface: API is a programmatic interface used for interprogram communications or for interfacing between protocol layers.
network address (MAC) corresponds to the IP address in the packet.

ASP Abstract Service Primitive: An implementation-independent description of an interaction between a service-user and a service-provider at a particular service boundary, as defined by Open Systems Interconnection (OSI).

ATM Asynchronous Transfer Mode: A transfer mode in which the information is organized into cells. It is asynchronous in the sense that the recurrence of cells containing information from an individual user is not necessarily periodic.

B-ISDN Broadband Integrated Services Digital Network (Broadband ISDN): A high-speed network standard (above 1.544 Mbps) that evolved Narrowband ISDN with existing and new services with voice, data and video in the same network.

(B-ISDN) A technology suite designed to transmit multimedia information (voice, data, video, image). Broadband ISDN specifications generally start at speeds of 150 Mbps using optical fibre (SONET) for the baseline transmissions. The two modes of B- ISDN are Synchronous Transfer Mode and Asynchronous Transfer Mode.

Call	A call is an association between two or more users or between a user and a network entity that is established by the use of network capabilities. This association may have zero or more connections.
CAS	Channel Associated Signaling: A form of circuit state signaling in which the circuit state is indicated by one or more bits of signaling status sent repetitively and associated with that specific circuit.
CBR	Constant Bit Rate: An ATM service category which supports a constant or guaranteed rate to transport services such as video or voice as well as circuit emulation which requires rigorous timing control and performance parameters.
CCR	Current Cell Rate: The Current Cell Rate is an RM-cell field set by the source to its current ACR when it generates a forward RM-cell. This field may be used to facilitate the calculation of ER, and may not be changed by network elements. CCR is formatted as a rate.
CDF	Cutoff Decrease Factor: CDF controls the decrease in ACR (Allowed Cell Rate) associated with CRM.
Cell	A unit of transmission in ATM. A Protocol Data Unit characterised by fixed, rather than variable length payloads. ATM cell consists of 5 octets header and 48 octests payload.
DA	Destination Address: Information sent in the forward direction indicating the address of the called station or customer.
DXI	Data Exchange Interface: A variable length frame-based ATM interface between a DTE and a special ATM CSU/DSU. The ATM CSU/DSU converts between the variable-length DXI frames and the fixed-length ATM cells.
FC	Feedback Control: Feedback controls are defined as the set of actions taken by the network and by the end-systems to regulate the traffic submitted on ATM connections according to the state of network elements.
GCAC	Generic Connection Admission Control: This is a process used during route selection to determine whether other nodes in the route are likely to have enough resources to support the connection, based on the advertised topology information.
GFC	Generic Flow Control: GFC is a field in the ATM header which can be used to provide local functions (e.g., flow control). It has local significance only and the value encoded in the field is not carried end-to-end.
HEC	Header Error Control: Using the fifth octet in the ATM cell header, ATM equipment may check for an error and corrects the contents of the header. The check character is calculated using a CRC algorithm allowing a single bit error in the header to be corrected or multiple errors to be detected.
IAM	Inverse Multiplexing over ATM: A process that allows multiple T1 or E1 communications facilities to be combined into a single broadband facility for the transmission of ATM cells.

IEEE Institute of Electrical and Electronics Engineers: A worldwide engineering publishing and standards-making body for the electronics industry.

IP Internet Protocol: Originally developed by the Department of Defense to support interworking of dissimilar computers across a network. This protocol works in conjunction with TCP and is usually identified as TCP/IP. A connectionless protocol that operates at the network layer (layer 3) of the OSI model.

ISDN Integrated Services Digital Network, a group of services based on evolving the telephony network to all digital services.

ISO International Standards Organisation: An international organization for standardization, based in Geneva, Switzerland, that establishes voluntary standards and promotes global trade of 90 member countries.

LAN Local Area Network: A network designed to move data between stations within a campus.

LGN Logical Group Node: LGN is the single node that represents a peer group (consisting of physical nodes and/or LGNs) in the next higher level peer group.

MAC Media Access Control: IEEE specifications for the lower half of the data link layer (layer 2) that defines topology dependent access control protocols for IEEE LAN specifications.

MAN Metropolitan Area Network: A network designed to carry data over an area larger than a campus such as an entire city and ts outlying area.Compare with WOMAN.Managed System An entity that is managed by one or more management systems, which can be either Element

MPEG Motion Picture Experts Group: An ISO Standards group dealing with video and audio compression techniques and mechanisms for multiplexing and synchronizing various media streams.

MTP Message Transfer Part: Level 1 through 3 protocols of the SS7 protocol stack. MTP 3 (Level 3) is used to support BISUP.

N-ISDN Narrowband Integrated Services Digital Network: Services include basic rate interface (2B+D or BRI) and primary rate interface (23B+D or PRI). Supports narrowband speeds at/or below 1.5 Mbps.

NE Network Element: A system that supports at least NEFs and may also support Operation System Functions/Mediation unctions. An ATM NE may be realized as either a standalone device or a geographically distributed system. It cannot be further decomposed into managed elements in the context of a given management function.

NEL Network Element Layer: An abstraction of functions related specifically to the technology, vendor, and the network resources or network elements that provide basic communications services.

NNI Network Node Interface: An interface between ATM switches defined as the interface between two network nodes.

NT	Network Termination: Network Termination represents the termination point of a Virtual Channel, Virtual Path, or Virtual Path/Virtual Channel at the UNI.
OSI	Open Systems Interconnection: A seven (7) layer architecture model for communications systems developed by the ISO for the interconnection of data communications systems. Each layer uses and builds on the services provided by those below it.
PAD	Packet Assembler and Disassembler: A PAD assembles packets of asynchronous data and emits these buffers in a burst to a packet switch network. The PAD also disassembles packets from the network and emits the data to the non-packet device.
PDU	Protocol Data Unit: A PDU is a message of a given protocol comprising payload and protocol-specific control information, typically contained in a header. PDUs pass over the protocol interfaces which exist between the layers of protocols (per OSI model).
PICS	Protocol Implementation Conformance Statement: A statement made by the supplier of an implementation or system stating which capabilities have been implemented for a given protocol.
SE	Switching Element: Switching Element refers to the device or network node which performs ATM switching functions based on the VPI or VPI/VCI pair.
SIP	SMDS Interface Protocol: Protocol where layer 2 is based on ATM, AAL and DQDB. Layer 1 is DS1 and DS3.
TC	Transaction Capabilities: TCAP (see below) plus supporting Presentation, Session and Transport protocol layers.
TCAP	Transaction Capabilities Applications Part: A connectionless SS7 protocol for the exchange of information outside the context of a call or connection. It typically runs over SCCP and MTP 3.
TE	Terminal Equipment: Terminal equipment represents the endpoint of ATM connection(s) and termination of the various protocols within the connection(s).
UDP	User Datagram Protocol: This protocol is part of the TCP/IP protocol suite and provides a means for applications to access the connectionless features of IP. UDP operates at layer 4 of the OSI reference model and provides for the exchange of datagrams without acknowledgements or guaranteed delivery.
VD	Virtual Destination. Refer to VS/VD.
VPI	Virtual Path Identifier: An eight bit (on the UNI) or twelve bit (on the NNI) field in the ATM cell header which indicates the virtual path over which the cell should be routed.
WAN	Wide Area Network: This is a network which spans a large geographic area relative to office and campus environment of LAN (Local Area Network). WAN is characterized by having much greater transfer delays due to laws of physics.

INTRODUCTION

This subject component is developed as part of the course module to provide you with information on advanced topics in ATM and Broadband Communications Technology.This component covers a range of topics that provides you with a comprehensive view of Asynchronous Transfer Mode (ATM) and Broadband Integrated Services Digital Network (B-ISDN) technology. The notes include the following main sections:

- Introduction,
- ATM Principle,
- ATM standards,
- B-ISDN Protocol Reference Model (PRM) for ATM,
- ATM Switches,
- ATM Interfaces and Networks,
- Traffic Management and Control in ATM Networks

If you have any comment or suggestion about this component, you are welcome to send it to me.

Asynchronous Transfer Mode (ATM) has been recommended and been accepted by industry as the transfer mode for Broadband Integrated Services Digital Networks (B-ISDN). ATM has been designed to be able to handle different types of services and applications such as voice, data, image,text and video and mixture of all these.

ATM provides a good bandwidth flexibility and can be used efficiently from local area networks (LANs) and wide area networks (WANs). ATM is a connection-oriented packet switching technique in which all packets are of fixed length of 53 octets where 5 octets for header and 48 octets for information payload field. No processing like error control is done on the information field of ATM cells inside the network and it is carried transparently in the network.

CHAPTER ONE : ATM and B-ISDN

1.1 ATM Fundamental Concept

From a technical point of view, the fundamental underpinning of ATM is:

- to support all existing services as well as emerging services in the future,
- fixed-size cells with VPI and VCI to minimises the switching complexity,
- statistical multiplexing to utilises network resources very efficiently,
- to minimise the processing time at the intermediate nodes and supports very high transmission speeds as well as very low speed by negotiate service contract for a connection with required quality of services,
- to minimise the number of buffers required at the intermediate nodes to bound the delay and the complexity of buffer management,
- guarantees performance requirements of existing and emerging applications,
- layered architecture, and
- Capable of handling bursty traffic.

1.2 Bursty Traffic

Consider voice. If you were to observe a speaker in a normal two-way conversation and observed the acoustic energy, you'd find patterns where there would be times when there's silence because a person is pausing, and of course, there would be times when a person is listening. It's been known since the early '60s that in a typical two-way conversation, one of the participants will only be generating acoustic energy about 40% of the time. This has actually been exploited now for a number of years for things like undersea telephone cables where the cables are very expensive.

By transmitting information only during the non-silent periods, you can effectively double the capacity. This illustrates an interesting observation. Voice is typically based on circuit-switched technology. The circuit switching is really an artifact of the way the technology evolved over the years. You could get significant efficiencies if during the "silent" periods you could use the transmission capabilities (bandwidth) for other information such as data or e-mail.

Consider other forms of communication. Data communications are very typically bursty and video will have very much the same properties. Note: 1 octet = 8 bits. In ITU-T (formally CCITT) recommendations octet is used to define the ATM specifications. But byte is also used in many publications instead. They are the same and exchangeable here.

1.3 ATM Cell

The traditional way to deal with bursty data is some form of packet switching technology. In ATM, the packet is called a cell because all packets in a given network are of fixed length. In ATM, the defined cell has 48 bytes of payload and 5 bytes of header for a total of 53 bytes, or octets. The 5 byte of header has to have enough information in them to allow the network to forward each cell to the proper destination. An ATM network is an infrastructure that is good at handling cells, and that's all it has to do. Containerised shipping is a good analogy.

1.4 ATM as Universal Transport Networks

All voice, data and video can be digitised. Then a device can converted the original digitised format into the cells. This device appropriately marks the cells, and the cells are multiplexed onto one link going into ATM networks. Then the network can switch and deliver the cells to the appropriate destinations. Notice that with the network, or the cell switching infrastructure, in place, if you want to add a new traffic type, you don't have to change the infrastructure. You only need to deal with the devices at the edge of the network. In fact, you only have to deploy those devices where you need the new traffic type. This is one of the really attractive aspects of ATM. Only minimal changes are needed when there is a need to change or to improve the network.

1.5 Broadband ISDN

Sometimes there's a bit of confusion between ATM and Broadband ISDN (B-ISDN). The two are related because ATM evolved from the standardisation efforts for B-ISDN. The exact relationship is that ATM is the technology upon which B-ISDN is based.

Usually the term "B-ISDN" is applied to wide area carrier services. Even though the technology foundation is ATM, as we go forward, the term B-ISDN is not usually applied to local area or campus networks. Nonetheless, because they're based on the same technology, the advantages of ATM will still apply.

1.6 Carrying a Bit Stream

Consider what the devices that convert information to an ATM format have to do, beginning with voice. Although voice is bursty, all of the standard encoding technologies today use a continuous bit output -- 64 Kbit/s, for example -- so that's what's represented up on this link. This continuous stream of information has to be converted into cells. These devices "slice" the information into discrete cells for transmission across the network.

In this example, the bits are pulled off piece by piece -- in this case, just eight bits per cell and put into cells and sent into the network. Notice that there are some empty cells because the rate of bits, at 64 Kbit/s, is typically going to be far less than the bit rate on the link. Consequently, there are "empty" cells. The cells traverse the network and come out on the output link.

Notice that the spacing between the cells has changed slightly. This reflects the statistical or asynchronous nature of ATM, and it's not particularly a problem so long as the variations are relatively minor. (In section 3.3 and 4, some of the effects of variation in delay of cells will be examined.).

The ATM network also keeps the cells in order. A cell may be dropped occasionally, but the cells will stay in order, and so it's easy to visualise on how the inverse device (at the receiving end) can take the input stream of cells and recreate the original bit stream.

1.7 Carrying Packet Data

For data communications, where you're probably already dealing with packets, the process is very similar. Here we see though, a little bit of difference because the data will come in big "chunks." Thus, the cells will tend to be close together when they go into the network, and they may spread out a bit again because of the statistical nature. Nevertheless, reassembling the cells back into the original packets is conceptually straightforward.

1.8 Layered Architecture

As with any new technology, though, nothing is really quite that simple. We must consider that ATM is based on a layered architecture. We'll begin with the physical layer because you still need a physical layer for ATM communications. (The various physical layers for ATM are examined in section 3.2). On top of the physical layer is the ATM layer. This is where we find the cells. On top of the ATM layer is what's called the ATM Adaptation Layer (AAL). It runs from end system to end system and is transparent to the ATM network.

The information from the AAL goes through the network in the sense that the protocol information associated with it is in the 48 bytes of the cell, but the network does not look at the AAL information at all.

Of course, as we implement new kinds of traffic on top of ATM, we will see more different kinds of AAL protocols. Several of these will be examined in section 3.4.

CHAPTER TWO : ATM PRINCIPLE

ATM is a fast packet oriented transfer mode based on asynchronous time division multiplexing and it uses fixed length (53 bytes) cells. Each ATM cell consists of a information field (48 bytes) and a header (5 bytes). The header is used to identify cells belonging to the same virtual channel and thus used in appropriate routing. Cell sequence integrity is preserved per virtual channel. ATM Adaptation layers (AAL) are used to support various services and provide service specific functions. This AAL specific information is contained in the information field of the ATM cell. Basic ATM cell structure is used for the following functions.

2.1 Routing

ATM is a connection oriented mode. The header values (i.e. VCI and VPI etc.) are assigned during the connection set up phase and translated when switched from one section to other. Signalling information is carried on a separate virtual channel than the user information. In routing, there are two types of connections, i.e., Virtual channel connection(VCC) and Virtual path connection(VPC).

A VPC is an aggregate of VCCs. Switching on cells is first done on the VPC and then on the VCC.

2.2 ATM Resources Management

ATM is connection-oriented and the establishment of the connections includes the allocation of a virtual channel identifier (VCI) and/or virtual path identifier (VPI). It also includes the allocation of the required resources on the user access and inside the network. These resources, expressed in terms of throughput and quality of service, can be negotiated between user and network either before the call-set up or during the call.

2.3 ATM Cell Identifiers

ATM cell identifiers including VPI, VCI and Payload Type Identifier (PTI) are used to recognise an ATM cell on a physical transmission medium. VPI and VCI are same for cells belonging to the same virtual connection on a shared transmission medium.

2.4 Throughput

Peak Cell Rate (PCR) can be defined as a throughput parameter which in turn is defined as the inverse of the minimum interarrival time T between two consecutive basic events and T is the peak emission interval of the ATM connection. PCR applies to both constant bit rate (CBR) and variable bit rate (VBR) services for ATM connections. It is an upper bound of the cell rate of an ATM connection and there is another parameter sustainable cell rate (SCR) allows the ATM network to allocate resources more efficiently.

2.5 Quality Of Service

Quality of Service (QOS) parameters include cell loss, the delay and the delay variation incurred by the cells belonging to the connection in an ATM network. QOS parameters can be either specified explicitly by the user or implicitly associated with specific service requests. A limited number of specific QOS classes will be standardised in practice.

2.6 Usage Parameter Control

In ATM, excessive reservation of resources by one user affects traffic for other users. So the throughput must be policed at the user-network interface by a Usage Parameter Control (UPC) function in the network to ensure that the negotiated connection parameters per VCC or VPC between network and subscriber is maintained by each other user. Traffic parameters describe the desired throughput and QOS in the contract. The traffic parameters are to be monitored in real time at the arrival of each cell. ITU-T (formerly CCITT) recommends a check of the peak cell rate (PCR) of the high priority cell flow (CLP = 0) and a check of the aggregate cell flow (CLP = 0 and 1), per virtual connection.

2.7 Flow Control

In order to control the flow of traffic on ATM connections from a terminal to the network, a General Flow Control (GFC) mechanism is proposed by ITU-T at the User to Network Interface (UNI). This function is supported by GFC field in the ATM cell header. Two sets of procedures are associated with the GFC field, i.e., Uncontrolled Transmission which is for use in point-to-point configurations and Controlled Transmission which can be used in both point-to-point and shared medium configurations.

CHAPTER THREE : ATM STANDARDS

3.1 B-ISDN Protocol Reference Model (PRM) for ATM

B-ISDN PRM consists of 3 planes. a User plane for transporting user information, a Control plane which is responsible for call control and connection control functions and it contains mainly signalling information, and a Management plane which contains layer management functions and plane management functions. There is no defined (or standardised) relationship between OSI layers and B-ISDN ATM protocol model layers. But the following relations can be found. The Physical layer of ATM is almost equivalent to layer 1 of the OSI model and it performs bit level functions.

ATM layer can be equivalent of the lower edge of the layer 2 of the OSI model. The ATM Adaptation Layer performs the adaptation of OSI higher layer protocols. The B-ISDN Protocol Reference Model Sublayers and functions can be shown by the following figure.

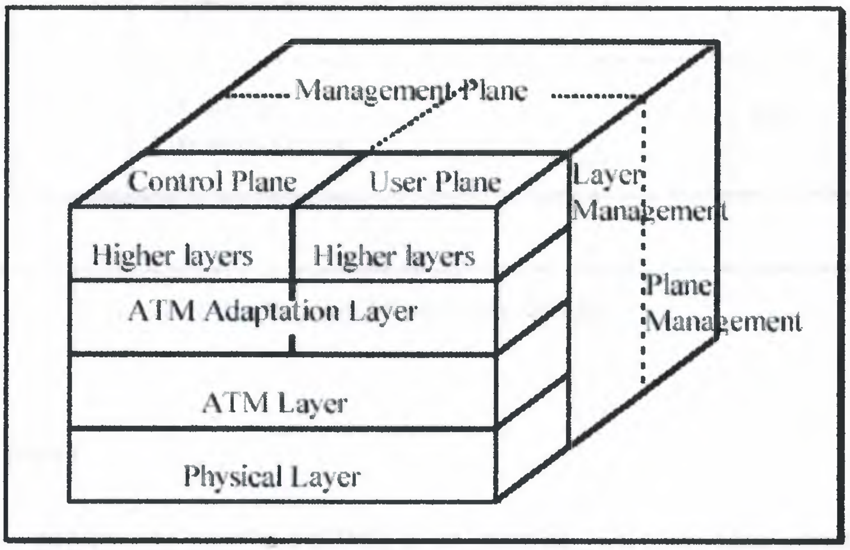

Figure 3-1. B-ISDN ATM Protocol Reference Model.

Layer Management	Higher Layer Functions		
	Convergence Sublayer	CS	AAL
	Segmentation and Reassembly	SAR	
	Generic Flow Control Cell header generation/extraction Cell VPI/VCI Translation Cell Multiplexing and Demultiplexing		ATM
	Cell rate decoupling HEC header generation/verification Cell delineation Transmission frame adaption Transmission frame generation/recovery	TC	Physical Layer
	Bit timing Physical Media	PM	

Figure 3-2. Layers of the Protocol Reference Model.

3.2 Physical layer

The User-to-Network Interface (UNI) is an interface between data terminal equipment (DTE) that implements ATM protocols and an ATM network. The first requirement for interpretability of the terminal equipment with the ATM network is to transmit information successfully at the physical level.The physical layers that have been specified by The ATM Forum or are under development. They include multi-mode fibre, single-mode fibre, shielded-twisted pair, coaxial cable and some unshielded-twisted pair. We'll explore the different interfaces from section 3.2.2 to 3.2.6. As shown in the above figure, Physical Layer is divided into two sub-layers: Physical Medium (PM) sub-layer, and Transmission Convergence (TC) sub-layer.

3.2.1 Two Sublayers

There is an important architecture difference in the physical layer for ATM from what we traditionally see as the OSI physical layer. If you think about the traditional physical layer, it essentially deals with the "atoms." That is, it deals with the bits and its job is to get a bit from place to place.

The PM sub-layer contains only the Physical Medium (PM) dependent functions (such as bit encoding; the characteristics of connectors; the property of the fibre optics, etc.). It provides bit transmission capability including bit alignment. It performs Line coding and also electrical/optical conversion if necessary. Optical fibre will be the physical medium and in some cases, coaxial and twisted pair cables are also used. It includes bit timing functions such as the generation and reception of waveforms suitable for the medium and also insertion and extraction of bit timing information. But in ATM, the "atom" is actually the cell. The smallest thing that an ATM network is ever going to deal with is the cell.

Therefore certain cell functions have been put into the physical layer, and those functions are called the transmission convergence (TC) sublayer.

You can kind of think of an ATM network as a very insistent seal, which is always barking for a fish, or in this case, a cell. And if you have a cell to send because you have data to send, you throw the cell to the network. If you don't, if there's nothing to send, you still have to send a cell into the network to keep the network healthy. It's an "empty" cell, and the TC sublayer's job also is to insert empty cells for transmission and remove empty cells when they get to the destination in order to keep the cell streams constant. Because of the different kinds of details in the coupling between the fibre or other physical medium, the TC sublayer is different, depending on the physical layer that's being discussed.

The TC sub-layer mainly does five functions as shown in the figure. The lowest function is generation and recovery of the transmission frame.

The next function, i.e. transmission frame adaptation takes care of all actions to adapt the cell flow according to the used payload structure of the transmission system in the sending direction. It extracts the cell flow from the transmission frame in the receiving direction. The frame can be a synchronous digital hierarchy (SDH) envelope or an envelope according to ITU-T Recommendation G.703.

Cell delineation function enables the receiver to recover the cell boundaries from a stream of bits. Scrambling and Descrambling are to be done in the information field of a cell before the transmission and reception respectively to protect the cell delineation mechanism.

The HEC sequence generation is done in the transmit direction and its value is recalculated and compared with the received value and thus used in correcting the header errors. If the header errors can not be corrected, the cell will be discarded.

Cell rate decoupling inserts the idle cells in the transmitting direction in order to adapt the rate of the ATM cells to the payload capacity of the transmission system. It suppresses all idle cells in the receiving direction. Only assigned and unassigned cells are passed to the ATM layer.

3.2.2 155 Mbps, SONET STS-3c

Let's start with SONET, which is probably the physical layer most often associated with ATM. The essential feature of SONET is to keep track of boundaries of streams that don't really depend on the particular medium. So, although we typically think about it as fibre, it will in fact operate over other media. Some of the work going on currently in **The ATM Forum** on a physical specification for using (copper) unshielded-twisted pair will be using the SONET type framing. This is the SONET frame at 155 Mbps. To read this chart, start in the upper left-hand corner. The

bytes are transmitted across the medium a row at a time, wrapping to the next row. By the time you go through all nine rows, the elapsed time is nominally 125 microseconds.

In Figure 3-3. STS-3C (also know as SDH STM-1) Frame., the first 9 bytes of each row have various overhead functions. For example, the first two bytes here are used to identify where the beginning of this frame is so the receiver can lock on to this frame.

In addition, although not shown here, there is another column of bytes which are included in the "Synchronous Payload Envelope" that are additional overhead, with the result that each row has 260 bytes of information. Consequently, 260 bytes per row times 9 rows times 8 bits divided by 125 microseconds, you get 149.76 Mbps of payload.

This is called the STS-3C. It is also known as the STM-1 because in the international carrier networks, this will be the smallest package that you'll see available in terms of the Synchronous Digital Hierarchy (SDH), the international flavour of SONET. The bit rates for SDH STM-n are three times the bit rates for SONET STS-n for the same "n."

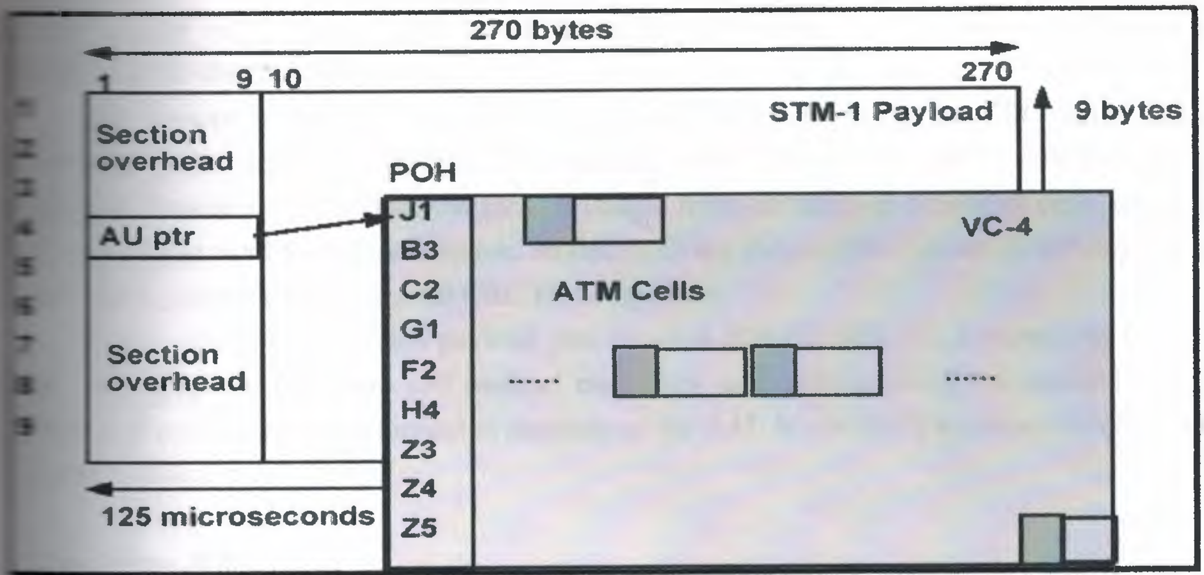

Figure 3-3. STS-3C (also know as SDH STM-1) Frame.

3.2.2.1 For Higher Bit Rate

SONET also has some nice features in that if you want to go to higher rates -- like 622 Mbps - it becomes basically a recipe of how you take four of these STM-1 structures and simply interleave the bytes to get to 622 Mbps (STM-4, or STS-12). There are additional steps up to 1.2 gigabits, 2.4 gigabits, etc. And -- at least in theory -- the recipe tells you how to get as high a speed interface as you would like.

3.2.2.2 SONET Cell Delineation

The cells within the SONET payload are delineated by using the Header Error Check (HEC) in the ATM cell.

The receiver, when it's trying to find the cell boundaries, takes five bytes and says, "I wonder if this five bytes is a header." It does the HEC calculation on the first four bytes and matches that calculation against the fifth byte. If it matches, the receiver then counts 48 bytes and tries the calculation again. And if it finds that calculation correct several times in a row, you can probably safely assume that in fact it's found the cell boundaries. If it tries the calculation and it fails, you just slide the window and try the calculation again.

This kind of process must be used because, of course, we don't really know what's in the 48 bytes of payload, but the chances that the user data would contain these patterns separated by 48 bytes is essentially zero for any length of time.

Consider for a moment what happens if you come across a series of empty cells. Then how do you determine the cell boundaries? This is especially important since the a CRC for an all zero (empty cell) header would be all zeros. Consequently, the HEC must be based on something other than a simple CRC.

The answer is that the HEC is calculated by first calculating the CRC value, then performing an "exclusive or" operation of the CRC value with a bit pattern called the coset, resulting in a nonzero HEC. Thus, the HEC is unique from the zeros in the empty cells, and the HEC may still be used for cell delineation. At the receiving end, another "exclusive or" operation is performed, resulting in the original CRC for comparison.

If you calculate how much payload you get on a SONET STS-3C, it comes out to 135 Mbps, assuming that the entire cell payload may carry user information. (The amount of the payload that actually carries information depends on the AAL in use. We'll examine the different AALs in section 3.4).

3.2.3 44 Mbps, DS3

DS3 is another important interface because it's probably going to be the dominant high-speed ATM interface in the public carriers in North America for the next few years. Although SONET is a very nice technology, SONET is still not ubiquitously deployed in typical carrier networks. It simply takes a long time to get new technology deployed, and DS3 is already widely deployed. As the early services roll out for the wide area, they will usually be on DS3.

This shows a transmission frame. These are bits now (rather than the SONET frame bytes), but you still read it the same way: You read across a row. There's an overhead bit, followed by 84 bits of data, followed by an overhead bit and so forth. This whole structure has a nominal time interval of 106.4 microseconds. If you go through the calculation, this turns out to be a payload of 44.21 Mbps.

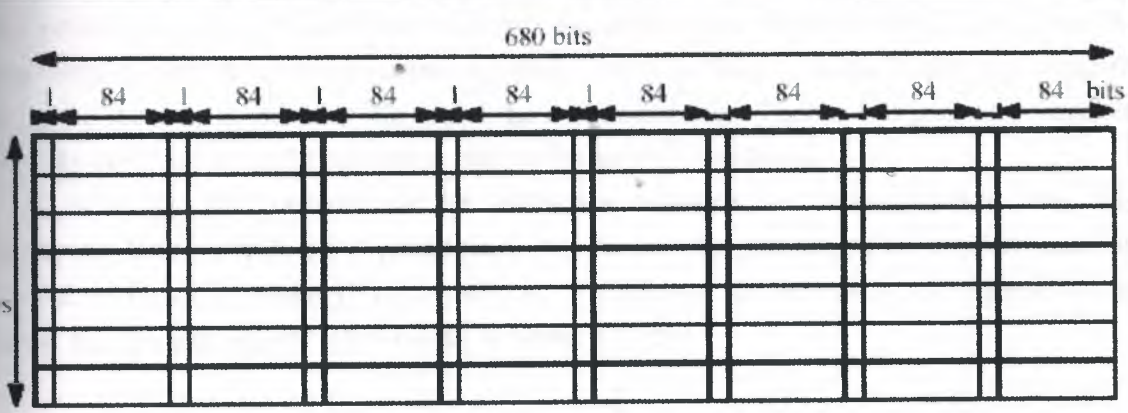

The frame time interval is 106.4 microseconds.

Figure 3-4. The DS3 frame.

3.2.3.1 DS3 Cell Delineation

Now we have the bits coming in 84-bit blocks. If we view that as a continuous stream of bits, we have to find the cell boundary. In addition, it turns out to be nice to have a frame which is 125 microseconds (for compatibility with SONET and with traditional digital telephony). The technique used for DS3 accomplishes both the cell delineation and the 125 microsecond frame.

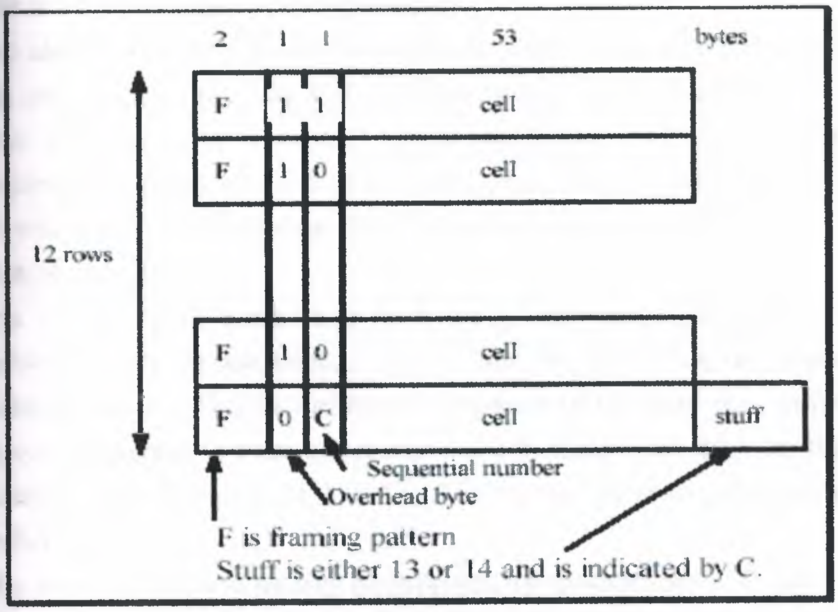

Figure 3-5. DS3 Cell Delineation.

This is accomplished by use of a framing pattern of two bytes inserted periodically followed by a sequence number and then there's an overhead byte then followed by a cell. A receiver trying to find the cell boundaries will be looking for this framing pattern. When it finds a framing pattern, it looks for a sequence number and then counts 54 bytes and looks for another framing pattern and the right sequence number. When it finds two such patterns separated properly, the receiver deems that it's found the cell boundaries. Then it can just keep checking this to make sure that the cell boundaries start at this point in the frame.

It turns out that because of the difference between 125 microseconds and 106.4 microseconds, there's a little bit of variation, so the amount that gets put on the end here can vary from frame to frame, between 13 or 14 nibbles. Once all of the various types of overhead are accounted for, the cell payload turns out to be about 37 Mbps.

2.2.4 155 Mbps, 8B/10B Coding

Another 155 Mbps interface is based on what's called 8B/10B coding. This rate is intentionally exactly the SONET STS3 rate, but the method of finding the cell boundaries and the method of encoding the bits is a bit different.

This physical interface is being used for multi-mode fibre, which can go up to 2 kilometers, and for shielded-twisted pair, which can go up to 100 meters. Obviously the shielded-twisted pair is an inbuilding kind of application.

The idea here is to use what's called the block code. For every opportunity to send 10 changes in the state of the physical medium (technically a "baud"), you only send 8 bits of information. Thus, you use "symbols," where you conceive of 10-baud blocks as symbols that stand for something. With 10 baud of symbol space, you get 1024 symbols. These symbols can be used to encode the 8 bits of data (256 data symbols), and some of the remaining symbols for special control purposes.

You can also pick symbols in such a way that your frequency spectrum is nice and it avoids problems with the emissions, etc. To find the cell boundaries, there's a special five-symbol pattern called a "frame delimiter." Because of the way the symbols are defined, this cannot appear in anybody's data. Consequently, it becomes very easy for the receiver to find the cell boundaries. Once the special OA&M cell is detected, next 26 cells are "user cells" containing the user information.

The last user cell will usually be followed by another special frame delimiter (fivesymbol pattern) to start the next "frame" of cells. Following the five-symbol frame delimiter, there are 48 data symbols reserved (but as of this time undefined) for OA&M functions. Thus, we have the repeating pattern of a special symbol; OA&M cell; 26 "user cells"; then another 5-symbol pattern.

The cell payload or the cell rate turns out to be 135.6 Mbps, which is exactly equal to the STS3. This can be very convenient. For example, maybe you want to use the shielded-twisted pair to go from a desktop to a wiring closet, but then you want to convert to SONET, say, to go across campus. This could be done with just a physical layer converter box. Because you have exactly the same cell rates in and out, you don't need any buffers for speed matching. It is interesting to note that this block code was not invented by The ATM Forum; it was used from another standard called Fibre Channel. The ATM Forum does not really want to invent new things.

If there are existing standards that can be simply referred to, then that's the way to go. The Forum has a very practical view of how to establish interpretability. If there are standards that are in place, it's much better to use them than to try to write new material.

3.2.5 100 Mbps, 4B/5B Coding

There's also a 100 Mbps physical layer. One of the reasons this exists is that it is basically using the FDDI technology. FDDI uses a block coding technique using 5-bit (baud) blocks to encode 4 its of information. There are 16 symbols used for data, and there are 16 remaining symbols used for control.

The technique for finding the cells is to define a symbol called the "TT" symbol -- and that's inserted in front of every cell. This becomes very easy. Since this symbol pair cannot appear in any data, it will not appear anyplace else in a stream. The receiver needs only to scan for the first "TT" symbol. It then has locked into the cell boundaries immediately. If you go through the calculations of cell payload, this is effectively 6 bytes of overhead for every 48 bytes of payload. Thus, this yields a little less than 89 Mbps for the cell payloads.

3.2.6 1.544 Mbps DS-1

DS1 is offered by the public carriers in North America. It is also specified by the ATM Forum to carry ATM traffic. The standard DS1 format consists of 24 consecutive bytes with a single overhead bit inserted for framing. There's a fixed pattern for these overhead bits to identify the framing bits and the frame structure. (see Figure 3-6).

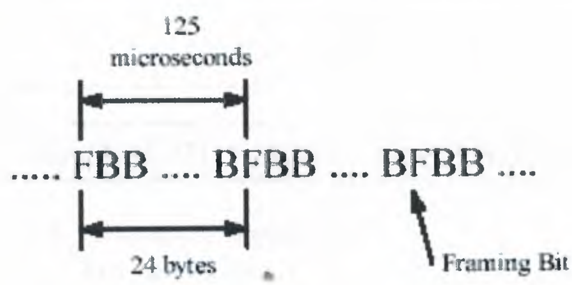

- (24 byte x 8 bit / byte) / 125 microsenconds = 1.536 Mbit/s of payload
- Cell delineation by HEC detection
- Cell payload = 1.536 Mbit/s x (48/53) = 1.391 Mbit/s

Figure 3-6. 1.544 Mbit/s, DS1 Frame structure.

Once that pattern is identified, you now know where the bytes within the DS1 physical layer payload are. Now, the question is how to find the cell boundaries.

The cells are going to be put into these physical layer payload bytes. Notice that there are only 24 bytes in each of these blocks, so the cell is actually going to extend across multiple blocks. There could be 24 bytes of a cell in the first block, 24 bytes of the same cell in the second block, and then the remaining five bytes of the cell in the third block. However, the cell actually can fall anywhere on the byte boundaries, so how do we find the cell boundaries?

The answer: Use the same mechanism as is used with SONET. Keep looking at five-byte windows and doing the CRC calculation, use the header error check approach. The actual payload that can be transported within a DS1 is about 1.4 Mbps.

3.2.7 2.048 Mbps E1

The 2.048 Mbps interface will be particularly important in Eurpoe, where this speed (E1), is the functional equivalent of North American DS1 interfaces. Note that in contrast to the DS1 format, there are no extra framing bits added. In fact, the 2.048 Mbps rate is an exact mulitple of 64 kbps.

The basic E1 frame consists of a collection of 32 bytes, recurring every 125 microseconds. Instead of using framing bits, this format uses the first (Byte 0) and seventeenth (Byte 16) for framing and other control information. The receiver uses the information within the framing bytes to detect the boundaries of the physical layer blocks, or frames. The remaining 30 bytes are used to carry ATM cells. Consequently, the physical layer payload capacity for the E1 interface is 1.920 Mbps. (see figure).

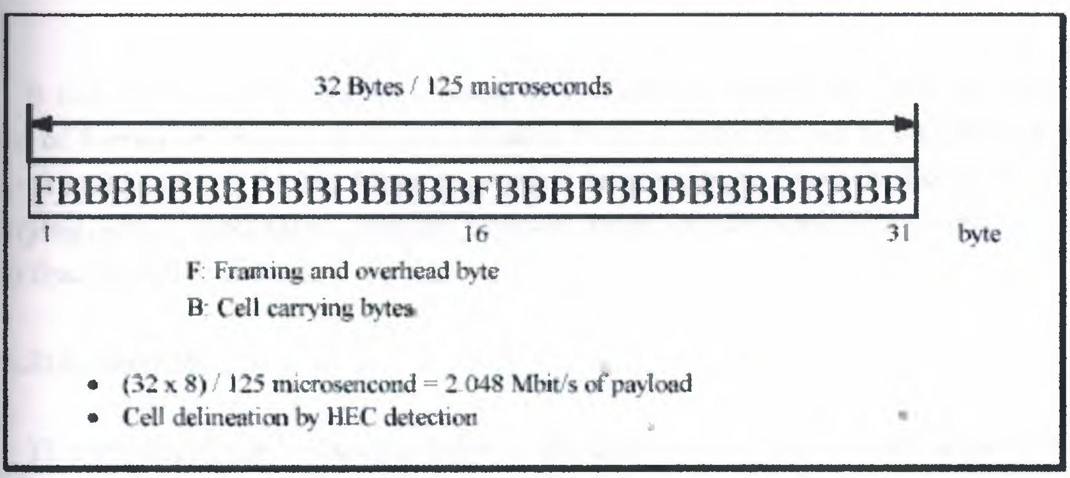

Figure 3-7. 2.048 Mbit/s E1 frame structure.

Just as in SONET and DS1, as previously discussed, the HEC is used to find the cell boundaries.

3.2.8 34 Mbps E3

Just as the E1 interface is the functional equivalent of the DS1 interface, the E3 interface is the functional equivalent of the DS3 interface.

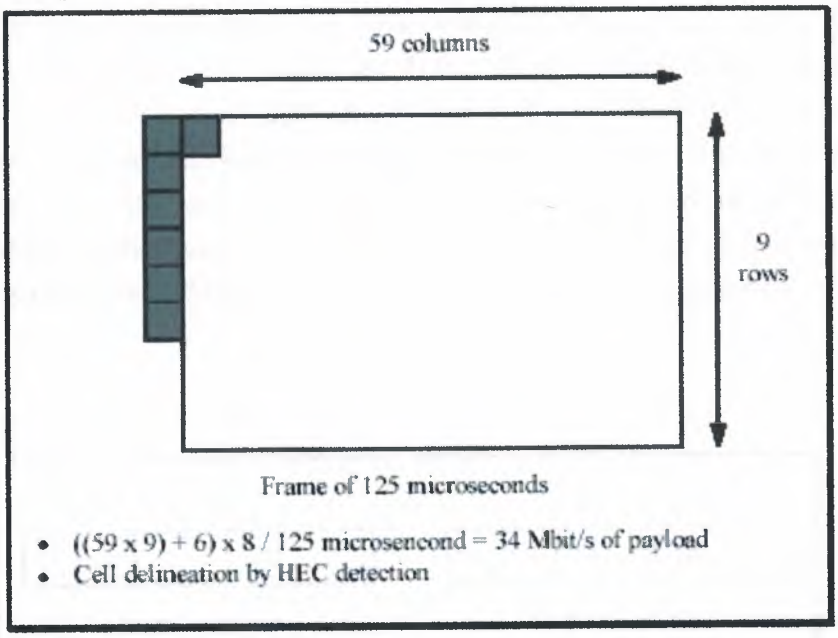

Figure 3-8. 34 Mbit/s E3 frame structure.

It is a single 125 microsecond frame, so this pattern recurs 8,000 time each second. It consists of 9 rows of 59 bytes each, plus 6 extra framing and overhead bytes. The result, once you do the arithmetic, is 34.368 Mbps of physical layer capacity. The actual capacity available for carrying cells is 33.92 Mbps, once the overhead bytes are subtracted. Once again, the HEC is used to find the cell boundaries.

3.2.9 6.312 Mbps J2

This physical layer is found in Japan in the carrier networks. Even though the bit rate is the same as the DS2 rate in the North American hierarchy, the actual assignment of the bits is different. It consists of four 125 microsecond frames. Each individual frame has 98 eight-bit slots (784 bits) plus 5 bits (F bits) of overhead. The last two slots are reserved leaving 6.144 Mbps for carrying cells.

As usual, the receiver looks for a well-known pattern in the F bits to find physical level frame boundary. Also as usual, the HEC is used to find the cell boundaries.

3.2.10 25.6 Mbps UTP-3

Turning to the private UNI, the lowest speed interface is the 25 Mbps interface over UTP-3. This is designed as a physical layer that can use the typical existing wiring within the office environment, such as between the wiring closet and the desktop. Thus, this is targeted at desktop ATM. In fact, this actually just takes the Token Ring physical layer and does a couple of interesting tricks with it. In particular what it does is it uses what's called a 4B/5B block code. Every five bits in the physical layer are considered a five-bit block, and this actually represents a four-bit pattern. Thus, we have 32 possible five-bit symbols. Sixteen of the symbols will be for data and 16 of the symbols can be used for other things such as control. The reason of doing this is it effectively takes the 16 Mbps Token Ring rate and makes it 25 megabits.

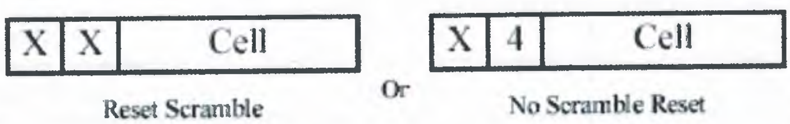

- Use IEEE 802.5 physical layer with 4B/5B coding plus scrabling
- 32 Mbaud x 4/5 = 25.6 Mbit/s of payload
- Cell delineated by special symble pairs

Figure 3-9. 25.6 Mbit/s UTP-3 frame structure.

Defining., or delimiting, cells is very easy here. You define a brand-new symbol called the X symbol, which will never show up in the cell because the cell is all data and always uses symbols from the 16 "data" symbols. So, whenever a receiver sees an X symbol -- actually sees two X symbols -- it knows that what follows is the symbols for the rest of the cell. (see figure). It turns out there's another technical detail here. There's a scrambling technique which helps make the spectrum of frequency a little smoother. There are two ways to do this scrambling. One is to not reset the scrambling. In this case, the transmitter will use two Xs. If the receiver is to reset the scrambler, you put an X followed by a 4 data symbol. But again, you'll never find the X within the cell because it's not a data symbol, so it's very easy to lock onto the cell boundaries.

3.2.11 Other Possibilities

The ATM Forum has tried to specify all the possibilities of physical layers for ATM. Probably the most important ongoing work in The ATM Forum in this area is defining ATM for unshieldedtwisted pair. There's a lot of unshielded- twisted pair deployed in existing buildings. ATM will be eaier to deploy if you can make that twisted pair work because you will not have to rewire the building.

There are actually two types of twisted pair cabling under consideration. The properties of Category 3 cabling will probably limit the speed to about 52 Mbps. Going to higher speeds on this kind of cable is probably not going to be practical, at least in an economic sense because of the amount of processing you have to do to make it work. Category 5, which is a little better quality cable, will be able to run at 155 Mbps.

Specifications are in progress for both Category 3 and Category 5. The schedules have them being completed soon.

Frame Format	Bit Rate (Mbit/s)	Transmission Media
DS1	1.544	Twisted pair
DS3	44.736	Coax pair
STS-3c, STM-1	155.520	SMF
E1	2.048	Twisted pair
E3	34.368	Coax pair
J2	6.312	Coax pair
N x T1 *	N x 1.544	Twisted pair
N x E1 *	N x 2.048	Twisted pair

Table 1. Physical Layer UNI Interface (ATM Forum).

3.3 ATM layer

ATM layer is the layer above the physical layer. As shown in the Figure 3-2, it does the 4 functions which can be explained as follows. Cell header generation/extraction: This function adds the appropriate ATM cell header (except for the HEC value) to the received cell information field from the AAL in the transmit direction.

VPI/VCI values are obtained by translation from the SAP identifier. It does opposite, i.e. removes cell header in the receive direction. Only cell information field is passed to the AAL. Cell multiplex and demultiplex: This function multiplexes cells from individual VPs and VCs into one resulting cell stream in the transmit direction. It divides the arriving cell stream into individual cell flows based on the cell header VCI or VPI in the receive direction.

VPI and VCI translation: This function is performed at the ATM switching and/or cross-connect nodes. At the VP switch, the value of the VPI field of each incoming cell is translated into a new VPI value of the outgoing cell. The values of VPI and VCI are translated into new values at a VC switch.

Generic Flow Control(GFC): This function supports control of the ATM traffic flow in a customer network. This is defined at the B-ISDN User-to-network interface (UNI).

3.3.1 ATM UNI Cell

In section 3.2, we considered moving cells across the physical link with ATM's physical layer. Now let's look explicitly at the ATM cells and how these cells get forwarded to the proper destination?

Looking at the detail of the cell, there are five bytes containing a number of fields, as illustrated in Figure 3-10.

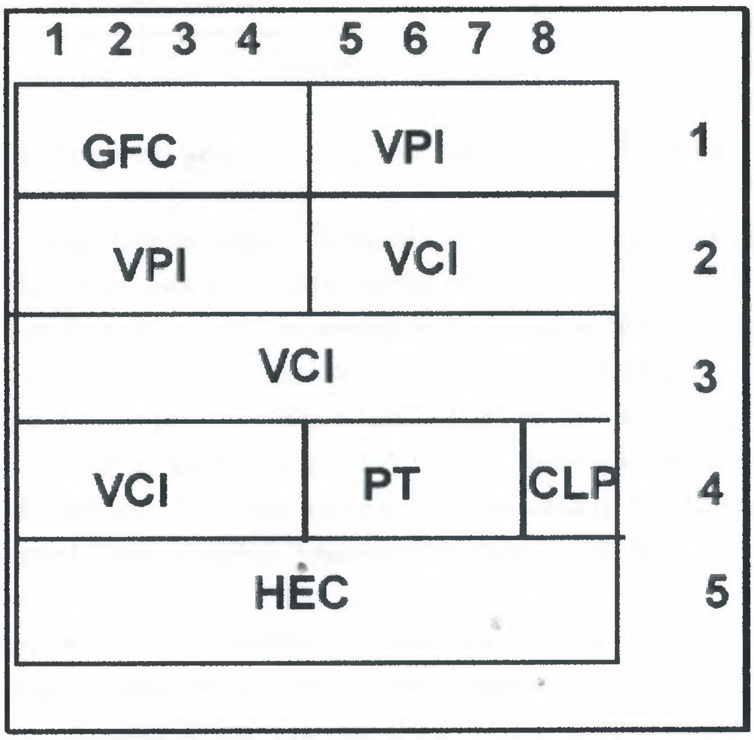

Figure 3-10. The ATM cell header format at the UNI.

The PT is for payload type, CLP for cell loss priority, and HEC for header error check.

3.3.2 Why 53 Bytes?

3.3.2.1 Packetisation Delay Advantage of Small Cells

The first question to ask about the cell size is, "Why 53 bytes?" Not only is it an unusual number, but it's also not very large. Let's look at a few of the trade-offs involved here (see figure Figure 3 - 11).

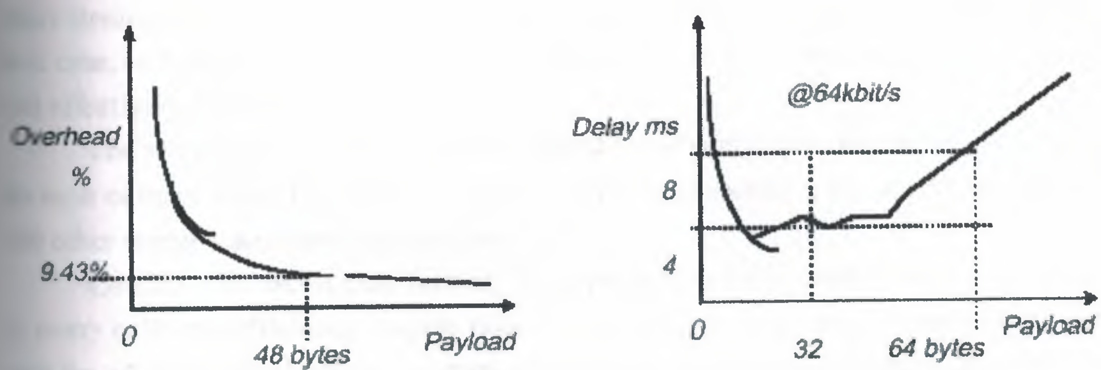

Figure 3-11. Trade-off between delay and cell payload efficiency.

One is packetisation delay. Think about standard digital (PCM) voice, where each individual conversation uses a constant stream off bits at 64 Kbit/s. The voice is encoded by taking 8,000 samples per second. The sample is 8 bits and reflects the energy level at that point. 8,000 eight-bit samples per second results in a data rate of 64 Kbit/s.

Now consider the task of filling a cell. If that cell has 40 bytes of payload, the first voice sample will sit around in the partially-filled cell for 40 sample times and then the cell will be launched into the network. That first voice sample is therefore 5 milliseconds old before the cell even gets launched. This is called "packetisation delay," and it's very important for real-time traffic like voice.

For example, consider satellite communications where the delay is on the order of 250 milliseconds in each direction. You've probably experienced the problems when you talk with somebody on a satellite. When you get these higher delays, it can interfere with normal onversational interactions.

Even lower delays on the order of 10 to 100 milliseconds may cause problems because of echo in a voice network when you do an analog-to-digital conversion. Since we want to keep delay to a minimum, small cells are desirable.

However, there must be some overhead on the cell so the cell can be forwarded to the right place. Using a five byte header, the percentage of the bandwidth that is used by overhead can be very high. Obviously, you can't make the cell too small because you start losing efficiency. The key is to try to balance the delay characteristics with the efficiency. A five byte header with a 48 byte payload results in about 10% overhead.

3.3.2.2 Queueing Advantage of Small Cells

Delay is important, but it also turns out that delay variation is important. Delay variation is the amount of delay difference that cells will experience as they traverse the network.

For example, consider a DS3 link with a 100-byte message to be transmitted. We're not considering the entire network -- just one link. Further assume that this link is shared with 100 other streams of data. What's the worst case and best case delay for the 100-byte message? The best case, of course, is that there's no other data to send when the message arrives. Just send it, and effectively the delay's almost zero.

The worst case would be when everybody else is also sending, in which case you'd wait for each of these 100 other streams to send a cell. In this scenario, you send one cell, you wait for 100 other streams, and then you send your cell.

Consider this worst case further. If the payload is very small, you end up having to send so many cells that efficiency is quite poor. If the cells get large, the amount of time you have to wait for all these other cells to go, goes up. This is somewhat like trying to determine how long you have to wait at a railroad crossing for the train. It depends on how long the train is. And of course, it would be almost no time if the cars were separated.

As you make the cells bigger and bigger, the time you have to wait for a cell before you can get access to the link goes up, and it essentially goes up linearly. We can see some variation at small cell sizes, though, reflecting the fact that 100 bytes, for example, will just fit into a 100-byte payload.

3.3.2.3 Compromise solution between North America and Europe

Naturally, the important question of picking a cell size involved extensive analysis of technical concerns. In Europe, in fact, one of the major concerns was the packetisation delay because the European telephone networks typically are not very large, Thus, they do not have to have a lot of echo cancellation technology deployed. In North America, that was not an issue. Coast-to-coast traffic has already caused telephone companies to deploy echo cancellation technology. Thus, we ended up with these two views. North Americans generally favoured a cell with 64 octets of payload and a five-octet header, while Europeans generally favoured 32 octets of payload and a four-octet header. One of the big differences was the concern about how to handle voice.

This actually turned out to be a political decision by the State Department. CCITT, now ITU-T, is a treaty organisation. The U.S. delegation is run by the State Department. So this compromise of 48 octets for the payload was the result intensive averaging between 64 and 32.

Of course, this turned out to be too long to avoid the use of echo cancellers while failing to preserve the efficiency of 64 octets.

3.3.3 Virtual Connections

Once the cell size is fixed at 53 bytes, the next issue is how to get the cells from place to place. The important fields for this in the header are the VPI/VCI fields, as shown in this example of a number of virtual connections through an ATM switch. Within the switch, there has to be a connection table (or routing table) and that connection table associates a VPI/VCI and port number with another port number and another VPI/VCI.

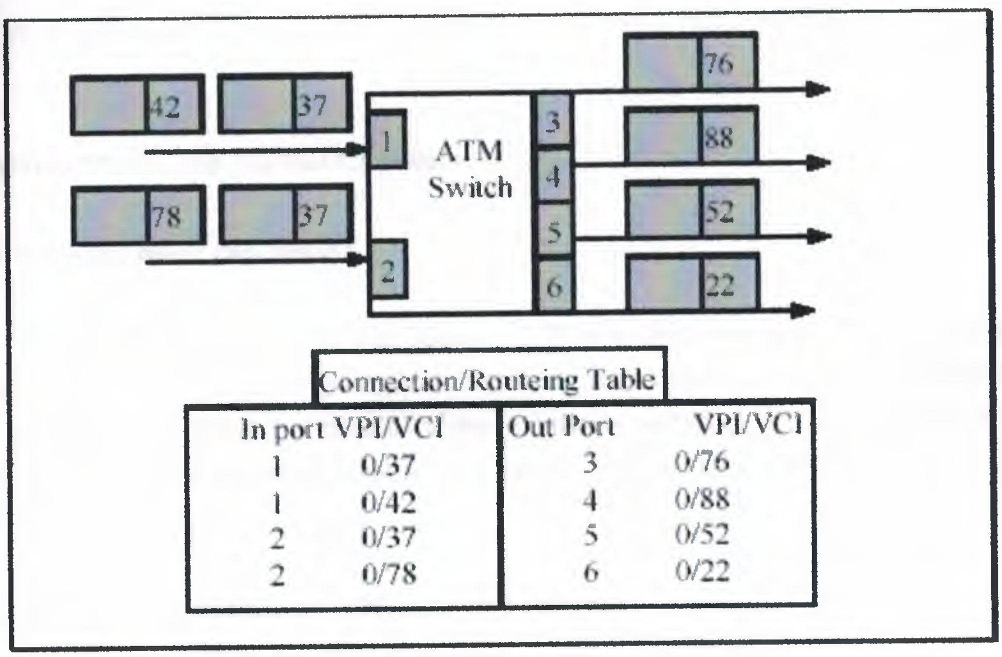

Figure 3-12. Connection/Routeing Table in ATM Switch.

When a cell comes into the switch, the switch looks up the value of the VPI/VCI from the header. Assume that the incoming VPI/VCI is 0-37. Because the cell came in on port one, the switch looks in the port one entries and discovers that this cell has to go to port three. And, by the way, when you send it out on port three, change the VPI/VCI value to 0-76. So as this cell goes through the switch, it pops out with a different header on it. Of course, the information content remains the same.

The VPI/VCI values change for two reasons. First, if the values were unique, there would only be about 17 million different values for use. As networks get very large, 17 million onnections will not be enough for an entire network.

Probably more important, though, is consideration of the administration of unique values across a large network. For instance, how would you guarantee that in "somewhere", the new connection you're about to establish was using a value which was unique compared to all the other connections that were existing in the world?

It is interesting to note that both of these considerations are becoming quite important in the Internet, where a limited number of TCP/IP addresses are available. If the address space were made large enough to serve as universal addresses, the overhead in comparison to the payload in the cell would become unacceptable. Consequently, the VPI/VCI value is only meaningful in the context of the given interface. In fact, in this example "37" is used on both interfaces but there's no ambiguity because they're in the context of different physical interfaces. There's a separate entry for 37 for port two, which of course goes to a different destination.

So the combination of the VPI/VCI values allow the network to associate a given cell with a given connection, and therefore it can be routed to the right destination.

3.3.4 Virtual Paths and Virtual Channels

3.3.4.1 Why are there two fields?

The answer is shown here. The idea of a VPI is to think about a VPI as a bundle of virtual channels. The VPI is 8 bits, so on a given interface, you could have up to 256 different bundles. In the example here. there a couple of virtual channels within each virtual path. Of course, the individual virtual channels have unique VCI values, but the VCI values may be reused in each virtual path.

3.3.4.2 Why are there separate VPI and VCI values?

One application is shown in the Figure 3-12. There are two different ways of getting connections to an ATM network. In the previous section, virtual channels were mapped. The VPI-0/VCI-37 was mapped to a given output with VPI-0/VCI-76 with the same virtual path number.

Suppose this is a carrier network and there are two locations to be connected with ATM. Obviously, it would be nice if you could buy a "bundle" of connections. Then you could establish and tear down virtual channels between your locations without telling the carrier and, more importantly, without having to go through service order processing.

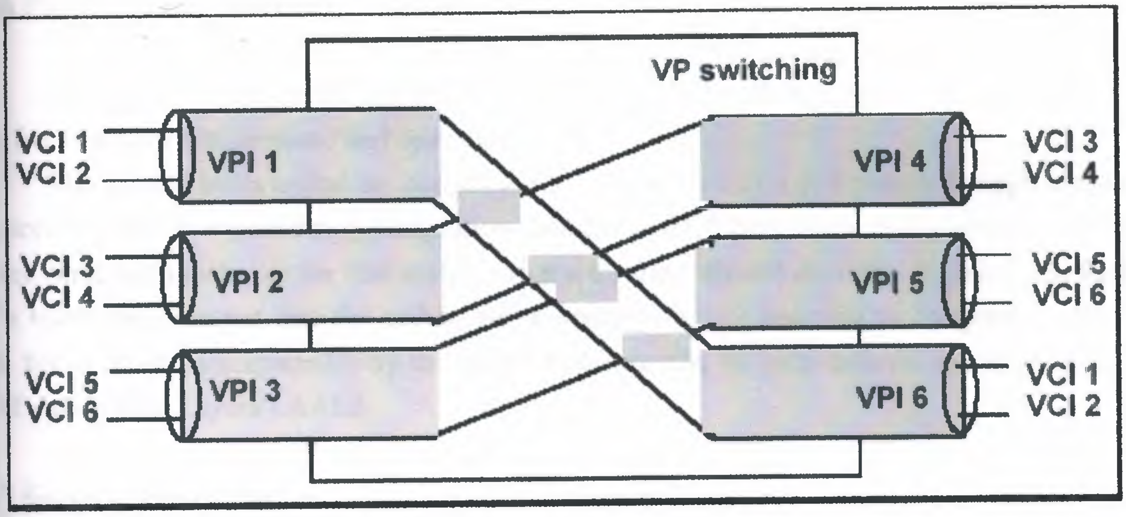

Figure 3-13. An Example of VP Switches.

This is the exact idea of a virtual path. In this example, virtual path VPI=3 is on the left side of the figure on this side of the network with a couple of virtual channels in it. These virtual channels are forwarded as a bundle and appear at the destination with the same VCI values. Since the network does not change (or care about) the VCIs, it only looks at the VPI field. The VPI field could (and in this example does) change, but this whole bundle exists as an entity and is carried through the network as a whole. If a third channel were to be added, only the equipment at each end must be co-ordinated. There is no need to tell the network provider. This is type of service is called Virtual Path Service.

3.3.5 Cell Loss Priority (CLP)

Perhaps no other single bit in the history of telecommunications has had so much function or been so important as the cell loss priority bit. In many ways, this may be the most important bit in the cell header.

The purpose of this bit is that cells with this bit set should be discarded before cells that do not have the bit set. One way to think about it is it's the way an expendable cell identifies itself as a "trash me first" cell.

Consider reasons that cells may be marked as expendable. First, this may be set this by the terminal. This may be desirable if, for example, you were using a wide area network service and you got a price break for these low-priority cells. This could also be used to set a kind of priority for different types of traffic when you were aware that you were overusing a committed service level. More likely, though, the more important use is that this bit will be set by the ATM network. In particular, this bit will reflect back into something called traffic management in the traffic contract.

3.3.6 Payload Type Identifier

The payload type identifier has three bits in it. The first bit is used to discriminate cells of data from cells of maintenance and operation.

The second bit is called the congestion experience bit. If a cell passes through a point in the network that is experiencing congestion, this bit is set. There is no really precise definition yet on what that means, or for that matter, what a terminal should do when it sees it set. Rather, it's a hook that was put into the cell to allow possibilities for reacting to congestion. And the third bit is carried transparently by the network. Currently, its only defined use is in one ofthe ATM Adaptation Layers - AAL5.

3.3.7 Generic Flow Control

The generic flow control (GFC) field occupies the first four bits in the header. It is only defined on the UNI. It is not used in the Network Node Interface (NNI), which is the interface between switches.

The GFC field is currently undefined, so it's set to zero and reserved. Potential uses are for flow control or for building a multiple access, a shared medium ATM network on the local access facilities.

3.3.8 Header Error Check

The last eight bits in the header are the header error check (HEC). HEC is needed because if a cell is going through a network and the VPI/VCI values get errors, it will get delivered to the wrong place. As a security issue, it was deemed useful to put some error checking on the header. Of course, the HEC also is used, depending on the physical medium, e.g. in SONET, to delineate the cell boundaries.

HEC actually has two modes. One is a detection mode where if there is an error with the CRC calculation, the cell is discarded. The other mode allows the correction of one-bit errors. Whether one or the other mode is used depends on the actual medium in use. If fibre optics is used, the onebit error correction may make a lot of sense because typically the errors are isolated. It may not be the right thing to do in a copper medium because errors tend to come in bursts. When the one-bit error correction is used, you increase the risk of a multiple-bit error being interpreted as a single-bit error, mistakenly "corrected," and sent someplace. So the error detection capabilities drop when the correction mode is used.

Notice that the HEC is recalculated link by link because it covers the VPI- VCI value, and the VPI/VCI value changes as you go through the network.

3.4 ATM Adaptation Layer (AAL)

3.4.1 Two Sublayers

AAL is divided into two sub-layers as shown in Figure 3-2: Segmentation and reassembly (SAR), and Convergence sublayer(CS).

SAR sublayer: This layer performs segmentation of the higher layer information into a size suitable for the payload of the ATM cells of a virtual connection and at the receive side, it reassembles the contents of the cells of a virtual connection into data units to be delivered to the higher layers.

CS sublayer: This layer performs functions like message identification and time/clock recovery. This layer is further divided into Common Part Convergence Sublayer (CPCS) and a Service Specific Convergence Sublayer (SSCS) to support data transport over ATM. AAL service data units are transported from one AAL Service Access Point (SAP) to one or more others through the ATM network. The AAL users can select a given AAL-SAP associated with the QOS required to transport the AAL-SDU. There are 5 AALs have been defined, one for each class of service.

3.4.2 Broadband Services and Applications

There are several practical applications using **ATM Technology**. ATM is going to be the Backbone Network for many broadband applications including Information SuperHighway. Some of the key applications include: video conferencing, desktop conferencing, multimedia communications, ATM over satellite communications, mobile computing over ATM for wire-less networks.

The ITU-T has classified broadband services into the following categories:

- Interactive services:
- Conversational services,
- Message services, and
- Retrieval services.
- Distribution services:
- Distribution services with user control, and
- Distribution services without user control.

All these services will be transported by ATM cells from sources to destinations. The role of the ATM Adaptation Layers (AALs) is to define how to put the into the ATM cell payload. The services and applications are different and therefore require different types of AAL. It is important to know what kinds of services are required.

	Class A	Class B	Class C	Class D
Timing relation	required		not required	
Bitrate	constant	variable		
Connection mode	connection-oriented			connection-less

Examples: A - Circuit emulation, CBR Video
B - VBR video and audio
C - CO data transfer
D - CL data transfer

Figure 3-14. Service Classes and Their Attributes.

Figure 3-14 illustrates the results of the ITU-T's efforts for defining service classes. To read the diagram, take a vertical slice under each of the letters.

Class A has these attributes:

- End-to-end timing is required.
- Constant bit rate.
- Connection oriented.

Thus, Class A is emulating a circuit connection on top of ATM. This is very important for initial multimedia applications because virtually all methods and technologies today for carrying video and voice assume a circuit network connection. Taking this technology and moving it into ATM requires supporting circuit emulation service (CES).

Class B is similar except that it has a variable bit rate. This might be doing video encoding but not playing at a constant bit rate. The variable bit rate really takes advantage of the bursty nature of the original traffic.

Class C and ***D*** have no end-to-end timing and have variable bit rates. They really are oriented toward data communications, and the only difference between the two is connection-oriented versus connection-less.

3.4.3 AAL type 1 for Class A

Figure 3-15 shows AAL1 for Class A, illustrating the use of the 48-byte payload. This illustrates that one of the bytes of the payload must be used for this protocol.

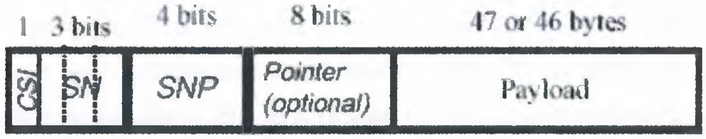

Figure 3-15. AAL 1 Packet Format for Class A.

SN - Sequence Number (Lost cell detect: used by Adaptive Clock Method),
CSI - Convergence sublayer indication (0=no pointer, 1=pointer),
SNP - Sequence Number Protection by CRC

There are a number of functions here, including detecting lost cells and providing time stamps to support a common clock between the two end systems. It is also possible that this header could be used to identify byte boundaries. For example, if this were emulating a DS-1 connection, one could identify the subchannels (the DS-0s) within that stream.

3.4.3.1 AAL1: Adaptive Clock Method

The primary objective is to obtain clock agreement, making sure that we can play out the original information stream.

Let's use as an example a 64 Kbits/s voice. The transmitter is collecting voice samples, filling up cells and throwing those cells into the network at about once every 5.875 msec (transmit 47 octets at a speed of one octets every 125 microseconds). The receiver is shown in Figure 3-16. The receiver's job is to play out the original bit stream at 64 kilobits per second. This is where we see the impact of variation and delay.

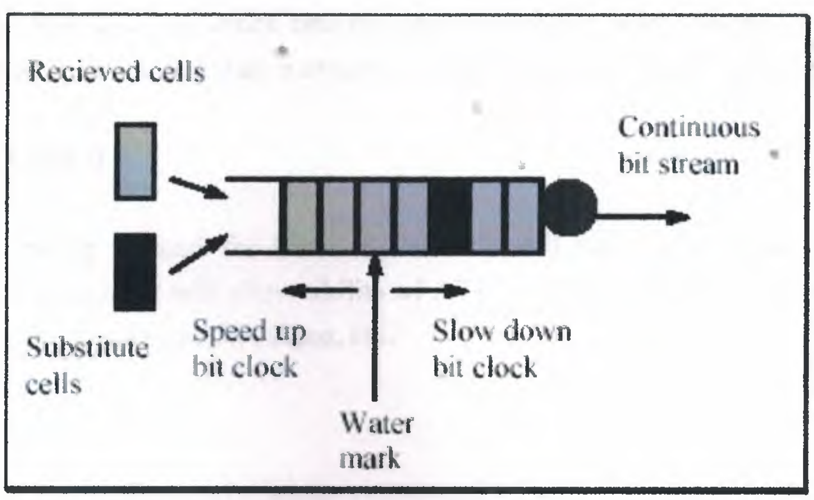

Figure 3-16. An illustration of Adaptive Clock Method.

Using the adaptive clock method, the receiver establishes a buffer based on the characteristics of the connection at 64 kilobits. It establishes a water mark (we'll see why) and then collects some cells up to about the water mark. Then the receiver can start unwrapping the bits in the payload and playing them out as a stream of bits at 64 kilobits per second.

Consider a couple of issues now. First, how do we know that the clock that's being used to play out the bit stream is in agreement with the transmitter's clock? If it's too fast, eventually we'll see the buffer empty because the cells will be arriving a little bit too slow compared to rate of emptying them. Thus, we'll have a buffer starvation problem. If the clock rate is a little bit too slow, the buffer will start to fill, and eventually it will overrun the buffer. Then we start losing cells. The solution is that the receiver observes the fill of the buffer relative to the water mark. If it starts to get empty, it slows the (output) clock down because the clock's going a little fast. If it starts to get too full, it speeds the (output) clock up. This way, the receiver's output clock rate stays centred around the transmitter's clock.

While we need this buffer to do this adaptation, the question of how big the buffer should be remains. It must be a function of how variable the arrival rate is for the cells. If the cells behave like the typical rental car company bus at the airport, you stand there for 20 minutes waiting for the bus and then five buses show up all at once. In this case, you obviously have a problem because you need a very large buffer (shelter at the airport) to ride out those long dry spells.

This is what happens if there is a lot of delay variation in the cells as they traverse the network. This translates into making this buffer bigger and increasing the delay from when a cell is first generated to when it finally gets played out. Cell delay variation is a very important factor in the quality of service (QoS), thus it is an important parameter in traffic management.

A second important factor is consideration of the effects of losing a cell. Part of the protocol is a sequence number, which is not meant to maintain sequence of the cells, but is maintained to detect loss. If a cell is lost, the receiver should detect the loss and essentially put in a substitute cell. Otherwise, the clock rate becomes unstable.It is interesting to notice that with this kind of scheme, one can maintain a circuit-like connection of virtually any speed over ATM.

3.4.4 AAL2 for Class B

AAL2 is being defined for Class B, but it's still under development. This will be important though, because it will allow ability of ATM to support the bursty nature of traffic to be exploited for packet voice, packet video, etc.

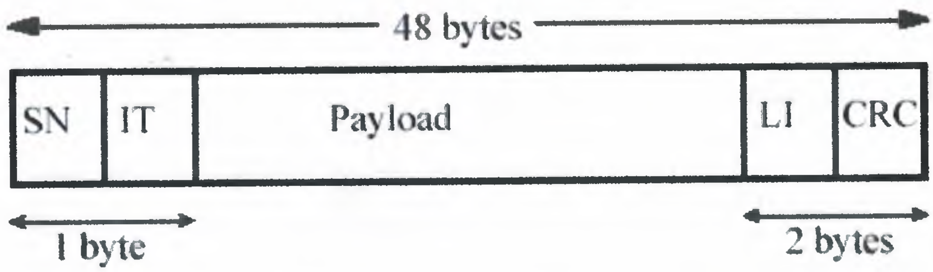

Figure 3-17. AAL 2 PDU.

3.4.5 AAL 3/4 for Classes C & D

In AAL 3/4, the protocol first puts error-checking functions before and after the original data. Then the information is chopped into 44-byte chunks. The cell payloads include two bytes of header and two bytes of trailer, so this whole construct is exactly 48 bytes.

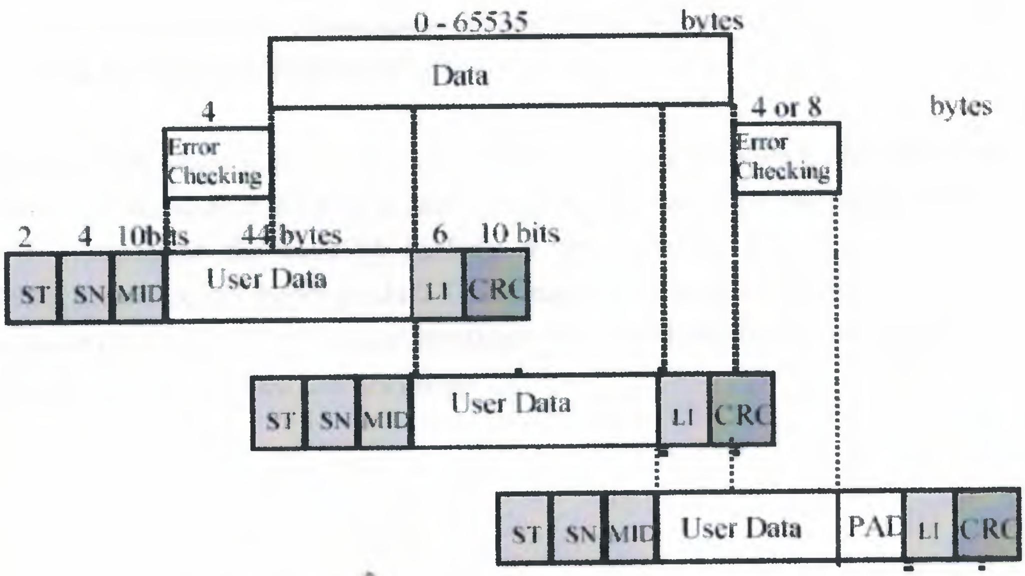

Figure 3-18. AAL 3/4 Packet Format for Class C & D.

Notice that there is a CRC check on each cell to check for bit errors. There is also an MID (Message ID). The MID allows multiplexing and interleaving large packets on a single virtual channel. This is useful in a context where the cost of a connection was veryexpensive since it would help to guarantee high utilisation of that connection.

3.4.6 AAL5 for Classes C & D

The other data-oriented adaptation layer is AAL5. Here, the CRC is appended to the end and the padding is such that this whole construct is exactly an integral number of 48-byte chunks.

This fits exactly into an integral number of cells, so the construct is broken up into 48-byte chunks and put into cells.

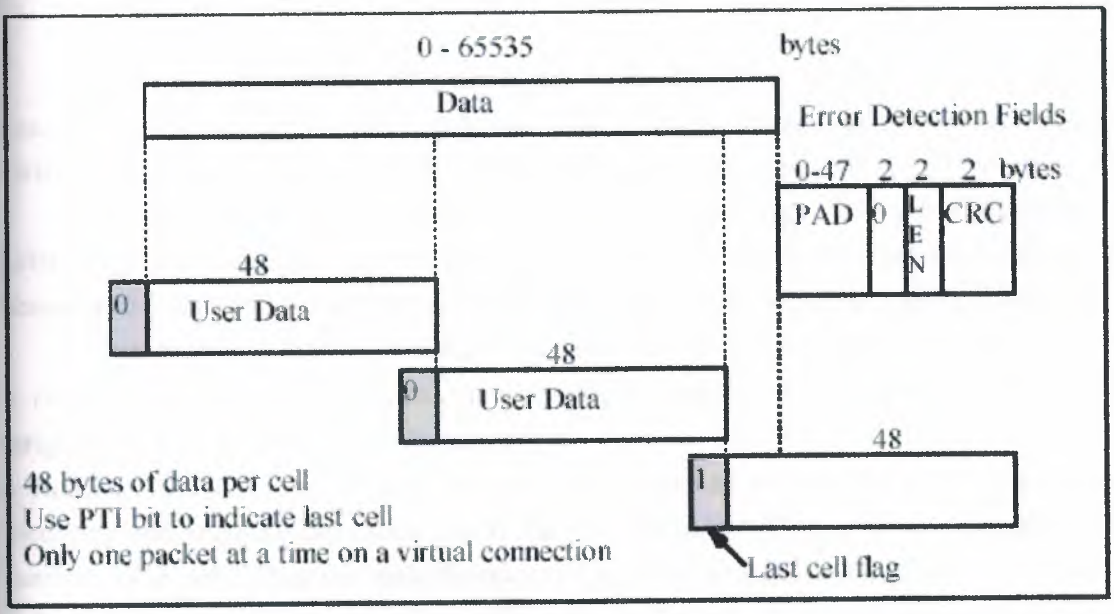

Figure 3-19. AAL 5 Packet Format for Class C & D.

To determine when to reassemble and when to stop reassembling, remember the spare bit for PT that was in the header. This bit is zero except for the last cell in the packet (when it is one). A receiver reassembles the cells by looking at the VPI-VCI and, for a given VPI-VCI, reassembling them into the larger packet. This means that a single VPI-VCI may support only one large packet at a time. Multiple conversations may not be interleaved on a given connection. This is attractive where connections are cheap.

CHAPTER FOUR : ATM SWITCHES

Having considered the physical layer, ATM layer, AAL and the header of the cell, let's consider how cells are actually switched. This is by no means an exhaustive dissertation on switching, but it gives an overview of the nomenclature involved with ATM switches.

ATM Switching is also known as fast packet switching. ATM switching node transports cells from the incoming links to outgoing links using the routing information contained in the cell header and information stored at each switching node using connection set-up procedure.

Two functions at each switching node are performed by a connection set up procedure.

- A unique connection identifier at the incoming link and a unique connection identifier at the outgoing link are defined for each connection.
- Routing tables at each switching node are set up to provide an association between the incoming and outgoing links for each connection. VPI and VCI are the two connection identifiers used in ATM cells. Thus the basic functions of an ATM switch can be stated as follows.
- Routing (space switching) which indicates how the information is internally routed from the inlet to outlet.
- Queueing which is used in solving contention problems if 2 or more logical channels contend for the same output.
- Header translation that all cells which have a header equal to some value j on incoming link are
switched to outlet and their header is translated to a value k.

4.1 Queueing Disciplines in an ATM Switching Element:

There are mainly three different buffering strategies available determined by their physical location as follows.

Input queueing: In this, the contention problem is solved at the input buffer of the inlet of the switching element. Each inlet contains a dedicated buffer which is used to store the incoming cells until the arbitration logic decides to serve the buffer. The switching transfer medium then switches the ATM cells from the input queues to the outlet avoiding an internal contention. The arbitration logic can be as simple as round-robin or can be complex such as taking into account the input buffer filling levels. However, this scheme has Head of Line (HOL) blocking problem i.e. if two cells of two different inlets contend for the same output, one of the cells is to be stopped and this cell blocks the other cells in the same inlet which are destined for different outlet. This queueing discipline can be shown by Figure 4-1.

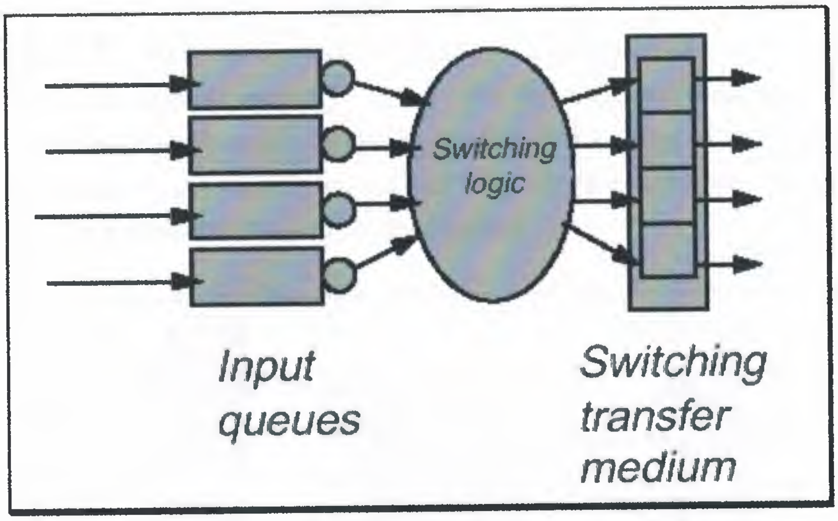

Figure 4-1. Input Queueing Disciplines.

Output Queueing: In this queueing discipline, queues are located at each outlet of the switching element and the output contention problem is solved by these queues. The cells arriving simultaneously at all inlets destined for the same output are queued in the buffer of the outlet. The only restriction is that the system must be able to write N cells in the queues during one cell time to avoid the cell loss where N is the total number of inlets of the switch. In this mechanism, noarbitration logic is required as all the cells can be switched to their respective output queue. The cells in the output queue are served on FIFO basis to maintain the integrity of the cell sequence.

Figure 4-2 illustrates this mechanism.

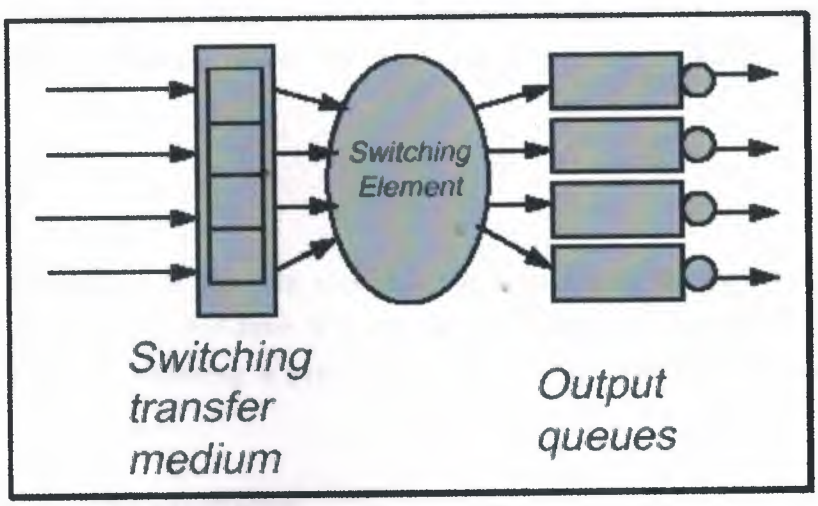

Figure 4-2. Output Queueing Disciplines.

Central Queueing: In this scheme, the queueing buffers are shared between all inlets and outlets. All the incoming cells are stored in the central queue and each outlet chooses the cells which are destined for it from this central memory. Since cells for different outlets are merged in this central queue, FIFO discipline is not followed in reading and writing of this queue. Cells can be written and read at random memory locations and this needs a complex memory management system for thisscheme. Figure 4-3 shows this mechanism.

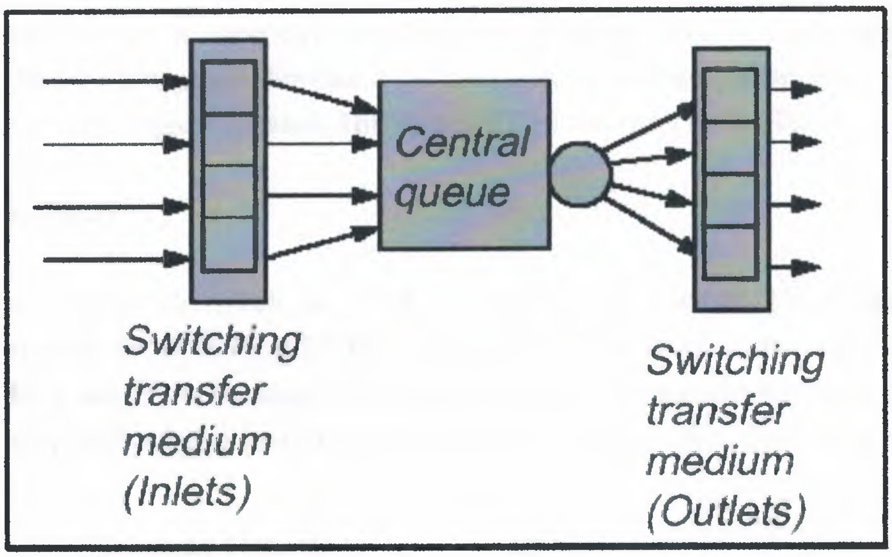

Figure 4-3. Central Queueing Disciplines.

4.2 Performance Issues

There are five parameters that characterise the performance of ATM switching systems: throughput, connection blocking probability, cell loss probability, switching delay, and delay jitter.

4.2.1 Throughput

This can be defined as the rate at which the cells depart the switch measured in the number of cell departures per unit time. It mainly depends on the technology and dimensioning of the ATM switch. By choosing a proper topology of the switch, the throughput can be increased.

4.2.2 Connection Blocking Probability

Since ATM is connection oriented, there will be a logical connection between the logical inlet and outlet during the connection set up phase. Now the connection blocking probability is defined as the probability that there are not enough resources between inlet and outlet of the switch to assure the quality of all existing as well as new connection.

4.2.3 Cell Loss Probability

In ATM switches, when more cells than a queue in the switch can handle will compete for this queue, cells will be lost. This cell loss probability has to be kept within limits to ensure high reliability of the switch. In Internally Non-Blocking switches, cells can only be lost at their inlets/outlets. There is also possibility that ATM cells may be internally misrouted and they reach erroneously on another logical channel. This is called Cell Insertion Probability.

4.2.4 Switching Delay

This is the time to switch an ATM cell through the switch. The typical values of switching delay range between 10 and 1000 microseconds. This delay has two parts:

- Fixed switching delay: it is because of internal cell transfer through the hardware.
- Queueing delay: this is because of the cells queued up in the buffer of the switch to avoid the cell loss.

4.2.5 Jitter on the Delay

This is denoted as the probability that the delay of the switch will exceed a certain value. This is called a quantile and for example, a jitter of 100 microseconds at a 10exp-9 quantile means the probability that the delay in the switch is larger than 100 Micro secs. is smaller than 10exp-9.

4.3 ATM Switching Technology

A generic switch consists of switching fabric, input processors and output processors. The input processor examines the header and identifies where the output port is for this cell. The switch uses the connection table discussed earlier; with one addition -- a tag which may or may not be present. Notice that the input processor, the fabric and the output processors can buffer information. When you hear about input buffered switches, output buffered switches, or even fabric buffered switches, this refers to where the main buffering is for the cells. These are implementation of the queueing disciplines, discussed in section 4.1, using different technologies.

4.3.1 Shared Backplane

Smaller ATM switches usually targeted at the local area often use a shared backplane. In Figure 4- 4, there are input processors, which are doing translations from VPI/VCI to tag. Then these processors use some form of bus arbitration because the bus bandwidth is shared among the input processors. When one of these processors gets access to the bandwidth, it sends the cell and all of the output processors can observe that cell.

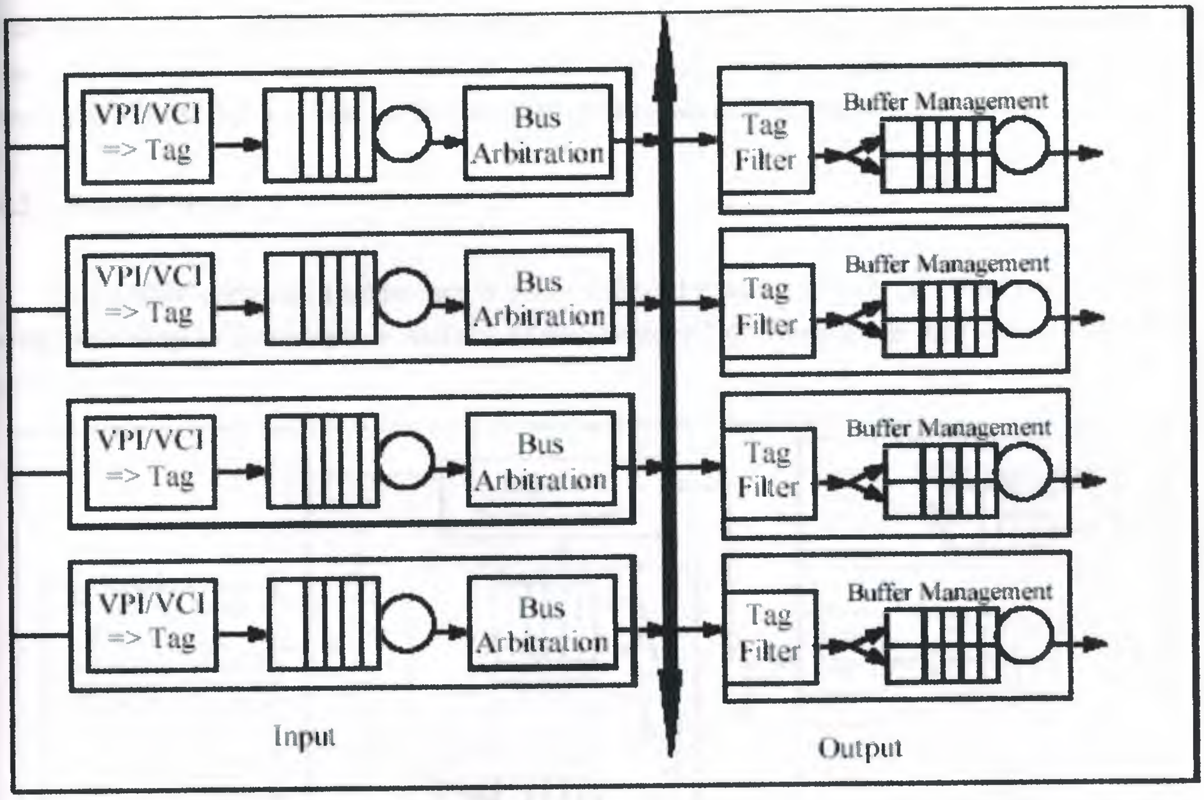

Figure 4-4. Shared Backplane Switch.

The output processors look at the tag, and that tag identifies whether the cell is meant for this particular output processor. If the cell is indeed intended for a given output processor, the processor copies the information from the bus -- that's what the tag filter is all about -- and buffers the cell to be sent out on the link. Notice that this has a natural multi-cast capability because every output processor sees all the cells.

In fact, it is possible to define a tag that identifies multiple output processors so cells can start to fan out through this switch. In ATM point-to-multipoint connections, one cell comes in, and multiple copies go out into the ATM network. This architecture facilitates building these point-to-multipoint connections. This might be used for video distribution, for example.

Buffer management on the output is important. Notice there are two output buffers shown in Figure 4-4. This is important because it's possible to have cells coming in from multiple inputs destined for the same output. Buffers are needed so some of the cells may be temporarily stored while there is a focused load on a particular output link.

Why two buffers? Primarily because there are really two kinds of quality of service. For voice, for example, you want the delay to be short. In fact, you're probably better off throwing a voice cell away rather than delivering it late. You might want to put this kind of real-time traffic cells off in one buffer, which tends to be pretty small.

Data has many opposite properties. You don't really want to lose cells for data, so you probably have a separate buffer. This buffer will be relatively large. The delay is increased, but the cell loss rate is decreased. While there is extensive discussion about the "best" ATM switching fabric, the way that buffers are managed is just as important.

4.3.2 Shared Memory

Another common architecture is called "shared memory." This is a dual-ported memory with the ability to move data in and out of the memory "N" times faster than the rate of arrivals on these ports.

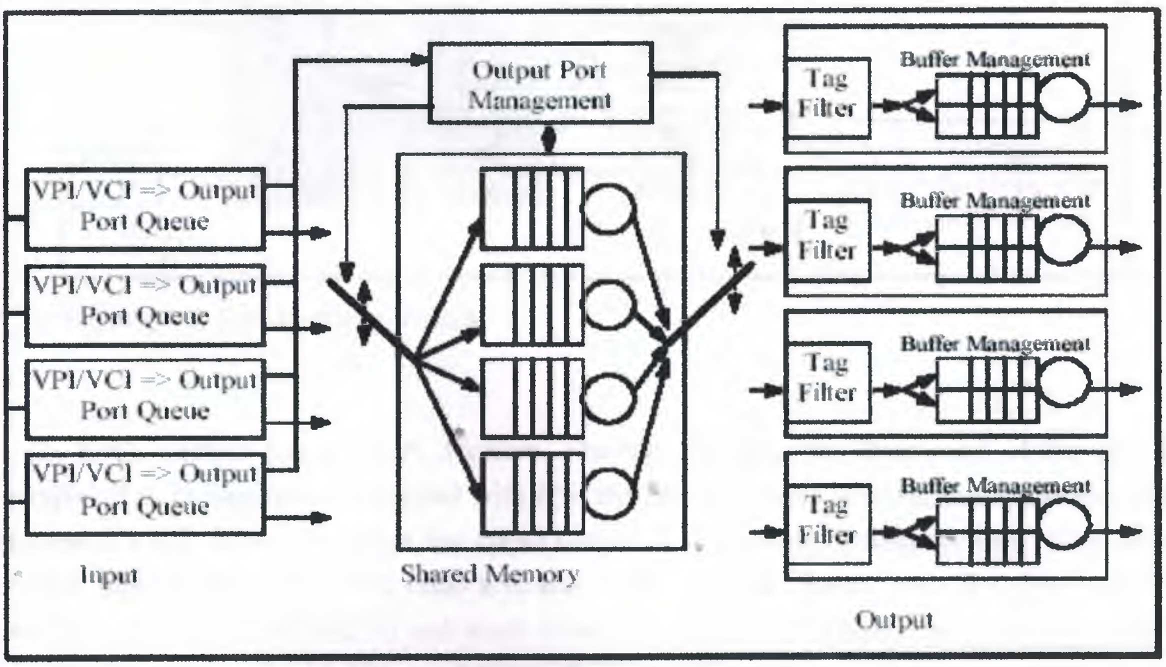

Figure 4-5. Shared Memory Switch.

Assuming the four ports shown in Figure 4-5 as an example, the rate at which cells can move into the memory must be four times faster than the individual cell rates on the ports because it must keep up with cell arrivals on all the ports simultaneously. In this example, the input processors are receiving cells, and the processors are communicating with a queue manager that controls both how the cells are placed into buffers in memory and how they're pulled out of buffers in memory to the right output processor.

In this case, multi-casting requires that the cells stay in memory be transferred multiple times to multiple output ports. Thus, in this case, switching fabric performs most or all of the buffering.

4.3.3 Self Routing Fabric

A third switching architecture is a self-routing fabric. This technique involves building switch elements of multiple little elements.

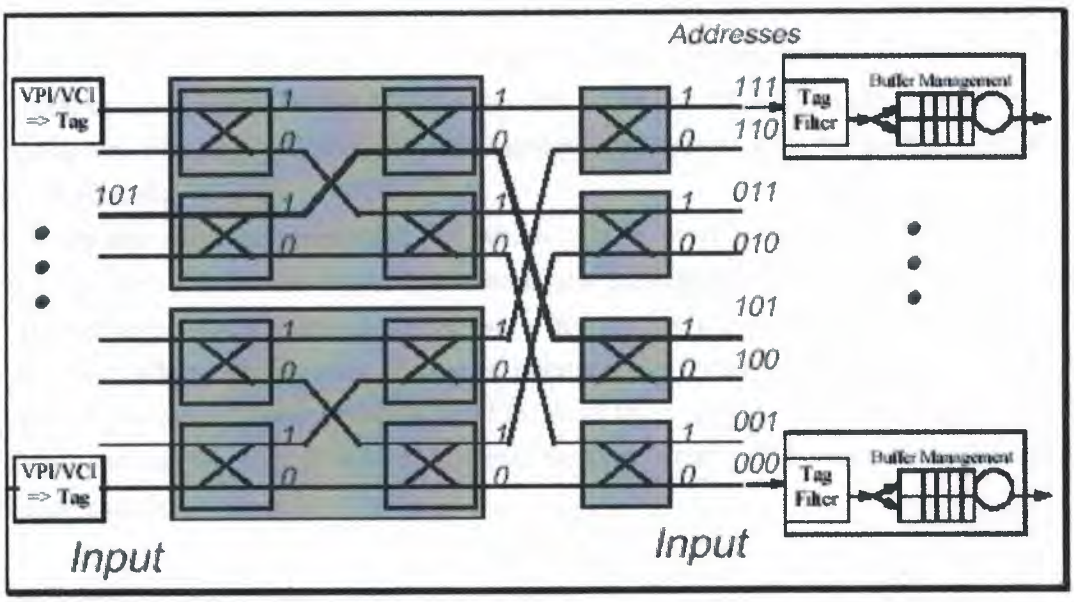

Figure 4-6. Self-Routing Switch.

Each element has a binary decision. The tag identifies the destination of the cell. For example, if an incoming cell is tagged with 101, the first bit would be examined in the first stage. Because it's a 1, it would go out the upper output to the next switching element. This second element looks at the second bit, finds a 0, and sends it out the lower port. The third switching element finds a 1 on the third bit and sends it out the upper port. Consequently, regardless of the input port, cells tagged with 101 come out on the correct destination port. Consequently, the tags identify precisely for which output port a particular cell is destined.

One must address, though, what happens if two cells simultaneously are destined for the same output port? The cells will collide in one of the elements. Typically there's a scheme to examine the cells ahead of time and hold them in the input and only let one go through, but there are many various designs.

One of the interesting aspects of this architecture is that if you start stacking up the switching elements in rows and columns, you get more and more ports. Thus, you get larger and larger switches. Occasionally you'll see somebody referring to this as an infinitely scaleable switching technology.

4.4 Virtual Connections

4.4.1 Permanent Virtual Connections

The connections just discussed involve routing through a switch only. How do we get a connection established through a network?

One technique is called a permanent virtual connection (PVC). This will be done through some form of service order process. Conceptually, there is some sort of network management system that communicates to the various devices what the VCI-VPI values are and what the translations are. For example, the network management system tells the switch what entries to make in its connection table.

There are some environments for which this is most reasonable. If there are a small number of devices attached to the ATM network, and these devices tend not to move around very much, this behaves much as telephone network private lines. This tends to make a lot of sense when you have a large community of interest between two locations. Because it takes a while to set up these connections, you typically want to leave them set up. You don't want to try to tear them down and set them up in a very dynamic fashion. That's why these are called permanent virtual connections.

4.4.2 Switched Virtual Connections

A second technique for establishing a connection through a network is called a switched virtual connection (SVC). This allows a terminal to set up and tear down connections dynamically.

The way SVCs operate is one of the VPI/VCI values predefined to be used for the signalling protocol that controls connections. The value is VPI-0/VCI-5, and this connection is terminated by the call processing function. Of course, the "receiving" terminal also has VPI-0/VCI-5 terminating at the call processing function for this (or another) switch. A protocol called the "signalling protocol" is used on the VPI-0/VCI-5 connection to communicate with the switch, passing information to allow the connection to be set up or to be torn down (or to even be modified while it's in existence.)

The result is dynamic connection configuration. Further, these connections will probably be established in less than a second. Note that the connection that is set up for actual information transfer is NOT VPI-0/VCI-5. The other connection passing around the call processing function, does not interact with the call processing functions within the switch.

4.4.3 An Example of Call Set Up

Let's look at the SVC call set-up procedure in more detail. Assume that "A" wants to talk to "B".

4.4.3.1 Step 1

First, a message is sent on the signalling channel called the set-up message. This message has a lot of information like address of the party being called, the calling party address, the traffic characteristics, and the quality of service being requested. Also, since it's possible to have many, many connections in various stages of processing at a given time all being serviced over this same link and using VPI-0/VCI-5, there must be some identifier that uniquely identifies the call in question.

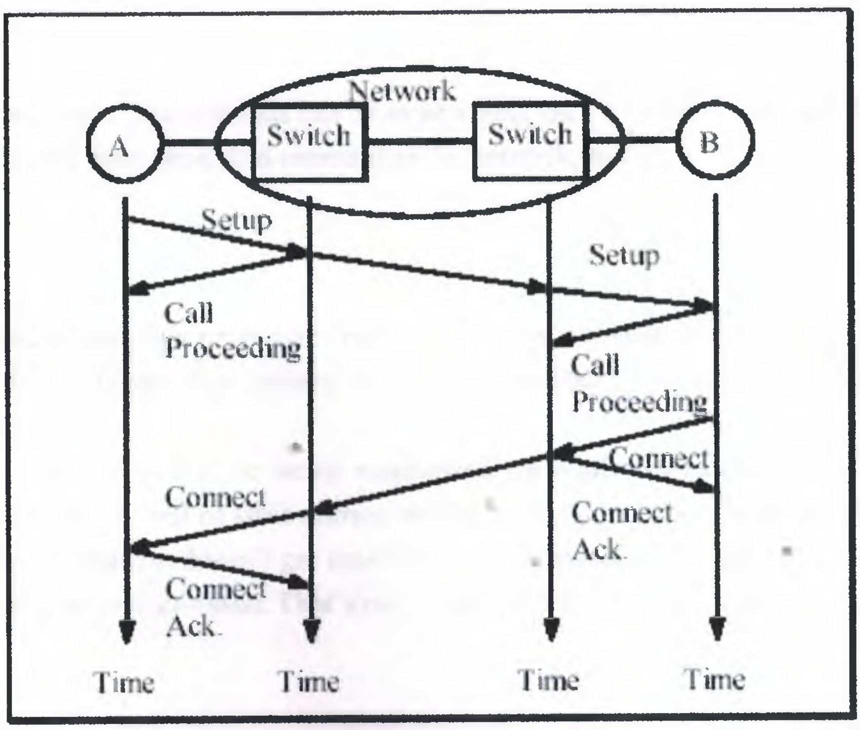

Figure 4-7. An example of Call Setup.

At this point. the first switch simply acknowledges that connection and assigns a VPI/VCI value for the connection. The connection doesn't really exist yet -- it can't be used for communications but the terminal does know the VPI/VCI value assigned.

4.4.3.2 Step 2

Then some magic happens. The network checks its resources, looks for paths, etc., to try to find out if there's a way to get this connection set up to the destination.

4.4.3.3 Step 3

At this point, the ATM equivalent of ringing occurs. A message comes into the destination terminal indicating an incoming call, including information about the traffic characteristics, etc. The called terminal acknowledges because, in a general case, this could be a public network. Rather than this just being a simple terminal at the destination, this could be a fairly large, complex private ATM network. Consequently, the called terminal is going to have to do some checking of its own resources and figure out if it can reach the ultimate destination. So this is an "Okay, I'm working on it" message.

4.4.3.4 Step 4

Assuming that the connection can be established, then a connect message that says "okay, I accept" is returned from the called terminal to the network, and...

4.4.3.5 Step 5

The connect message propagates back through the network to the calling terminal with a message that says, "Okay, that connection's been accepted. You can now start using it for communicating."

These connections that are being established are really the establishment of connection tables in the switches. A trail of table entries are left as this connection is set up. Of course, if for some reason the connection doesn't get established, the network must go back and clean up all of this context that's been established. That's one reason that this is fairly challenging technically.

CHAPTER FIVE : ATM INTERFACES AND NETWORKS

5.1 User-Network Interface (UNI) & Protocols

Two elements can be used to describe a reference configuration of the User-Network access of BISDN: Functional groups and Reference points. The following figure gives the reference configuration.

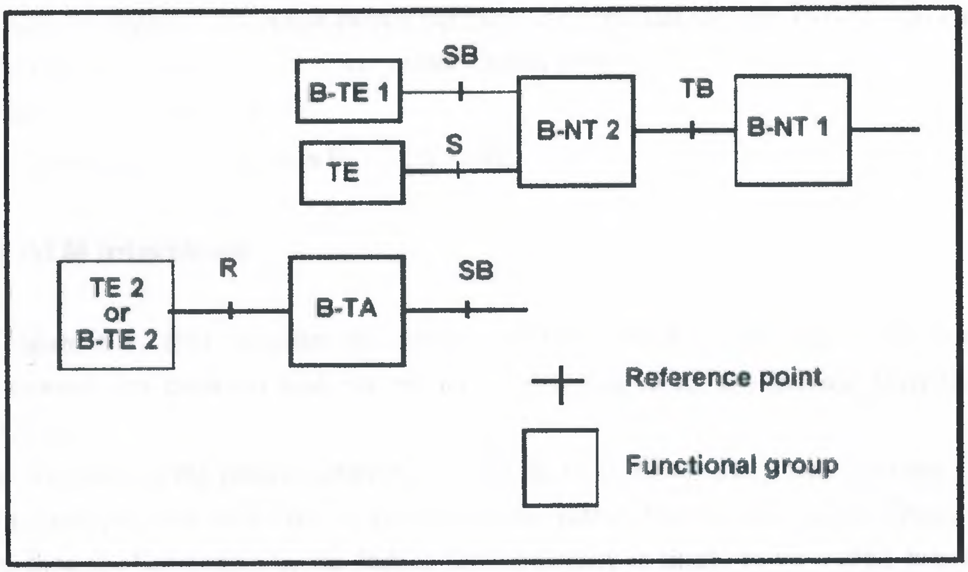

Figure 5-1. Reference Configuration.

The B-NT1 and B-NT2 are Broadband Network Terminators. The B-NT2 provides an interface allowing other type of TEs rather than the broadband TEs to be connected to the broadband network.

B-NT1 functions are similar to Layer 1 of the OSI Reference model and some of the functions are:

- Line Transmission Termination,
- Transmission Interface handling, and
- OAM functions.

B-NT2 functions are similar to layer 1 and higher layers of the OSI model. Some functions of BNT2 are:

- Adaptation functions for different interface media and topology,
- Multiplexing & demultiplexing and concentration of traffic,
- Buffering of ATM cells,

- Resource allocation & Usage parameter control,
- Signalling protocol handling,
- Interface handling,
- Switching of internal connections.

B-TE1 and B-TE2 are Broadband Terminal Equipment. B-TE1 can be connected directly to the network from the reference SB and TB. B-TE2 can only be connected to the network via a broadband adapter. B-TA is Broadband Terminal Adapter. It allows the B-TE2, which can not be connected directly, to be connected to the Broadband network.

SB and TB indicate reference points between the terminal and the B-NT2 and between B-NT2 and B-NT1 respectively. Reference points Characteristics

- TB and SB: 155.520 and 622.080,
- R: allow connection of a TE2 or a B-TE2 terminal.

5.2 Other ATM Interfaces

In Figure 5-2, first consider the private ATM network in the upper left corner. The interface between the terminal and the switch is referred to as the private User-to-Network Interface (UNI).

The interface to the public network is a public UNI. Now, these two interfaces are quite similar. For example, the cell size is the same; the cell format is the same. There are some differences, though. For example, the Public UNI interface is likely to be a DS3 interface early on, but it's very unlikely that one would deploy a DS3 across the campus. Consequently, we'll probably see some differences at the physical layer.

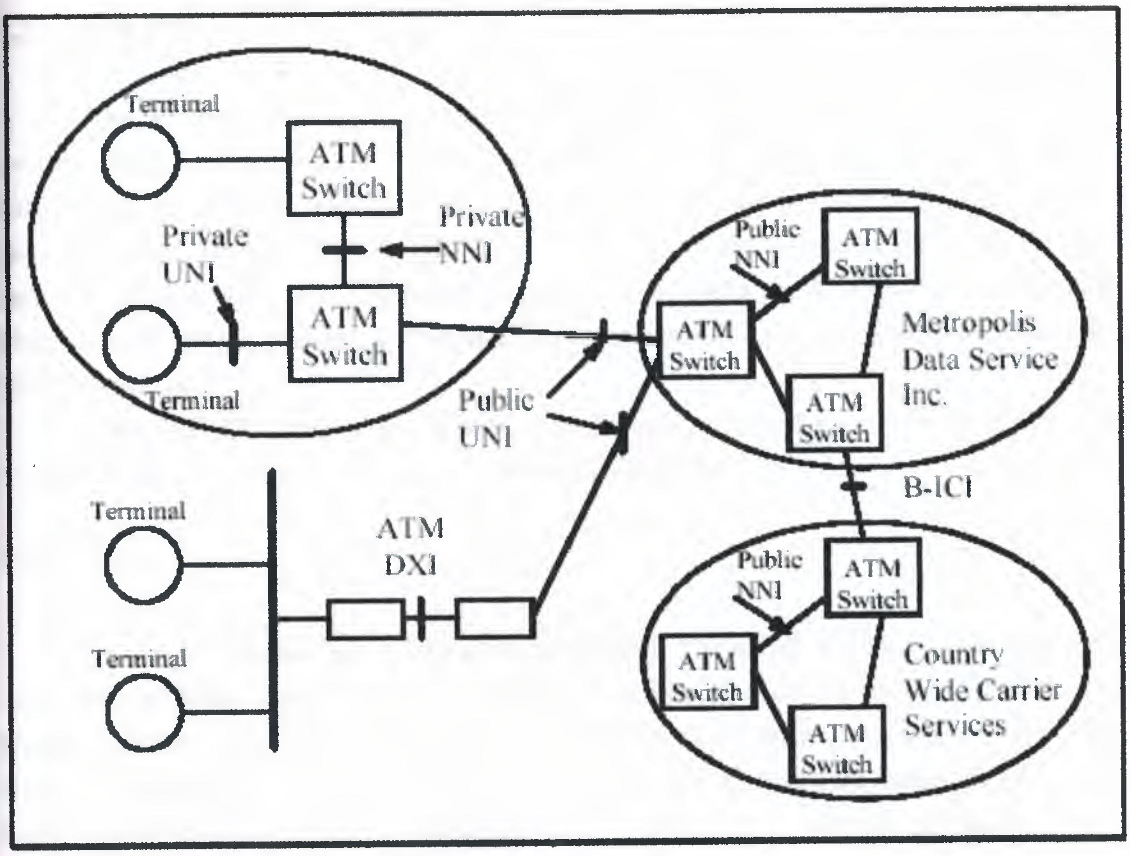

Figure 5-2. ATM Interfaces and Networks.

Within a private ATM network, there is the issue of connecting multiple switches together into an ATM network. This is referred to as the Network Node Interface (NNI). In some ways, the NNI is misnamed because it's really more than an interface. It is a protocol that allows multiple devices to be interconnected in somewhat arbitrary topologies and still work as one single network.

There's a corresponding protocol in the public arena called the public NNI. It has basically the same function, but, because of the context of the problem that's being addressed, it ends up in detail to be quite different.

The private NNI protocol is being specified by The ATM Forum and the public NNI is being specified by ITU. One of the major differences is that in the case of the public NNI, there's going to be a strong dependence on the signalling network.

The B-ICI specifies how two carriers can use ATM technology to multiplex multiple services onto one link, thereby exchanging information and co-operating to offer services. This is discussed in more detail in a later section.

5.2.1 ATM DXI

The ATM Data Exchange Interface (DXI) allows a piece of existing equipment -- in this case, a router -- to access the ATM network without having to make a hardware change. The hardware impact is in a separate Channel Service Unit / Data Service Unit (CSU/DSU). Typical physical layers for the DXI are like V35 or high-speed serial interface (HSSI). Since this is a data-oriented interface, the frames are carried in HDLC frames. All that is required is a software change in the router and the CSU-DSU to perform the "slicing" segmentation and reassembly (SAR) function.

The CSU-DSU takes the frames, chops them up into cells, does traffic shaping if it's required to abide by the traffic contract, and ends up with a UNI.

5.2.2 NNI

Figure 5-3 shows the cell structure at the NNI. The difference in the header at the NNI, as compared with the UNI, is found in the first four bits. Instead of being a GFC field (also see figure 3.6), they have been dedicated to the virtual path identifier, extending it to 12 bits. For a connection between two switches, the total number of virtual path connections goes up by a factor of 16, making this a very large number. This is most desirable at places in the network where a lot of connections go over one physical path.

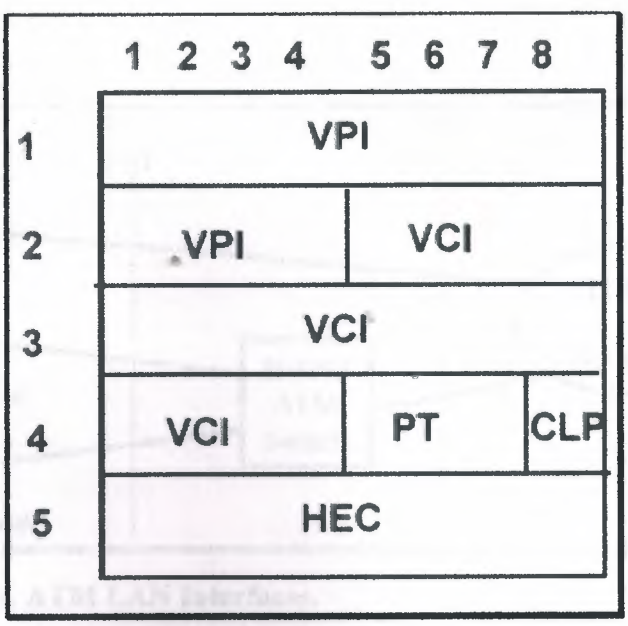

The NNI also provides some other functions, like distribution of topology information. Also, in the case of network failure, the switches in the network need to know that the failure happened, which connections have broken, and which ones need to be re-established.

5.2.3 B-ICI

The Broadband Inter-Carrier Interface (B-ICI), in its initial version, is a multiplexing technique. It specifies how two carriers can use ATM technology to multiplex multiple services onto one link, thereby exchanging information and co-operating to offer services. The services specified in the B-ICI are: cell relay service, a circuit emulation service, frame relay, and SMDS. Users of the carrier network don't "see" this interface, but it is important because it will help provide services across carriers.

5.3 ATM for LANs

Traditional Local Area Networks(LANs) like Ethernet, Token Ring and Token Bus are limited in speed (10 Mbits/s) and thus are limited to particular type of (mainly data) application. For multimedia applications, the bandwidth requirement is high and the information is a combination of voice, video and data and it requires a transfer mode capable of transporting and switching these different types of information. ATM satisfies this requirement for LANs. ATM Forum was created to specify the interfaces for ATM LANs. ATM LANs may be used to interconnect multiple LANs and also to connect to powerful workstations and servers. ATM Forum is a non-profit organisation founded in 1991 and is actively involved in the definition of ATM LAN and also in the eventual deployment of a universal B-ISDN network. A simple ATM LAN can be shown in Figure 5-4.

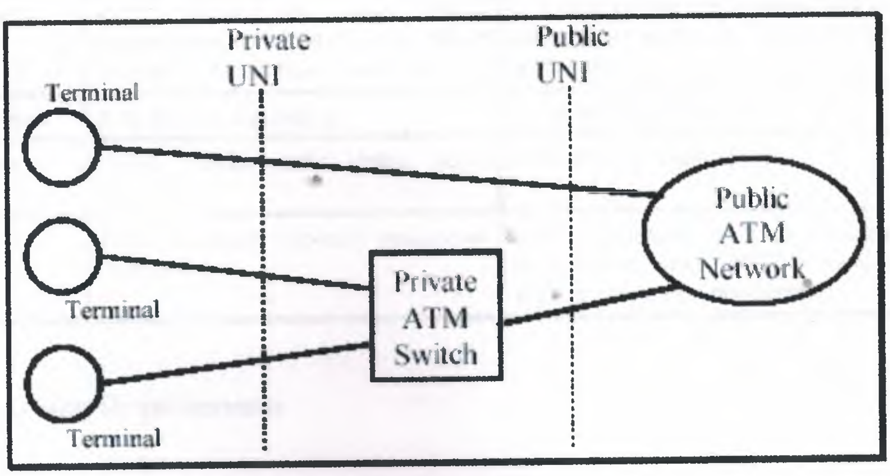

Figure 5-4. ATM LAN Interfaces.

ATM allows transmission capacities up to 622 Mbits/s which is enough for most LAN applications. ATM LANs are expected to be mainly star configured and they have simpler network management functions. Deployment of ATM in LANs will also stimulate faster deployment of the public BISDN.

5.4 ATM interworking

5.4.1 The Issues

Interworking between Frame Relay Bearer service (FRBS) and ATM needs to take the following issues into account:

- Needs proper mapping of frame relaying loss priority and congestion control indications.
- Requires procedures of negotiation for frame relaying frame size.
- Able to provide message-mode unassured operation without flow control, and
- Able to transfer user data immediately upon the establishment of the connection using AAL parameter negotiation procedures.

Frame relay service and Cell relay service (ATM) can be compared quickly by using the following table.

	Cell Relay Service (ATM)	Frame Relay Service
Connection type	Virtual Paths & Virtual Channels	Virtual Channels
Local address	VPI/VCI	DLCI (data link connection identifier)
PDU (protocol data unit) length	Fixed (48+5=53 octets)	Variable
Delineation method	HEC cell delineation method	Flag delineation method
Traffic descriptor	PCR (peak cell rate), SCR (sustainable cell rate), and Burst tolerance or maximum burst size	CIR (committed information rate), Bc (committed burst size), and Be (exceed burst size)
Priority indication	CLP (cell loss priority)	D/E (Discard/Eligibility) bit
Error protection	ATM header only (plus AAL function)	CRC-16 over entire frame
Congestion indication	EFCI (Explicit forward congestion indication)	FECN (forward explicit congestion notification) and BECN (backward explicit congestion notification)

5.4.2 Interworking Requirements

The following is the interworking arrangement between FRBS and B-ISDN Class C, message mode, non-assured operation. For Interworking in the **C (Control) plane**, Call Control mapping is provided in such a way that **U (User) plane** connections are made and released in both interworking networks, with interconnection in the interworking function. C-Plane procedures must provide for the negotiation of U-Plane parameters like throughput, maximum frame size etc. The mapping between frame relay and B-ISDN traffic descriptors are to be standardised. For Interworking in the U plane, there are two sets of service conditions for interworking the B-ISDN Class C services (message mode, unassured operation) and FRBS which are: B-ISDN directly supports FRBS, and B-ISDN supports Cell Relay Service (CRS) which can interwork with FRBS.

5.5 ATM Signalling

The Signalling capability for ATM Networks has to satisfy the following functions.

- Set up, maintain and release ATM virtual channel connections for information transfer.
- Negotiate the traffic characteristics of a connection (CAC algorithms are considered for these functions.)

Signalling functions may also support multi-connection calls and multi-party calls. Multiconnection call requires the establishment of several connections to set up a composite call comprising various types of traffic like voice, video, image and data. It will also have the capability of not only removing one or more connections from the call but also adding new connections to the existing ones. Thus the network has to correlate the connections of a call. A multi-party call contains several connections between more than two end-users like conferencing calls. Signalling messages are conveyed out-of band in dedicated signalling virtual channels in broadband networks. There are different types of signalling virtual channels that can be defined at the B-ISDN user-to-network interface. They can be described as follows.

- Meta-signalling virtual channel is used to establish, check and release point-to-point and selective broadcast signalling virtual channels. It is bi-directional and permanent.
- Point-to-point signalling channel is allocated to a signalling endpoint only while it is active.

These channels are also bi-directional and are used to establish, control and release VCCs to transport user information. In a point-to-multipoint signalling access configuration, etasignalling is needed for managing the signalling virtual channels.

5.6 ATM Addressing

A signalling protocol needs some sort of addressing scheme. Private networks will probably use OSI NSAP type addressing, primarily because an administrative process exists.The public carriers will probably use E164 numbers, which is the fancy way to say telephone numbers. A major reason is that the telephone networks have become very mechanised.

5.6.1 Private Address Format

In order for an addressing scheme to be useful, there must be a standardized address format that is understood by all of the switches within a system. For instance, for making phone calls within a given country, there is a well-defined phone number format. When calling between countries, this format is usually modified to include information like a "country code." Each call setup message will contain the information in these fields twice -- once identifying the party that is being called (destination) and once identifying the calling party (source).

Figure 5-5 shows the three address formats that have been defined by the ATM Forum. The first byte in the address field identifies which of the addess formats is being used. (Values for this field other than the three listed here are reserved and/or used for other functions.)

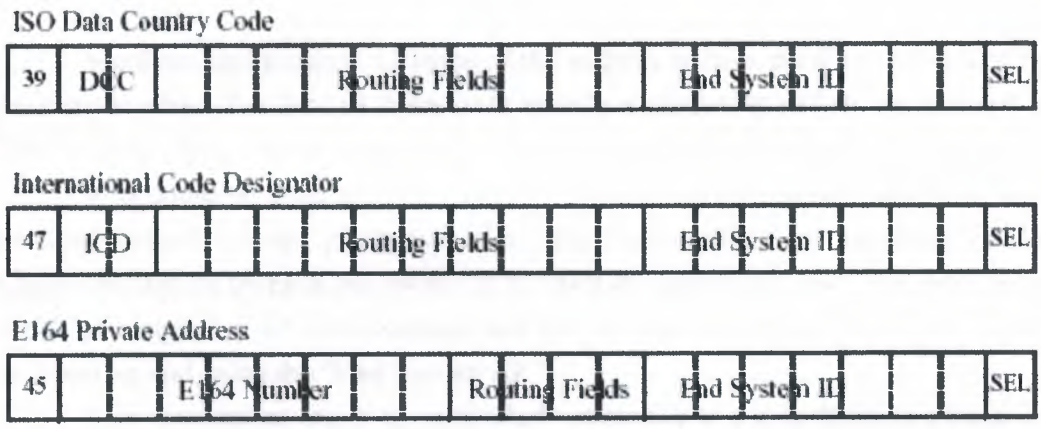

Figure 5-5. ATM address format.

The three address formats are:

(1) Data Country Code: Data Country Code numbers are administered by various authorities in each country. For instance, ANSI has this responsibility in the US. The DCC identifies the authority that is responsible for the remainder of the "Routing Fields."

(2) International Code Designator (ICD): ICDs are administered on an international basis by the British Standards Institute.

(3) E.164 Private Addresses: E.164 addresses are essentially telephone numbers that are administered by telephone carriers, with the administering authority identity encodes as a part of the E.164 number.

Regardless of the numbering plan used, it is very important that an ATM network implementer obtain official globally uniqe numbers to prevent confusion later on when ATM network islands are connected together.

Following the DCC or ICD fields -- or immediately following the E.164 in the case of the E.164 format -- is the "Routing Field." For DCC and IDC, this is the information that contains the address that is being called (or is placing the call).

This "Routing Field" can be thought of as an address space. The term "Routing Field" implies that there is more to the field than a simple address though. In particular, the addressing mechanism will very probably be hierarchical to assist in the routing. In the E.164 option, the use of the "routing field" is not defined at this time.

Each address in the routing field may refer to a particular switch, or it may even refer to a particular UNI on a switch. If it refers only to a switch, then more information will be needed to find the exact UNI that is specified. On the other hand, if it specifies a UNI, then this is sufficient to serve as a unique, globally significant address.

5.6.2 Address Registration

We discussed the first 13 bytes of the address field in the Figure 5-5. Let's consider now the case in which the first 13 bytes only specify a particular switch, as opposed to a particular UNI. In this case, the switching system must still find the appropriate UNI for the call.

This could be done using the next 6 bytes, called the "End System ID." End systems, or terminals, could contain additional addressing information. For instance, the terminal could supply the last six bytes to the switch to identify the particular UNI. This way, and entire switch could be assigned a 13 byte address, and the individual switch would then be responsible for maintaining and using the "End System ID."

This mechanism might be particularly attractive to a user desiring a large "virtual private network," so that the user would obtain "switch addresses" from an oversight organization and then locally administer the the end-system IDs. This would have the advantage of allowing the user organization to administer the individual addresses without involving the outside organization.

However, anyone outside the organization desiring to call a given UNI would have to know values for both the Routing Field and the End System ID.

The six bytes of End System ID are not specified, so the use can be left up to implementor. A common anticipated use of the End System ID is to use the 6 bytes (48 bits) for the unique 48 bit MAC address that is assigned to each Network Interface Card (NIC). Of course, both the ATM switch and the ATM terminal must know these address in order to route calls, send signalling messages etc. This information can be obtained automatically using the ILMI (Integrated Link Management Interface). The switch typically will provide the 13 most significant bytes (Routing Field) while the terminal provides the next 6 bytes (End System ID). The Selector (SEL) byte is not used by the ATM network, but it is passed transparently through the network as a "user information field." Thus, it can be used to identify entities in the terminal, such as a protocol stack.

5.7 ATM Network Management

The basic management capability for ATM is the Interim Local Management Interface (ILMI) specification of the UNI 3.0 document (now it is called Integrated Link Management Interface).

There are two basic parts to the ILMI. The first is a manager-type application called the UNI Management Entity (UME) that performs the function of an SNMP management application. The second is an SNMP agent and ATM management MIB.

Both the end station and the ATM switch have both parts -- the UME and the ILMI agent and MIB. This allows either to request or supply information or operations to the other. This is unlike standard SNMP where there is a hierarchical manager to agent relationship. This design was chosen by The ATM Forum as a solution that would work well in the majority of product environments. The SNMP protocol is transported directly over the AAL5 layer of ATM between the two devices.

The ILMI MIB provides to management applications the capability to control and monitor the ATM link and physical layers of the protocol. In addition, address registration is supported by the MIB. At the present time the ILMI management capability is the primary ATM management function. Based on the UNI 3.0 specification, the ILMI can be used for managing the basic ATM network resource relationships. ILMI can be used to manage the functional requirements between an ATM end station and the ATM switch or hub. It can also be used to manage the relationship between an ATM switch in the private network and a switch in the public network. In the case of the end station it is referred to as a "private" UNI, and in the case of the private to public switch it is referred to as a "public" UNI.

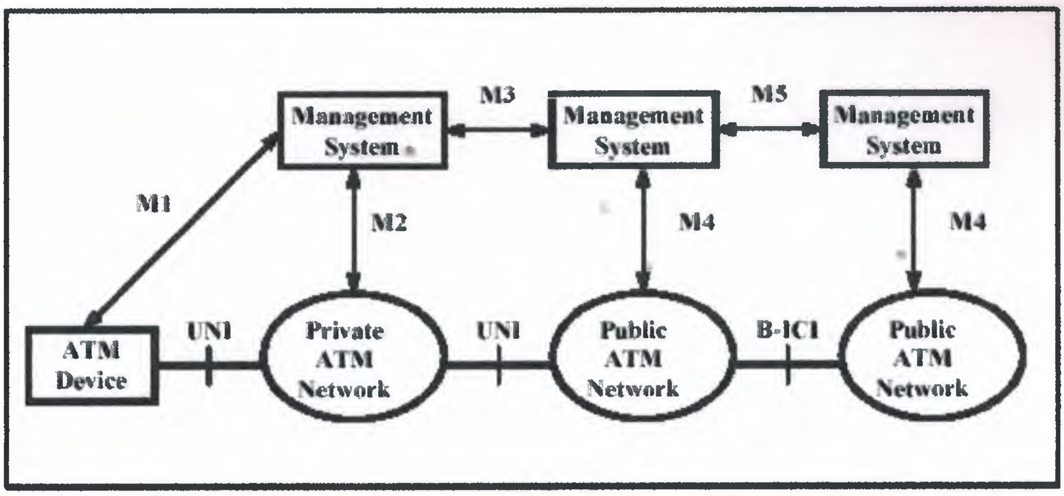

Figure 5-6. Network Management Model

Figure 5-6 shows the overall model for ATM network management. Note that the ILMI is for management of the UNI. The management functions between each of the entities in the network is designated as Mn for M1 through M5, as shown in the graphic.

In this model, the M1 and M2 functions are based on SNMP. M3, although it involves a carrier, is also based on SNMP. Other interfaces will use OSI management protocols. Work in the ATM Forum has focused on ILMI, M3, M4, and M5.

CHAPTER SIX : TRAFFIC MANAGEMENT AND CONTROL IN ATM NETWORKS

ATM networks must fairly and predictably allocate the resources of the network. In particular, the network must support various traffic types and provide different service levels. For example, voice requires very low delay and low delay variation. The network must allocate the resources to guarantee this. The concept used to solve this problem is called traffic management. When a connection -- a channel or a path -- is to be set up, the terminal initiating the service specifies a traffic contract. This allows the ATM network or network operator to examine the existing network utilisation and determine whether in fact a connection can be established that will be able to accommodate this usage. If the network resources are not available, the connection can be rejected. If it does look possible, the network will find a path through the network such that there is enough capacity to meet this traffic characteristic. While this all sounds fine, the problem is that the exact traffic characteristics for a given application are seldom known exactly.

Consider a file transfer. While we may think that we understand that application, in reality we're not certain ahead of time how big the file's going to be, or even how often a transfer is going to happen. Consequently, one cannot necessarily identify precisely what the traffic characteristics are.

Thus, the idea of traffic policing. The network "watches" the cells coming in on a connection to see if they abide by the contract. Those that violate the contract have their CLP bit set, which means they're going to be subject to loss. It doesn't mean these cells necessarily will be discarded, but it does mean that if the network starts to get into a congested state, these cells will be discarded first.

In theory, if the network resources are allocated properly, discarding all the cells with a cell loss priority bit marked will result in maintaining a level of utilisation at a good operational point in the network. Consequently, this is critical in being able to achieve the goal of ATM: to guarantee the quality of service (or the different kinds of quality of service) you need for the different traffic types.

There are many functions involved in the traffic control of ATM networks which are given as follows.

6.1 Connection Admission Control

This can be defined as the set of actions taken by the network during the call set-up phase to establish whether a VC/VP connection can be made. A connection request for a given call can only be accepted if sufficient network resources are available to establish the end-to-end connection maintaining its required quality of service and not affecting the quality of service of existing connections in the network by this new connection.

There are two classes of parameters which are to be considered for the connection admission control. They can be described as follows.

- Set of parameters that characterise the source traffic i.e. Peak cell rate, Average cell rate, burstiness and peak duration etc.
- Another set of parameters to denote the required quality of service class expressed in terms of cell transfer delay, delay jitter, cell loss ratio and burst cell loss etc.

6.2 Usage Parameter Control (UPC) and Network Parameter Control (NPC)

UPC and NPC perform similar functions at User-to-Network Interface and Network-to-Node Interface respectively. They indicate the set of actions performed by the network to monitor and control the traffic on an ATM connection in terms of cell traffic volume and cell routing validity. This function is also known as "Police Function". The main purpose of this function is to protect the network resources from malicious connection and to enforce the compliance of every ATM connection to its negotiated traffic contract. An ideal UPC/NPC algorithm meets the following features.

- Capability to identify any illegal traffic situation.
- Quick response time to parameter violations.
- Less complexity and much simplicity of implementation.

6.3 Priority Control

CLP (Cell Loss priority) bit in the header of an ATM cell allows users to generate different priority traffic flows and the low priority cells are discarded to protect the network performance for high priority cells. The two priority classes are treated separately by the network Connection Admission Control and UPC/NPC functions to provide two requested QOS classes.

6.4 Network Resource Management

Virtual Paths can be employed as an important tool of traffic control and Network resource management in ATM networks. They are used to simplify Connection Admission Control (CAC) and Usage/Network parameter control (UPC/NPC) that can be applied to the aggregate traffic of an entire virtual path. Priority control can also be implemented by segregating traffic types requiring different quality of service (QOS) through virtual paths. VPs can also be used to distribute messages efficiently for the operation of particular traffic control schemes like congestion notification. Virtual paths are also used in statistical multiplexing to

separate traffic to prevent statistically multiplexed traffic from being interfered with other types of traffic, for example guaranteed bit rate traffic.
Quality of Service (QoS) parameters includes:

CTD: It is the extra delay added to an ATM network at an ATM switch, in addition to the normal delay through network elements and lines. The cause of the delay at this point is the statistical asynchronous multiplexing. Cells have to queuing in a buffer if more than one cell contenting for the same output. It depend on the amount of traffic within the switch and thus the probability of contention.

CDV: The delay is depend on the switch/network design (such as buffer size), and the trafficcharacteristic at that moments of time. This results cell delay variation. There are two performance parameters associated with CDV: 1-point CDV and 2-point CDV. The 1-point CDV describes variability in the pattern of cell arrival events observed at a single boundary with reference to the negotiated 1/T. The 2-point CDV describes variability in the pattern of cell arrival events observed at an output of a connection with the reference to the pattern of the corresponding events observer at the input to the connection.

CLR: There are two basic cause of cell loss: error in cell header or network congestion. It is defined as: lost cells divided by the total transmitted cells.

CER: It is defined as: errored-cells / (successful-transferred-cells + errored-cells). The calculation excludes severely errored cell block. A cell block is a cell quence of N cells transmitted consecutively on a given connection.

6.5 Traffic Shaping

Traffic shaping changes the traffic characteristics of a stream of cells on a VPC or VCC by properly spacing the cells of individual ATM connections to decrease the peak cell rate and also to reduce the cell delay variation. Traffic shaping must preserve the cell sequence integrity of an ATM connection. Traffic shaping is an optional function for both network operators and end users. It helps the network operator in dimensioning the network more cost-effectively and it is used to ensure conformance to the negotiated traffic contract across the user-to-network interface in the customer premises network.
Traffic characteristics can be described by using the following parameters:

- **PCR**: Peak cell rate is the maximum rate at which the send is palnning to send.
- **SCR**: Sustained cell rate is the expected or required cell rate averaged over a long time interval.
- **MCR**: Minimum cell rate is the minimum number of cells/second that the customer considers as acceptable.

• **CDVT**: Cell delay variation tolerence tells how much variation will be presented in cell transmission times.

6.6 Congestion Control in ATM

Congestion control plays an important role in the effective traffic management of ATM networks Congestion is a state of network elements in which the network can not assure the negotiated quality of service to already existing connections and to new connection requests. Congestion may happen because of unpredictable statistical fluctuations of traffic flows or a network failure.

Congestion control is a network means of reducing congestion effects and preventing congestion from spreading. It can assign CAC or UPC/NPC procedures to avoid overload situations. To mention an example, congestion control can minimise the peak bit rate available to a user and monitor this. Congestion control can also be done using explicit forward congestion notification (EFCN) as is done in Frame Relay protocol. A node in the network in a congested state may set an EFCN bit in the cell header. At the receiving end, the network element may use this indication bit to implement protocols which will lower the cell rate of an ATM connection during congestion.

6.7 Generic Cell Rate Algorithm (GCRA)

The traffic contract is based on something called the Generic Cell Rate Algorithm (GCRA). Sometimes it's referred to as a "continuous leaky bucket," and this analogy is used here. The algorithm specifies precisely when a stream of cells either violates or does not violate the traffic contract. Consider a sequence of arrivals of cells. This sequence is run with the algorithm to determine which cells (if any) violate the contract.

The algorithm is defined by two parameters: the increment parameter "I" and the limit parameter "L." Think about this as a bucket with a hole in it. One unit of liquid leaks out of the bucket for every cell time. Consequently, on a given interface, cells have a specific time because they're fixed length.

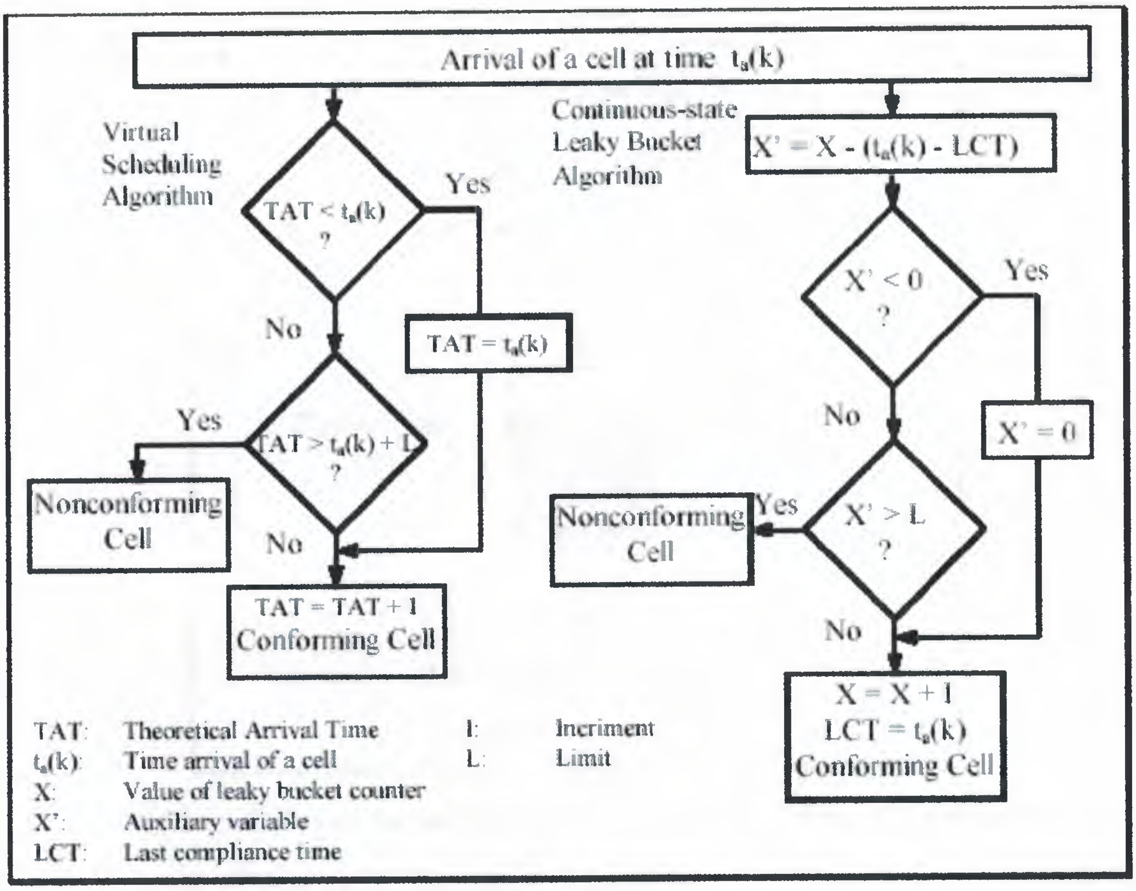

Figure 6-1. The generic cell rate algorithm (GCRA).

The GCRA can be implemented by either of the two algorithms defined by ATM forum: virtual scheduling algorithm or continuous-state leaky bucket. Figure 6-1 shows a flow chart of the algorithm.

The two algorithms served the same purpose: to make certain that cells are conforming (arrival within the bound of an expected arrival time), or nonconforming (arrival sooner than an expected arrival time). To make this a little more concrete, assume that "water" is being poured into the bucket and that it leaks out at one unit of water per cell time. Every time a cell comes into the network that contains data for this connection, "I" units of water are poured into the bucket. Of course, then the water starts to drain out. Figure 6-2 is the leaky bucket illustrating the GCRA.

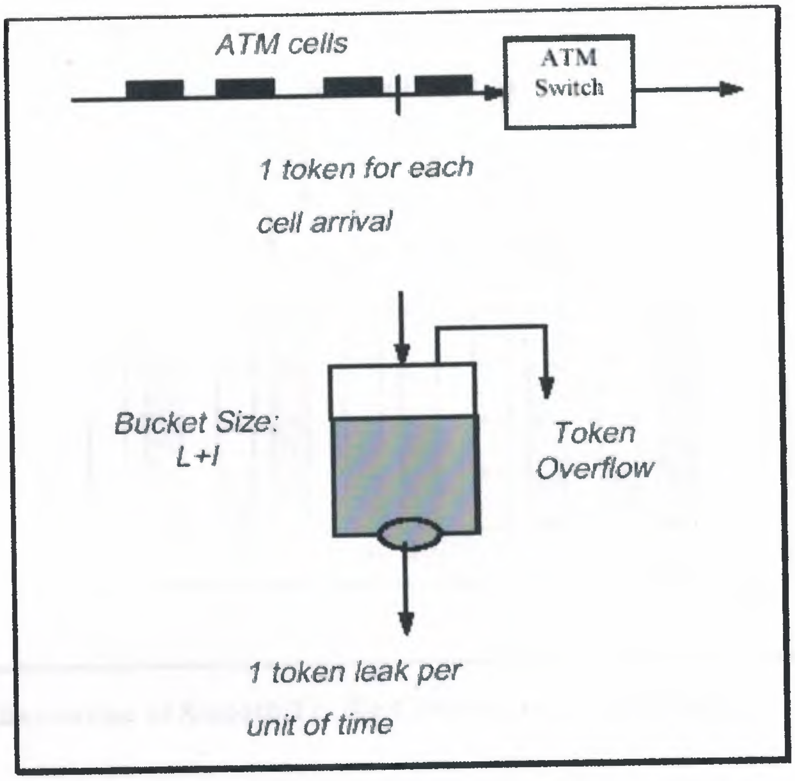

Figure 6-2. Leaky Bucket Illustrating Generic Cell Rate Algorithm (GCRA).

The size of the bucket is defined by the sum of the two parameters. Any cell that comes along that causes the bucket to overflow when the "I" units have poured in violates the contract.If the bucket was empty initially, a lot of cells can go into the bucket, and the bucket wouldeventually fill up. (Then you'd better slow down.) In fact, the overall rate that can be handled is the difference between the size of "I" and the leak rate. "I" affects the long-term cell rate "L" because it affects the size of the bucket. This controls how cells can burst through the network. The next two sections illustrate this.

6.7.1 Smooth Traffic

Consider this generic cell rate algorithm with a smooth traffic example.In Figure 6-3, the cell times separated left to right equally in time. The state of the bucket just before the cell time is represented by t-, and the state of the bucket just afterwards, a is represented by t+.

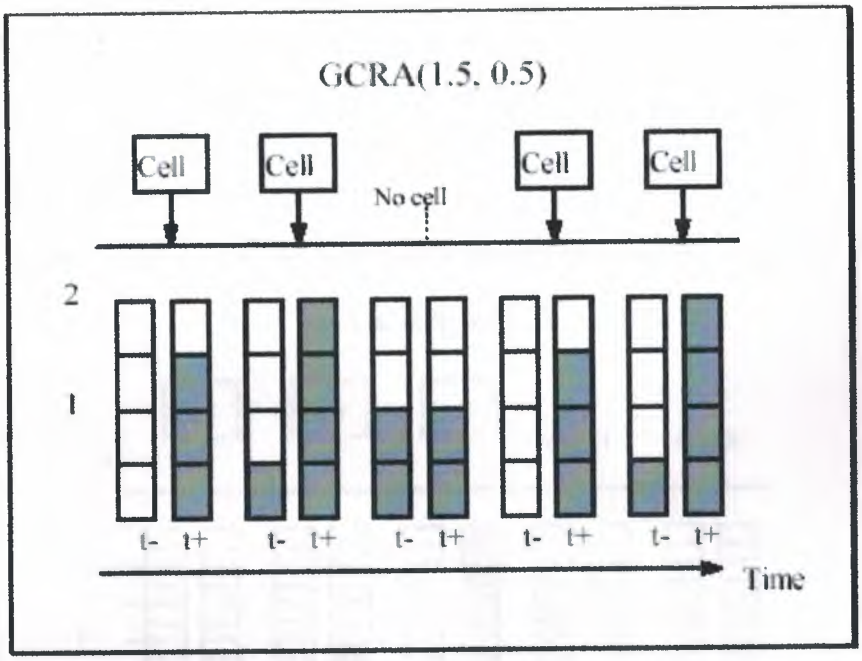

Figure 6-3. An Illustration of Smooth Traffic Coming to the Leaky Bukect - GCRA(1.5, 0.5).

Assume the bucket is empty and a cell comes in on this connection. We pour one-and-a-half units of water into the bucket. (Each cell contains one-and-a-half units of information. This is the increment parameter "I". However, we can only leak one unit per cell-time.) By the time we get to the next cell time, one unit has drained out, and, of course, by carefully planning this example, another cell comes in so you put the "I" units in. Now the bucket is one-half plus one and a half - It's just exactly full.

At the next time, if a cell came in, that cell would violate the contract because there's not enough room to put 1.5 units into this bucket. So let's assume that we're obeying the rules. We don't send a cell and this level stays the same and then it finally drains out, and of course, you can see we're back where we started.

The reason this is a "smooth" traffic case is because it tends to be very periodic. In this case, every two out of three cell times a cell is transmitted, and we assume that this pattern goes on indefinitely. Of course, two out of three is exactly the inverse of the increment parameter, 1.5. This can be adjusted with the "I" and the leak rate so that the parameter can be any increment desired -- 17 out of 23, 15 out of 16, etc. There's essentially full flexibility to pick the parameters to get any fine granularity of rate.

6.7.2 Bursty Traffic

Now consider a bursty traffic example. To make this bursty, increase the limit parameter to 7, and just to slow things down, the increment parameter is 4.5, so the bucket is 11.5 deep.

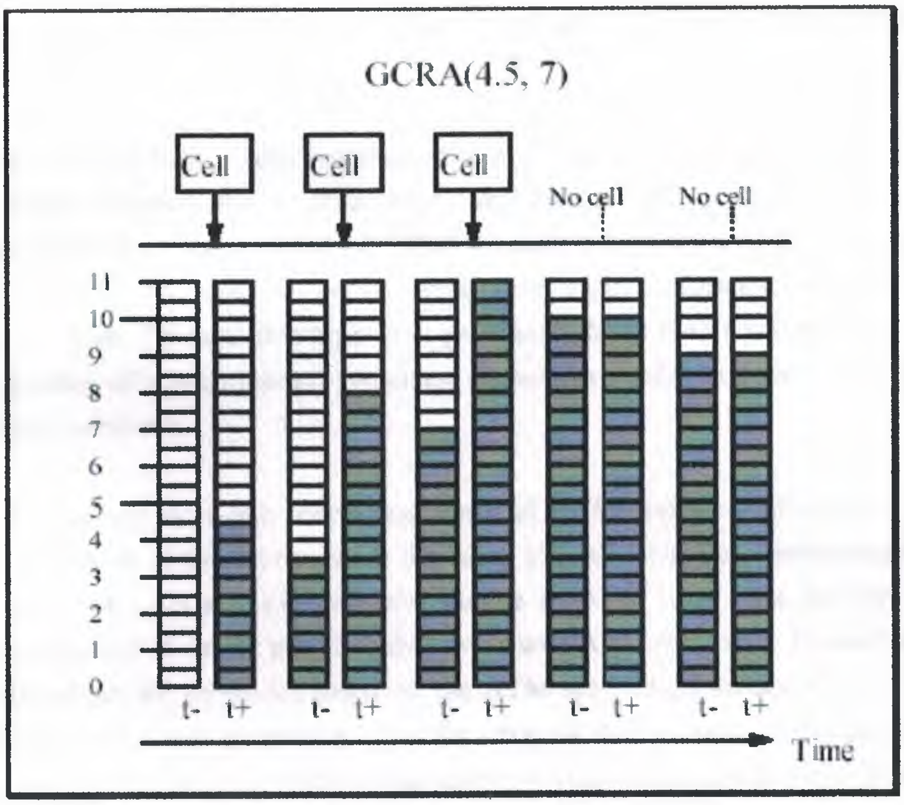

Figure 6-4. An Illustration of Bursty Traffic Coming to the Leaky Bukect - GCRA(4.5, 7).

As this example sends three cells, the information builds up and the bucket is exactly full after three cells. Now the rate is still only draining one unit of water per time but the increment is 4.5.

Obviously, you're going to have to wait quite a while here before you can send another cell. If you wait a long enough for the bucket to empty completely, another burst of three cells may be accepted. This illustrates the effect of increasing the limit parameter to allow more bursty type of traffic. Of course, this is especially critical for a typical data application.

CONCLUSION

The comunication requirementsof the network involve data, images, voice and video, possibly all integrated together in some instances. Clearly, since many network transaction (calls) involve voice and video information, its time-sensitive nature means that, for each call a path must be provided thorough the network that has a guaranteed transfer delay associated with it.

Therfore, the bandwidth requirements of a multiservice workstations are significantly higher than those required for a data only workstation. Hence in a network for the interconnection of multiservice workstations, shared-medium topologies of the type used for data are inappropiate since the cost of providing a very high performance interface in each workstation is too high. To meet this type of requirement, ATM LANs utilize a mesh topology comprising a number of interconnected switching exchanges similar in principle to thad used in existing telephone networks.

As with existing data only networks, users of multiservices workstations, as well as needing to communicate with other users at the same site, also want to communicate with user connected to an ATM LAN at a different site. Hence an ATM LAN has getways to existing (legacy) internets an also to newer private/public wide area ATM networks. To meet these needs, new generations of private networks based on the ATM are being introduced. Also, the public carriers are introducing a new generation of public network that is based on the same tecnology. This is comprised of ATM MANs which, in the future, will be connected at ATM WANs.

The sum up, Atm networks is the most important service in comunication system. Also in big network systems, data transfer will be very fast. Ther is a new technology wich is going to develop ATM systems.

References

1. Fred Halsall "Data Comunications, Computer Networks And Open Systems" Fourth Edition
2. Dervis Z. Deniz, "ISDN *and Its Applications to LAN Interconnection*", Mc-Graw Hill, 1994, ISBN: 0-07-707883-7.
3. Oliver G. Ibe , "Essentials of ATM Networks and Services"
4. G.C. Sacked and C.Y. Metz, " ATM and Mulriprotocol Networking", Mc-Graw Hill,1996.
5. H.Dutton and Peter Lenhard, "Asynchronus Transfer Mode (ATM) Technical Oveview", Second Ed, Prentice Hall, 1995.
6. B.Dorling, D.Freedman and J.Burger, "Internetworking Over ATM : An Introduction", Prentice Hall, 1996.