YAKIN DOĞU ÜNİVERSİTESİ

Faculty of Engineering

Computer Engineering

COM-400
Graduation Project

ATM NETWORKS

Submitted by: Gökhan UYSAL
20000704

Submitted to: Halil ADAHAN

Lefkoşa 2004

Preamble

This tutorial is an attempt to qualitatively present the ATM concepts, and to introduce it gently to readers unfamiliar with the subject. The purpose of this writeup is to bring interested readers up to speed on what the heck is this "ATM". There are also some "opinions" presented in this writeup. Please feel to disagree. And if I may have mis-stated some fact or something is in error, I would be happy to learn about it.

If the reader is interested in knowing the exact bit locations in the ATM header, or how to design and implement an ATM interface card, or other facts and figures about interconnect speeds etc, he/she is directed towards the copious ATM standards committees documents where the latest and greatest information is available in the most excruciating detail.

Why do we need ATM?

ATM was developed because of developing trends in the networking field. The most important parameter is the emergence of a large number of communication services with different, sometimes yet unknown requirements. In this information age, customers are requesting an ever increasing number of new services. The most famous communication services to appear in the future are HDTV (High Definition TV), video conferencing, high speed data transfer, videophony, video library, home education and video on demand.

This large span of requirements introduces the need for one universal network which is flexible enough to provide all of these services in the same way. Two other parameters are the fast evolution of the semi - conductor and optical technology and the evolution in system concept ideas - the shift of superfluous transport functions to the edge of the network. Both the need for a flexible network and the progress in technology and system concepts led to the definition of the Asynchronous Transfer Mode (ATM) principle.

The situation in the telecommunication world before ATM

Today's telecommunication networks are characterized by specialization. This means that for every individual telecommunication service at least one network exists that transports this service.

Each of these networks was specially designed for that specific service and is often not at all applicable to transporting another service.

When designing the network of the future, one must take into account all possible existing and future services.

The networks of today are very specialized and suffer from a large number of disadvantages:

Service Dependence

Each network is only capable of transporting one specific service.

Inflexibility

Advances in audio, video and speech coding and compression algorithms and progress in VLSI technology influence the bit rate generated by a certain service and thus change the service requirements for the network. In the future new services with unknown requirements will appear. A specialized network has great difficulties in adapting to new services requirements.

Inefficiency

The internal available resources are used inefficiently.

Resources which are available in one network cannot be made available to other networks.

It is very important that in the future only a single network will exist and that this network is service independent.

This implies a single network capable of transporting all services, sharing all its available resources between the different services.

It will have the following advantages:

Flexible and future safe

A network capable of transporting all types of services that will be able to adapt itself to new needs.

Efficient in the use of its available resources, all available resources can be shared between all services, such that an optimal statistical sharing of the resources can be obtained.

Less expensive

Since only one network needs to be designed, manufactured and maintained the overall costs of the design, manufacturing, operations and maintenance will be lower.

Progress in technology - ATM is possible

The definition of a service independent network has been influenced by an evolution in technology and system concepts.

System Concept Progress

The ideal network in the future must be flexible. The most flexible network in terms of bandwidth requirements and the most efficient in terms of resource usage, is a network based on the concept of packet switching.

Any bandwidth can be transported over a packet switching network and the resources are only used when useful information has to be transported.

The basic idea behind the concept changes is the fact that functions must not be repeated in the network several times if the required service can still be guaranteed when these functions are only implemented at the boundary of the network.

Progress in Technology in recent years large progress has occurred both in field of electronics and in the field of optics.

Broadband communication systems can be developed based on different technologies, the most promising being CMOS. (Complementary Metal Oxide Semiconductor) Cmos allows high complexity and reasonably high speed (up to 200 to 300 Mbits/s). The low power dissipation of Cmos is particularly important, and allows the realization of high complexity, high speed systems on a very small chip surface.

With the increased complexity per chip, the system cost can easily be reduced since the large integration will continuously allow the volume of the system to shrink or to increase the functionality at a constant cost.

Optical technology is also evolving quite rapidly.

Optical fiber has been installed for transmission services for several years.

Performance Requirements from ATM

In the future broadband network a large number of services have to be supported. These services are:

Low speed like telemetry, low speed data, telefax,

Medium speed like hifi sounds, video telephony, high speed data,

Very high speed like high quality video, video library...

A single typical service description does not exist. All services have different characteristics both for their average bit rate and burstiness.

To anticipate future unknown services we must try to characterize as general a service as possible.

The optimal transfer mode should support the communication of various types of information via an integrated access. Ideally the transfer mode must provide the capability to transport information, whatever type of information is given at the network, very much like the electricity network, which provides power to it's customers without regarding the way the customer uses his electricity.

Two other important factors are:

Semantic transparency - determines the possibility of network to transport the information error free.

The number of end to end errors introduced by the network is acceptable for the service

No system is perfect. Most of the imperfections of telecommunication systems are caused by noise. Other factors contribute to a reduced quality: limited resources causing blocking; any system errors. One of the most important parameters used to characterize imperfections is the BER (bit error rate) - the ratio between erroneous bits and transmitted bits. *Time transparency - determines the capability of the network to transport the information through the network from source to destination in a minimal time acceptable for the service. Time transparency can be defined as the absence of delay and delay jitter(different part of the information arrive at the destination with different delay). The value of end to end delay is an important parameter for real time services, such as voice and video. If the delay becomes too large echo problems may arise in a voice connection.

SERVICE	BER	DELAY
Telephony	10^(-7)	25 - 500 ms
Data transmission	10^(-7)	1000 ms
Broadcast Video	10^(-6)	1000 ms
Hifi Sound	10 ^(-5)	1000 ms

The basic principles of ATM

Information transfer

ATM is considered a packet oriented transfer mode based on:

Asynchronous time division multiplexing

The use of fixed length cells

An ATM cell structure is displayed in the following figure:

Each cell consists of an information field and a header.

The header is used to identify cells belonging to the same virtual channel and to perform the appropriate routing.

To guarantee a fast processing in the network, the ATM header has very limited function. Its main function is the identification of the virtual connection by an identifier which is selected at call set up and guarantees a proper routing of each packet. In addition it allows an easy multiplexing of different virtual connections over a single link.

The information field length is relatively small, in order to reduce the internal buffers in the switching node, and to limit the queuing delays in those buffers - small buffers guarantee a small delay and a small delay jitter as required in real time systems. The information field of ATM cells is carried transparently through the network. No processing is performed on it inside the network.

All services (voice, video, and data) can be transported via ATM, including connectionless services.

Routing

ATM is connection oriented. Before information is transferred from the terminal to the network, a logical/virtual connection is set.

The header values are assigned to each section of a connection for the complete duration of the connection, and translated when switched from one section to another. Signalling and user information are carried on separate virtual channels Two sorts of connections are possible:

Virtual Channel Connections VCC

Virtual Path Connections VPC

When switching or multiplexing on cells is to be performed, it must first be done on VPC, then on the VCC.

Virtual Channels

This function is performed by a header sub field - VCI. Since the ATM network is connection oriented each connection is characterized by a VCI which is assigned at call set up. A VCI has only a local significance on the link between ATM node and will be translated in the ATM nodes. When the connection is released, the VCI values on the involved links will be released and can be reused by other connections.

An advantage of this VCI principle is the use of multiple VCI values for multicomponent services. For instance video telephony can be composed of 3 components: voice, video and data each of which will be transported over a separate VCI. This allows the network to add or remove components during the connection. For instance, the video telephony service can start with voice only and the video can be added later.

Virtual Path

The network has to support semi-permanent connections, which have to transport a large number of simultaneous connections. This concept is known as virtual path.

All ATM switches can be schematically described as follows. A number of incoming links (I1, I2, In) transport ATM information to the switch, where depending on the value of the header this information is switched to outgoing link (O1, O2, On). The incoming header and the incoming link number are used to access a translation table. The result of the access to the table is an outgoing link and a new header value.

Resources

As ATM is connection oriented, connections are established either semi-permanently, or for the duration of a call, in case of switched services.

This establishment includes the allocation of a VCI (Virtual Channel Identifier)and/or VPI (Virtual Path Identifier), and also the allocation of the required resources on the user access and inside the network. These resources are expressed in terms of throughput and Quality of Service.

They may be negotiated between user and network for switched connection during the call set up phase

Let's look at the following topics

ATM Cell Identifiers

ATM cell identifiers are:

Virtual Path Identifier

Virtual Channel Identifiers

Payload Type Identifiers

They support recognition of an ATM cell on a physical transmission medium. Recognition of a cell is a basis for all further operations.

VPI and VCI are unique for cells belonging to the same virtual connection on a shared transmission medium. As such they are limited resources. Within a particular virtual circuit, cells may be further distinguished by their PTI, which cannot be allocated freely but depends on the type of payload carried by the cell. This field indicates whether the cell is carrying user information to be delivered transparently through the network or special network information. In case the field indicates network information, part of the information field indicates the type of network control whereas the remaining part of information field may be processed inside the network.

Throughput

Bandwidth has to be reserved in the network for each virtual connection. ATM offers the possibility to realize resources saving in the total bandwidth needed when multiplexing traffic of many variable Bit Rate connections.

The amount which can be saved depends heavily on the number of multiplexed connections, on the burstiness of the traffic they carry, on the correlation between them and on the quality of service they require.

Quality of Service

The quality of service of a connection relates to the cell loss, the delay and the delay variation incurred by the cells belonging to that connection in an ATM network. For ATM, the quality of service of a connection is closely linked to the bandwidth it uses. When providing limited physical resources using more bandwidth increases the cell loss, the delay, and the delay variation incurred, i. e. decreases the QOS for cells of all connections which share those resources.

Usage Parameter Control

In ATM there is no physical limitation on the user access rate to the physical transmission medium, apart from the physical cell rate on the medium itself. Multiplexing equipment will do its utmost to avoid cell loss to offer the highest possible throughput whatever the user chooses to send.

As virtual connections share physical resources, transmission media and buffer space, unforeseen excessive occupation of resources by one user may impair traffic for other users. Throughput must be monitored at the user - network interface by a Usage Parameter Control function in the network to ensure that a negotiated contract per VCC or VPC between network and subscriber is respected.

It is very important that the traffic parameters which are selected for this purpose can be monitored in real time at the arrival of each cell.

Flow control

In principle, no flow control will be applied to information streams at the ATM layer of the network. In some cases it will be necessary to be able to control the flow of traffic on ATM connections from a terminal to the network. In order to cope with this a GFC (general flow control) mechanism may be used. This function is supported by a specific field in the ATM cell header. Two sets of procedure are used: Uncontrolled Transmission - for the use of point to point configuration.

Controlled Transmission - can be used in both point to point and shared medium configuration.

Another principle is no error protection on link by link basis. If a link in the connection, either the user to network link or the internal links between the network nodes, introduces an error during the transmission or is temporarily overloaded thereby causing the loss of packets, no special action will be taken on that link to correct this error (= no requesting for retransmission).This error protection can be omitted since the links in the network have a very high quality

Signalling

The negotiation between the user and the network with respect to the resources is performed over a separate signalling virtual channel. The signalling protocol to be used over the signalling virtual channel is an enhancement of those used in ISDN signalling.

ATM - The Layered Model

The OSI model is very famous and used to model all sorts of communication systems. We can model the ATM with the same hierarchical architecture - however only the lower layers are used. The following relations can be found:

The Physical layers are more or less equivalent to Layer 1 of OSI model, and mainly perform functions on the bit level.

The ATM layer can be located mainly at the lower edge of the layer 2 of the OSI model.

The adaptation layer performs the adaptation of higher layer protocols, be it signalling or user information, to the fixed ATM cells.

These layers can then further be divided into sublayers. Each sublayer performs a number of functions, to be explained in the following sections. Click on the section you are interested in.

LAYER SUBLAYERS AAL-

Adaptation layer CS

SAR ATM layer Physical Layer TC

PM PM - Physical Medium Sublayer

This sublayer is responsible for the correct transmission and reception of bits on the appropriate physical medium. At the lowest level the functions that are performed are medium dependent: optical, electrical...

In addition this sublayer must guarantee a proper bit timing reconstruction at the receiver. Therefore the transmitting peer will be responsible for the insertion of the required bit timing information and line coding. Transmission Convergence Sublayer In this sublayer bits are already recognized, as they come from the PM sublayer. This sublayer performs the following functions:

Adaptation to the transmission system used

Generation of the HEC (Header Error Check) of each cell at the transmitter, and its verification at the receiver

Cell delineation - the mechanism to perform cell delineation is based on the HEC. If a correct HEC is recognized for a number of consecutive cells it is assumed that the correct cell boundary is found. To avoid erroneous cell delineation on user information, the information field of each cell is scrambled at the transmitting side and descrambled at the receiving side. This ensures that the probability of finding a correct HEC in the information field is very low

Once the cell delineation has been found an adaptive mechanism uses the HEC for correction or detection of cell header errors. Isolated single bit errors are corrected.

Cell uncoupling - the sublayer ensures insertion and suppression of unassigned cells to adapt the useful rate to the available payload of the transmission system ATM Layer The following main functions are performed by the layer:

The multiplexing and demultiplexing of cells of different connections into a single cell stream

A translation of cell identifiers, which is required in most cases when switching a cell from one physical link to another in an ATM switch or cross connect. This translation can be performed either on the VCI or VPI separately, or on both simultaneously.

Providing the user of a VCC or VPC with one QOS class out of a number of Classes supported by the network.

MANAGEMENT FUNCTIONS: the header of user information cells provides for a congestion indication and an ATM user to ATM user indication.

Extraction (addition) of the cell header before (after) the cell is being delivered to (from) the adaptation layer

Implementation of flow control mechanism on the user network interface. ATM Adaptation Layer

This layer enhances the service provided by the ATM layer to a level required by the next higher layer. It performs the functions for the user, control and management planes and supports the mapping between the ATM layer and the next higher layer. The functions performed in the AAL depend on the higher layer requirements.

The AAL layer is divided into two sublayers:

SAR - the segmentation and reassembly sublayer

The main purpose of the SAR sublayer is segmentation of higher layer information into a size suitable for the payload of the consecutive ATM cells of a virtual connection, and the inverse operation, reassembly of contents of the cells of a virtual connection into data units to be delivered to the higher layer. CS - the convergence sublayer. This sublayer performs functions like message identification, time/clock recovery etc. AAL Service Data Units (SDU) are transported from one AAL Service Access Point to one or more access points through the ATM network. The AAL users will have the capability to select a given AAL - SAP associated with the QOS required to transport the SDU. Up to now four AALS have been defined-one for each class of service.

The classes of ATM services

The services which will be transported over the ATM layer are classified in four classes, each of which has its own specific requirements towards the AAL. The services are classified according to three basic parameters:

1. Time relation between source and destination:

For real time applications like phone conversation, a time relation is required. Information transfer between computers does not require a time relation.

2. Bit Rate

Some services have a constant bit rate; others have a variable bit rate.

3. Connection mode:

Connectionless or connection oriented

Four types of AAL protocols have been recommended up to now: AAL 1, AAL 2, AAL 3/4, AAL 5.

AAL 1 - Adaptation for constant bit rate services

Recommended for services such as digital voice and digital video. It is used for applications that are sensitive for both cell loss and delay. Constant Bit Rate (CBR) services require information to be transferred between source and destination at a constant bit rate after virtual connection has been set up. The Layer services provided by the AAL 1 to the user are:

Transfer of Service Data Units with a constant bit rate and their delivery with the same bit rate

Transfer of data structure information

Transfer of timing information between source and destination

Indication of lost or corrupted information which is not recovered by the AAL itself when needed

CSI - CS indication - 1 bit

SN - sequence Number - 3 bits

SNP - sequence number protection 4 bits

The SAR sublayer accepts a 47 octet block of data from the CS and then adds a one octet SAR-PDU header to each block.

At the receiving end, the SAR sublayer gets a 48 byte block from the ATM layer, and then separates the SAR PDU header. The SAR sublayer receives a sequence number value from the CS. At the receiving end this number is passed to the CS. IT may be used to detect loss and incorrect insertions of SAR payloads. The SNP is used for protection against bit errors. It is capable of single bit error correction and multiple bit error detection. The convergence Sublayer functions:

Handling of cell delay Variation * Handling of cell payload assembly delay

AAL 2 - Adaptation for variable bit rate services

These type AAL offers a transfer of information with a variable bit rate. In addition, timing information is transferred between source and destination. Since the source is generating a variable bit rate, it is possible those cells are not completely filled and that the filling level varies from cell to cell. Therefore more functions are required from the SAR .

The SN field (Sequence Number) contains the sequence number to allow the recovery of lost or misrouted cells.

The IT (Information Type) indicates the beginning of a message (BOM), continuing of a message (COM), the end of a message (EOM) or that the cell transports timing or other information.

BOM, COM or EOM indicate that the cell is the first, middle or last cell of a message, i.e. an information unit as defined in the CS layer with possibly a variable length.

The LI (length indicator) field indicates the number of useful bytes in partially filled cells.

The CRC field allows SAR to detect bit errors in the SAR SDU

In the CS sublayer the following functions have to be performed:

Clock recovery by means of insertion and extraction of time information.

Handling of lost or incorrectly delivered cells.

Forward error correction for video and audio services

AAL 3/4 - Adaptation for data services

This AAL is recommended for transfer of data which is sensitive to loss, but not to delay. The AAL may be used for connection oriented as well as for connectionless services, since functions like routing and network addressing are performed on the network layer. Two modes of AAL 3/4 are defined: * Message Mode

The AAL SDU is passed across the AAL interface in exactly one AAL Interface Data Unit (IDU) . This service is provided for the transport of fixed or variable length AAL SDU.

Streaming mode

The AAL SDU is passed in one or more AAL IDU. Transfer of these IDUs may occur separate in time. The service provided for long variable length AAL SDUs. The SAR sublayer functions:

Segmentation and reassembly of variable length CS PDUs. The SAR PDU contains two fields for this purpose:

1. ST Segment Type - indicates which part of the CS PDU is carried by the SAR PDU:

First, middle or last

2. LI Length Indicator

Error Detection - using CRC field

Multiplexing of multiple CS PDUs on a common bearer in the ATM layer. Multiplexing is supported by a multiplexing identifier.

The CS functions

Delineation and transparency of SDUs

Error detection and handling - Corrupted SDUs are either discarded or optionally delivered to the service specific convergence sublayer.

Buffer allocate size- each PDU carries up front an indication to the receiving entity of the maximum buffer required to receive the PDU * Abort - a partially transmitted PDU can be aborted

AAL 5 - Adaptation for data services

This AAL is recommended for high speed connection oriented data service. This AAL offers a service with less overhead and better error detection.

The SAR sublayer functions:

The SAR sublayer accepts variable length SAR SDUs which are multiples of 48 octets from the CS sublayer, and generates SAR PDUs containing 48 octets of data.

The CS functions

The functions implemented by the AAL5 are the same as the ones offered by the AAL 3/4 except that the AAL 5 does not give a buffer allocation size indication to the receiving peer entity. Also error protection in the AAL 5 is fully handled at the CS layer itself, instead of being shared between SAR and CS.

Why ATM?

ATM stands for Asynchronous Transfer Mode. ATM is a connection-orientated technique that requires information to be buffered and then placed in a cell. When there is enough data to fill the cell, the cell is then transported across the network to the destination specified within the cell. We can see that ATM is very similar to packet-switched networks, but there are several important differences:

1. ATM provides cell sequence integrity. i.e. cells arrive at the destination in the same order as they left the source. This may not be the case with other packet-switched networks.

2. Cells are much smaller than standard packet-switched networks. This reduces the value of delay variance, making ATM acceptable for timing sensitive information like voice.

3. The quality of transmission links has lead to the omission of overheads, such as error correction, in order to maximise efficiency.

4. There is no space between cells. At times when the network is idle, unassigned cells are transported.

It is these techniques that allow ATM to be more flexible than Narrow-band ISDN (N-ISDN), and hence ATM was chosen as the broadband access to ISDN by the CCITT (now ITU-TSS). The broadband nature of ATM allows for a multitude of different types of services to be transported using the same format. This makes ATM ideal for true integration of voice, data and video facilities on the one network. By consolidation of services, network management and operation is simplified. However, new terms of network administration must be considered, such as billing rates and quality of service agreements.

The flexibility inherent in the cell structure of ATM allows it to match the rate at which it transmits to that generated by the source. Many new high bit-rate services, such as video, are variable bit rate (VBR). Compression techniques create bursty data which is well suited for transmission using ATM cells.

The Protocol Reference Model.

In a similar way to the OSI 7-layer model, ATM has also developed a protocol reference model, consisting of a control plane, user plane and management plane. The model also incorporates SAPs, SDUs and PDUs that are also mentioned in the OSI layered approach. As the diagram below shows, the User plane (for information transfer) and Control plane (for call control) are structured in layers. Above the Physical Layer rests the ATM Layer and the

ATM Adaptation Layer (AAL). Management provides network supervision.

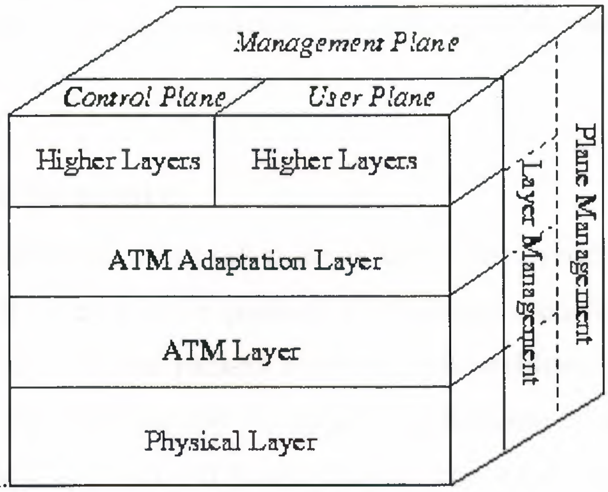

ATM Layer.

Layer/Sublayer	Function
ATM Adaptation Layer	
Convergence Sublayer	Convergence
Segmentation & Reassembly Sublayer	Segmentation & Reassembly
ATM Layer	Generic Flow Control Cell header generation/extraction Cell VPI/VCI translation Cell multiplex & demultiplex
Physical Layer	
Transmission Convergence Sublayer	Cell-rate decoupling HEC header generation/check Cell delineation
Physical Medium Sublayer	Bit timing Physical medium

Responsibilities

The ATM layer is responsible for transporting information across the network. ATM uses virtual connections for information transport. The connections are deemed virtual because although the users can connect end-to-end, connection is only made when a cell needs to be sent. The connection is not dedicated to the use of one conversation. The connections are divided into two levels:

- The Virtual Path
- The Virtual Channel

It is the properties of the VP and VC that allow cell multiplexing. There is a complication in that cell switching requires only the value of the VP identifier, VPI to be known.

Cell Structure.

The structure of the cell is important for the overall functionality of the ATM network. A large cell gives a better payload to overhead ratio, but at the expense of longer, more variable delays. Shorter packets overcome this problem, however the amount of information carried per packet is reduced. A compromise between these two conflicting requirements was reached, and a standard cell format chosen. The ATM cell consists of a 5-octet header and a 48-octet information field after the header. This is shown below.

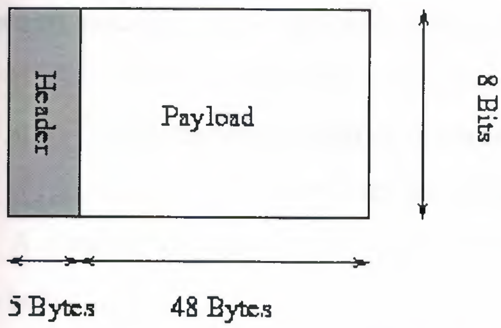

The information contained in the header is dependent on whether the cell is carrying information from the user network to the first ATM public exchange (User-Network Interface - UNI), or between ATM exchanges in the trunk network (Network-Node Interface - NNI). The formats of the two types of header are shown below. Notice the similarity between the two, with only the UNI having a Generic Flow Control, GFC, field.

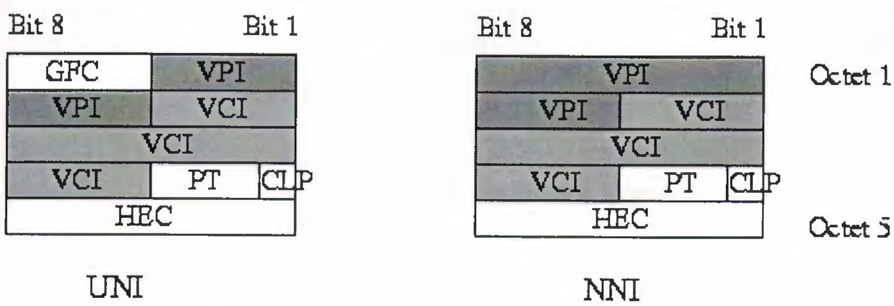

Virtual Channels

The connection between two endpoints is called a Virtual Channel Connection, VCC. It is made up of a series of Virtual channel links that extend between VC switches. The VC is identified by a Virtual Channel Identifier, VCI. The value of the VCI will change as it enters a VC switch, due to routing translation tables. Within a virtual channel link the value of the VCI remains constant. The VCI (and VPI) are used in the switching environment to ensure that channels and paths are routed correctl The VCI (and VPI) are used in the switching environment to ensure that channels and paths are routed correctly. They provide a means for the switch to distinguish between different types of connection.

There are many types of virtual channel connections, these include:

- User-to-user applications. Between customer equipment at each end of the connection.
- User-to-network applications. Between customer equipment and network node.
- Network-to-network applications. Between two network nodes and includes traffic management and routing.

Virtual channel connections have the following properties:

- A VCC user is provided with a quality of service, QoS, specifying parameters such as cell-loss ratio, CLR, and cell-delay variation, CDV.
- VCCs can be switched or semi-permanent.
- Cell sequence integrity is maintained within a VCC.
- Traffic parameters can be negotiated, using the Usage Parameter Control, UPC.

A detailed diagram showing the relationship between virtual channels and paths is shown below.

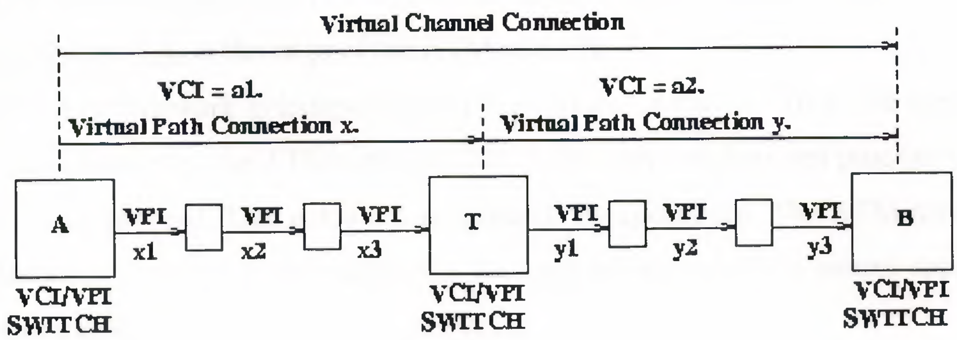

VIRTUAL CHANNEL VIEW

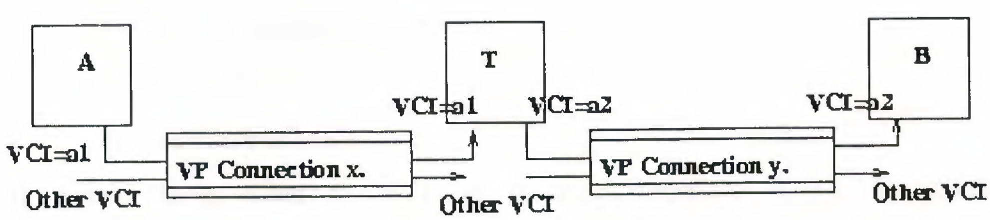

Virtual Paths

A virtual path, VP, is a term for a bundle of virtual channel links that all have the same endpoints. As with VCs, virtual path links can be strung together to form a virtual path connection, VPC. A VPC endpoint is where its related VPIs are originated, terminated or translated.

Virtual paths are used to simplify the ATM addressing structure. VPs provide logical direct routes between switching nodes via intermediate cross-connect nodes. A virtual path provides the logical equivalent of a link between two switching nodes that are not necessarily directly connected on a physical link. It therefore allows a distinction between logical and physical network structure and provides the flexibility to rearrange the logical structure according to traffic requirements. This is best shown in the diagram above.

As with VCs, virtual paths are identified in the cell header with the Virtual Path Identifier, VPI. Within an ATM cross-connect, information about individual virtual channels within a virtual path is not required, as all VCs within one path follow the same route as that path.

ATM Adaptation Layer

Responsibilities

The ATM Adaptation Layer, AAL, performs the necessary mapping between the ATM layer and the higher layers. This task is usually performed in terminal equipment, or terminal adaptors, TA, at the edge of the ATM network.

The ATM network is independent of the services it carries. Thus, the user payload is carried transparently by the ATM network. The ATM network does not process, or know the structure of the payload. This is known as semantic independence. The ATM network is also time independent, as their is no relationship between the timing of the source application and the network clock.

All of this independence must be built into the boundary of the ATM network, and falls into the realm of the AAL. The AAL must also cope with:

- Data flow to application
- Cell delay variation, CDV
- Loss of cells
- Misdelivery of cells

It would have been possible to develop seperate AALs for each type of telecommunication service offered, however the many common factors between services has meant that a small set of AAL protocols is sufficient to cover the envisaged possibilities. A telecommunication service is defined on the following parameters:

- Timing relationship between source and destination.
- Bit-rate.
- Connection mode.

Parameters such as communication assurance are treated as quality of service parameters. As a result, four classes of service have been defined.

Class:	A	B	C	D
Timing relationship between source and destination	required		not required	
Bit rate	constant	variable		
Connection mode	connection-orientated			connectionless

The class of service are general concepts, but these they are mapped onto different specific AAL types.

- Class A: AAL 1.
- Class B: AAL 2.
- Class C & D: AAL 3/4.
- Class C & D: AAL 5.

The AAL is organised on two sublayers:

- The Convergence Sublayer.
- The Segmentation and Reassembly Sublayer.

Information pertaining to the CS and SAR is found in the ATM dictionary.. The CS, which performs the tasks of processing cell delay variation, synchronisation and handling cell loss, is broken up into two parts:

- The Service Specific CS
- The Common Part CS

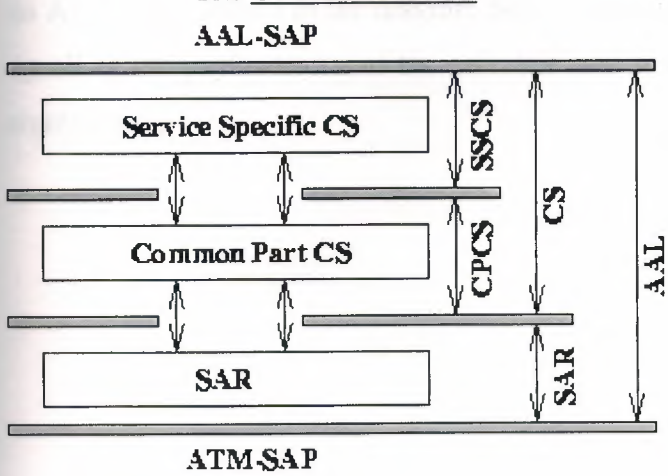

Again, information about these two sublayers is found in the dictionary. A diagram below shows the relationship between the layers and sublayers of the AAL.

Information that moves between layers of the AAL follows a naming convention. Protocol Data Units, PDUs, contain the information between peer layers, while Service Data Units,SDUa, pass data across Service Access Points, SAPs. This is shown clearly in a diagram in the ATM dictionary.

Below is a list of the defined AAL types. Contained with each type is a list of applications suited to that particular AAL.

AAL type 1.

- Circuit transport to support synchronous (e.g. 64KBit/s) and asynchronous (e.g. 1.5, 2 MBit/s) circuits.
- Video signal transport for interactive and distributive services.
- Voice band signal transport.
- High quality audio transport.

AAL type 2.

AAL 2 has not currently been defined, but services for this type may include:

- Transfer of service data units with a variable source bit-rate.
- Transfer of timing information between source & destination.

AAL types 3/4.

AAL 3 was designed for connection-orientated data, while AAL 4 for connectionless-orientated data. They have now been merged to form AAL 3/4. The structure of the layers for an AAL 3/4 is shown in the diagram below. Note how the user data for payload does not take up all of the payload area of the cell. We will see later that this reduces the usable bit rate significantly.

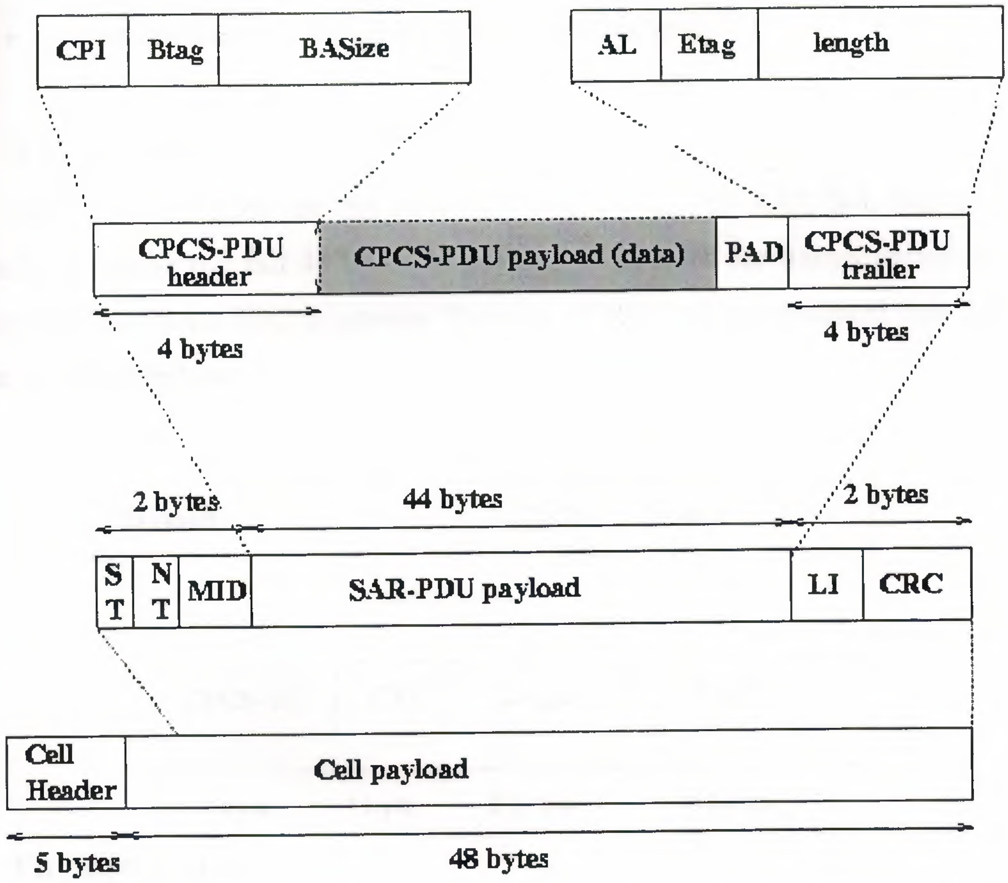

- ST Segment type (2 bits). Indicates whether segment is beginning, continuation, end or single segment message.
- SN Sequence Number (4 bits). Allows sequence of SAR-PDUs to be numbered modulo 16.
- MID Multiplexing identification (10 bits). Allows for more than one connection over a single ATM-layer connection. The value of the MID must be unique over the current VP only.
- LI Length indicator (6 bits). Indicates the number of bytes of CS-PDU information in the SAR-PDU, as the amount of information may not fill the 44 bytes available.
- CRC Cyclic redundancy check code (10 bits). Used to detect errors in the SAR-PDU. This include the CS_PDU and user data.
- CPI Common part indicator (1 octet).
- Btag Beginning tag (1 octet).
- BASize Buffer Size allocation (2 octets).
- PAD Padding (0 to 3 octets).
- AL Alignment (1 octet).
- Etag End tag (1 octet).

- length length of CPSU-PDU payload (2 octets).

AAL type 5.

AAL 5 is designed for the same class of service as AAL 3/4, but contains less overheads. It allows the full 48 bytes of payload to be used for transportation of CS-PDU segments, not just SAR-PDU segments. There is a CRC field incorporated into the CS-PDU field, as indicated below.

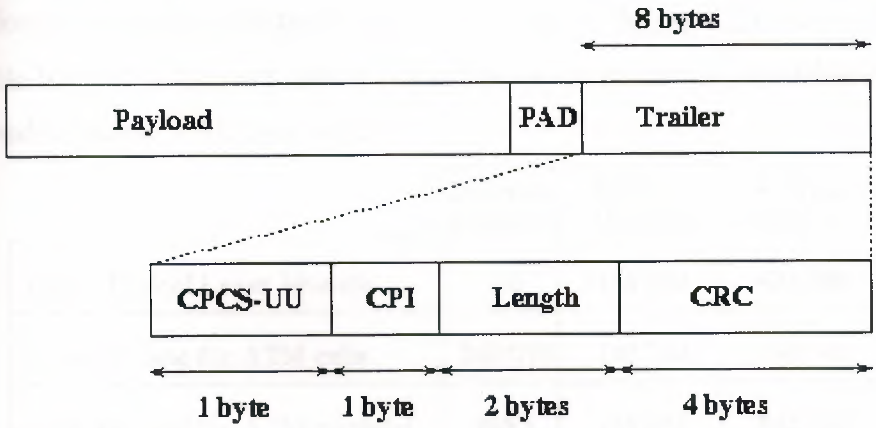

Physical Layer

Responsibilities.

The physical layer has two sublayers:

- Physical Medium sublayer.
- Transmission Convergence sublayer.

The Physical layer is responsible for the transmission of the data across a physical link, in much the same way as the physical layer of the OSI reference model. The diagram below shows the role of the interface between the ATM layer and the physical layer.

Transfer Capacity

The CCITT Recommendation I.432 defines two bit-rates for the physical layer:

□ 155 520 KBit/s.

□ 622 080 KBit/s.

The transportation medium may either be electrical or optical, and can use SDH-based or cell-based framing. Telecom Australia are currently introducing SDH into their network and so this tutorial will concentrate on this framing for ATM.

The bit rates mentioned above are the gross bit rates of the physical layer and hence contain transportation overheads of the carrier, and also of the layers above the physical layer

(ATM Adaptation Layer and ATM Layer). This causes the actual user data bit rate to be less than the gross rate by a significant amount. The values shown in the table below are based on a SDH frame structure. The column "fraction available" shows the ratio of payload to (payload plus header). Thus, a SDH frame (see below) allows 260 bytes of payload and 10 bytes of overhead and pointers. This gives a fraction of 260/270. Similarly for ATM cell formats, payload is 48 bytes and overheads 5 bytes, giving a fraction of 48/(48 +5) = 48/53. The final value of cell-payload bit rate does not allow for space taken up in the payload by AAL format types and related headers, (e.g. CRCs, MIDs, CPIs, etc). Thus the maximum available bit-rate to the user cannot reach that of the maximum available for cell-payload and is dependent on the AAL type used.

	Fraction available	STM-1 (KBit/s)	STM-4 (KBit/s)
Gross Pysical Layer bit-rate	1.0	155 520	622 080
Max. bit-rate for ATM cells	260/270	149 760	599 040
Max. bit-rate for ATM payload	48/53	135 631	542 526

ATM over SDH. SDH is a transmission hierarchy that allows ATM cells to be mapped into "containers", particularly the container C-4. These containers are then linked to a particular SDH frame using a pointer in the SDH overheads. Thus the C-4 containers are deemed "virtual" (e.g. VC-4) as they can swap frames. Notice in the diagram below that ATM cells may cross SDH frame boundaries due to non-integer multiples of ATM cells per frame.

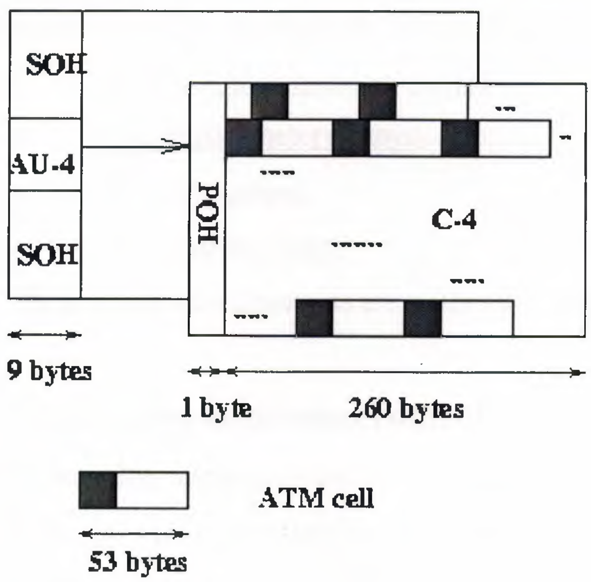

Connection-Orientated Service

Signalling Principles

ATM is a connection-orientated technique. As outlined in Section 2.1.1, A connection within the ATM layer consists of one or more links, each which is assigned an identifier.

A lot of applications, such as constant bit rate services (CBR) and X.25 data service are best handled by connection-orientated communications. With ATM and other connection-orientated techniques, a connection has to be established before information transfer takes place. ATM uses an out of band signalling system in dedicated signalling virtual channels, SVCs. There are different types of SVC for different requirements:

- The Meta-signalling virtual channel, MSVC, is bi-directional and permanent. It is used to establish, check and release point-to-point and selective broadcast SVCs.
- The point-to-point signalling channel is bidirectional and is used to establish, control and release VCCs and VPCs that carry user data.
- Broadcast SVCs are unidirectional and can send signalling messages to all, or select endpoints.
- General SVCs are like Broadcast SVCs, but do not allow selected groups.

Traffic Control

In order for a broadband network based on ATM to achieve a high level of performance, traffic control capabilities have to be introduced. The CCITT in Recommendation I.311 highlighted the following:

- Connection admission control.
- Usage parameter control.
- Priority control.
- Congestion control.

These control mechanisms are outlined below.

Connection admission control

Connection admission control is the set of actions taken by the network at the call setup phase in order to establish whether a VC/VP connection can be established. A connection can only be established if the network resources are available to provide the required quality of service. The introduction of a new connection should not affect the QoS of other established connections. Source traffic can be identified by parameters such as,

- Peak duration
- Average bit rate
- Burstiness
- Peak bit rate

Usage parameter control.

Usage parameter control is the set of actions taken by the network to monitor and control user traffic volume and cell routing validity. Its main purpose is to protect network resources from malicious as well as unintentional misbehaviour which can affect the QoS parameters of existing connections by detecting violations of negotiated parameters. UPC includes monitoring the following functions:

- Validity of VPI/VCI values.
- Monitoring VP/VC traffic volumes to check for violations.
- Monitoring total traffic volumes on links.

Priority Control.

Priority Control is determined using the cell loss priority bit in the cell header. Information can be broken into more and less important parts. Thus different components of the same signal will be treated differently by the network control mechanisms.

Congestion Control

Congestion is defined as a state of network elements in which, due to traffic overload, the network is not able to guarantee a QoS to already established connections and to new connection requests.

Congestion control tries to minimise congestion effects and avoid the problem spreading. Congestion control could, for example, reduce the peak bit rate available to a user.

Cell Delay Variation and Queues

As explained above, the small sized cells allow for small delay variation, CDV. This is useful for the transportation of isochronous media, which requires data (especially voice) to be sent at fixed intervals. Small delay variation allows for "virtual" isochronous transmission.

Traffic shaping schemes try to shape traffic into isochronous flow, with regular time intervals at the output. The leaky bucket is an example of a traffic shaping scheme. The leaky bucket algorithm uses a buffer of finite size that incoming traffic is placed into. Traffic is allowed to drain out of the bucket and sent on the network at a rate, p. Excess data that cannot fit into the buffer is discarded. The leaky bucket algorithm has the effect of shaping bursty traffic into a flow of equally spaced cells, each being emitted 1/p units of time after the previous cell. The size of the buffer limits the cell delay. Hence to limit CDV, a small buffer is required.

Connectionless Service.

ATM is connection-orientated communication. However, there are many applications, such as mail services and other data services that are characterised by small amounts of data, sent sporadically. To save time and expense, no connection is established - i.e. a connectionless service. User information is sent in a message containing all necessary addressing and routing information. This is used in local area networks that employ carrier sense multiple access with collision detection (CSMA/CD) network structures (e.g. ethernet). It is possible for ATM to be used in a connectionless configuration.

An ATM connectionless data service allows the transfer of information among service subscribers without the need for end-to-end call establishment. A connectionless data service will require the introduction of connectionless servers. The connectionless servers route cells to their destination according to the routing information contained in the cells.

The connectionless service sits on top of ATM, i.e. it is not integrated into the functionality of the ATM switch. This requires a direct connection between each user and the connectionless server. These connections can be semi-permanent or switched. The use of direct connection means that only n connections are required for n users. The diagram below indicates the provision of a connectionless service on ATM.

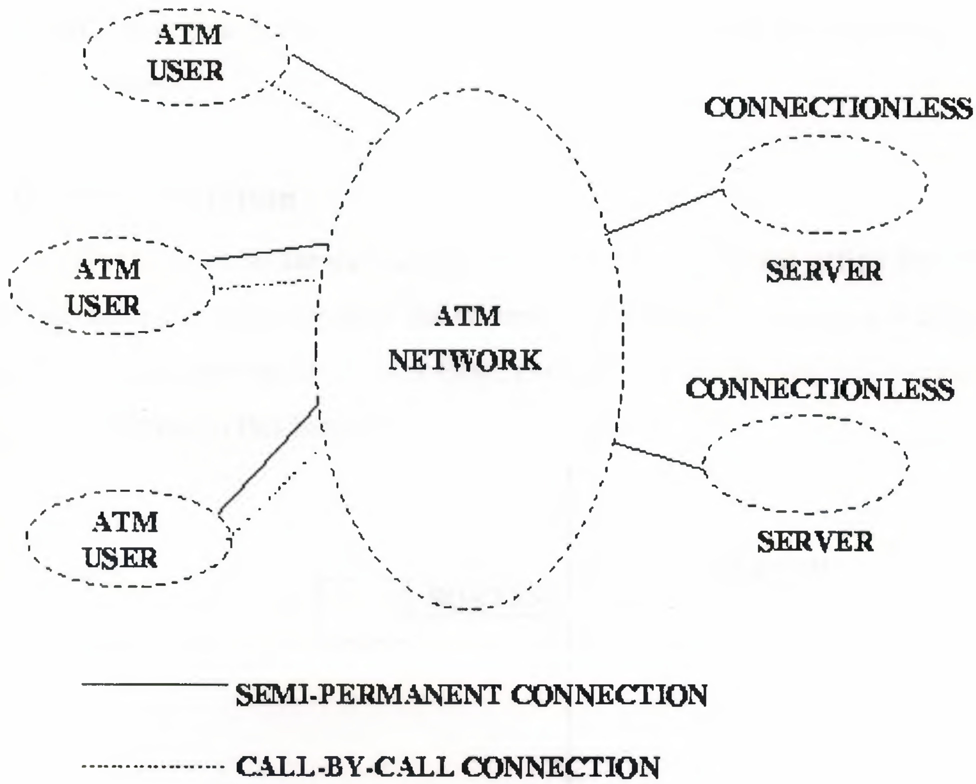

LAN Traffic over ATM

The first ATM networks are likely to be installed by companies that have a specific high bandwidth need. These could include single locations, between buildings (across a campus), or across a high speed (E3) link. Other solutions to the joining of LANs exist, such as FDDI, however these solutions are not suitable for the wide area networks (WANs), and the data must be transformed into something else for transmission. ATM, on the other hand, if used throughout the LAN, then the transition to a MAN or WAN would be "seamless", as the same language and technologies would be used throughout. This is an example of the scalability of ATM - the ability to handle different bit rates for different situations, and being able to upgrade to higher rates as technology progresses.

ATM LAN Network Configurations

This section of the tutorial covers topologies of different ATM network configurations. This includes some migrations towards ATM based solutions, as well as highlighting problems of other non ATM-based solutions.

As LANs increasingly require communication with each other, due to multimedia and other bandwidth hungry services, the connections between the LANs become overloaded and create a bottleneck. Although there are alternative solutions to this problem, ATM is the most "future-proof". Note that in the examples given below, ethernet is just one of the services that ATM can interface.

Current Situation

The diagram below shows the current situation in a typical office environment. The ethernet backbone that joins together the ethernet subnetworks becomes a bottleneck, as only one user can access the backbone at a time, even if they do not require the services of the entire line. A solution to this bottleneck must be found.

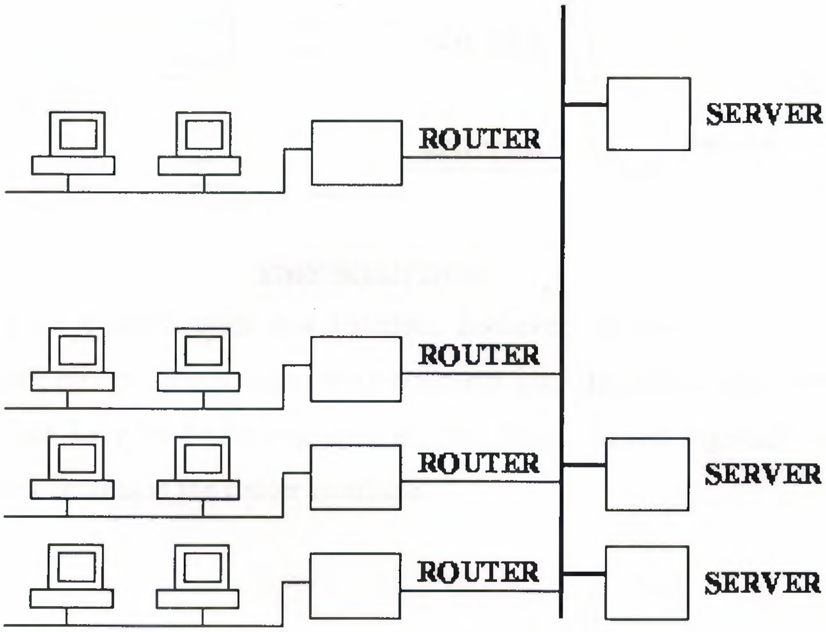

CURRENT SITUATION

FDDI Solution

FDDI provides a solution to the bottleneck problem, by increasing the speed of the backbone from 10MBit/s (in the ethernet case), to 100MBit/s (FDDI).

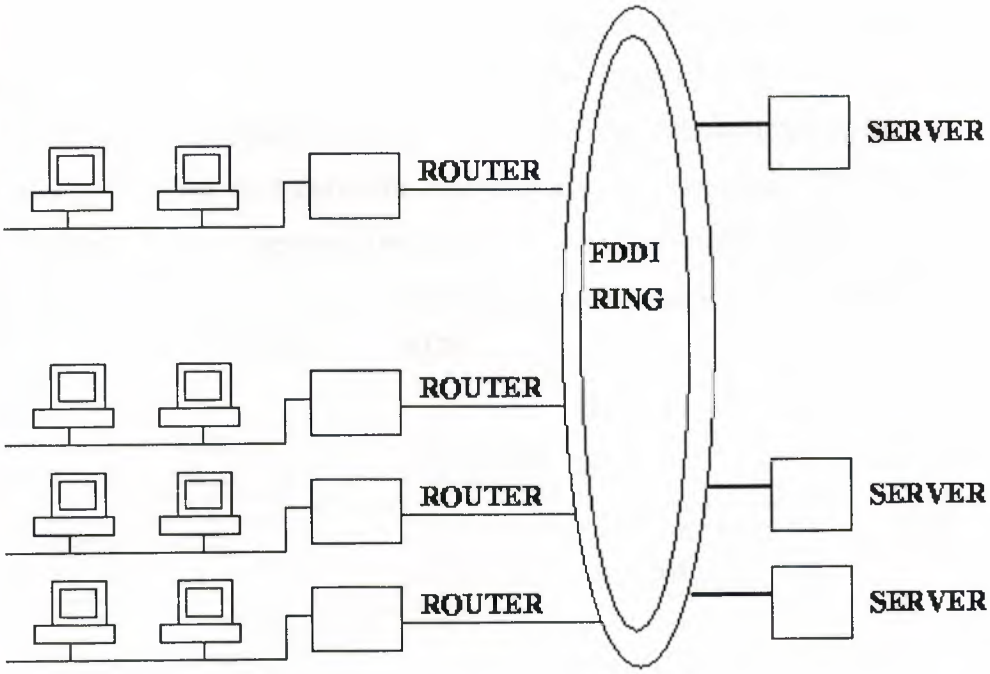

FDDI SOLUTION

The problem with this solution, however, is that it is not "future-proof". Future bandwidth hungry applications may soon eat into the 100MBit/s bandwidth, and the FDDI ring would have to be broken into smaller rings, linked together with routers. Hence the bottleneck returns at the router interface.

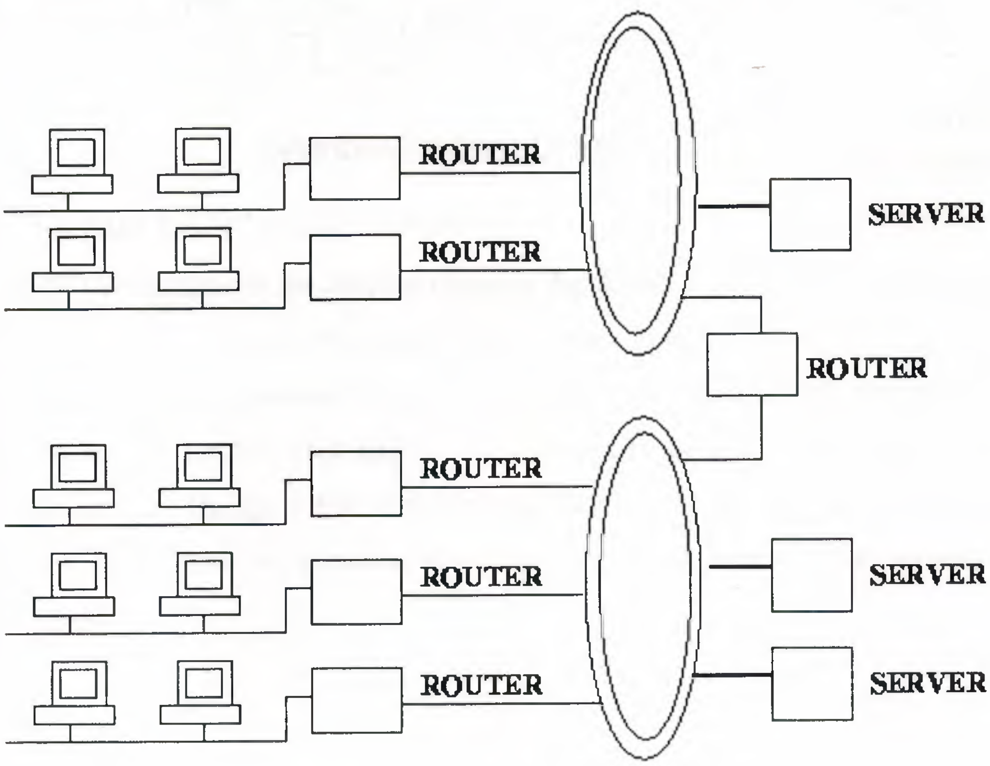

FDDI EXPANSION

Interim ATM solution

ATM can provide a solution to the bottleneck problem by replacing the backbone architecture with an ATM switch. The switch allows higher bandwidths to pass through, as it is not a single access system. I.e. multiple p[arties can communicate at the same time. This has a cumulative effect on bandwidth, allowing greater throughput. The ATM solution also allows different protocols between the routers and ATM switches (e.g. SDH, E1, E3), so that these connections are upgradable as the demand on that connection increases.

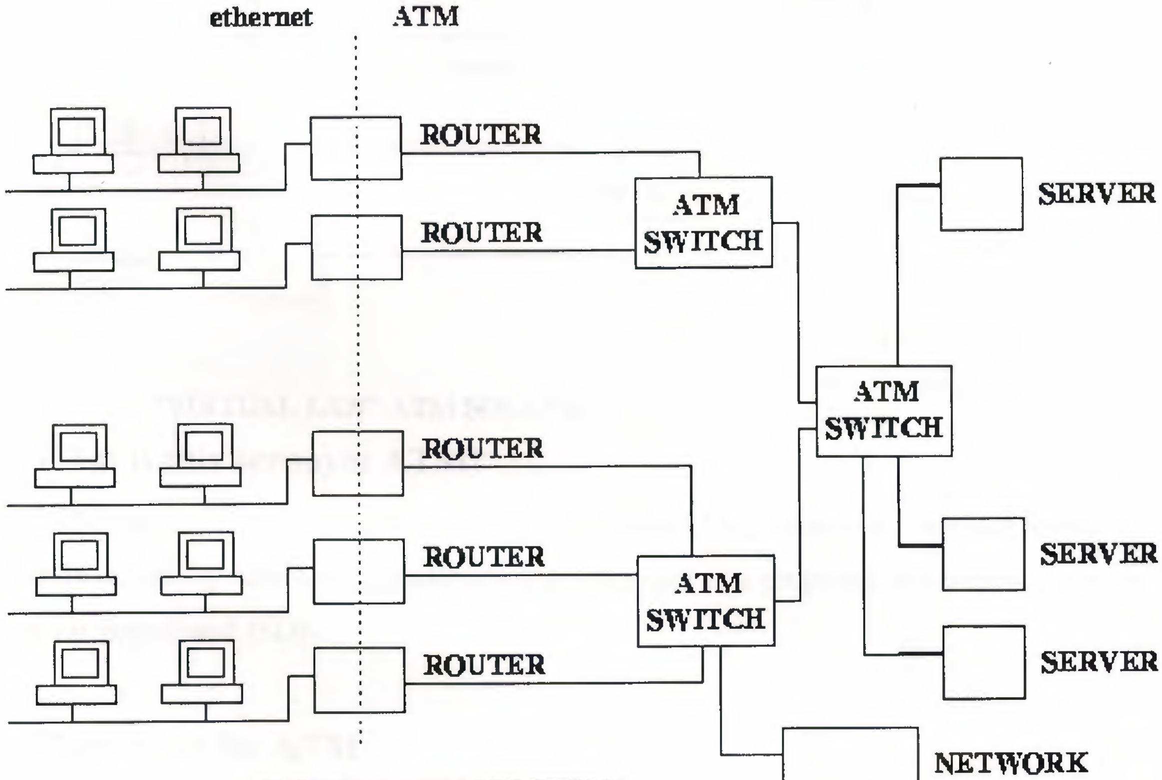

INTERIM ATM SOLUTION

"Virtual LAN" ATM solution

With the installation of adaptor cards in the ATM switch, (in this example ethernet), virtual LANs can be created. This means that workstations may be grouped to form a LAN, even though they are separated by physical links. The ATM switch provides the logical connections for the LANs. This allows workstations to be able to move physical location without the need to change LAN. The functionality of this system is provided by network management, that allows the administrator easy access to the entire ATM network through a remote terminal.

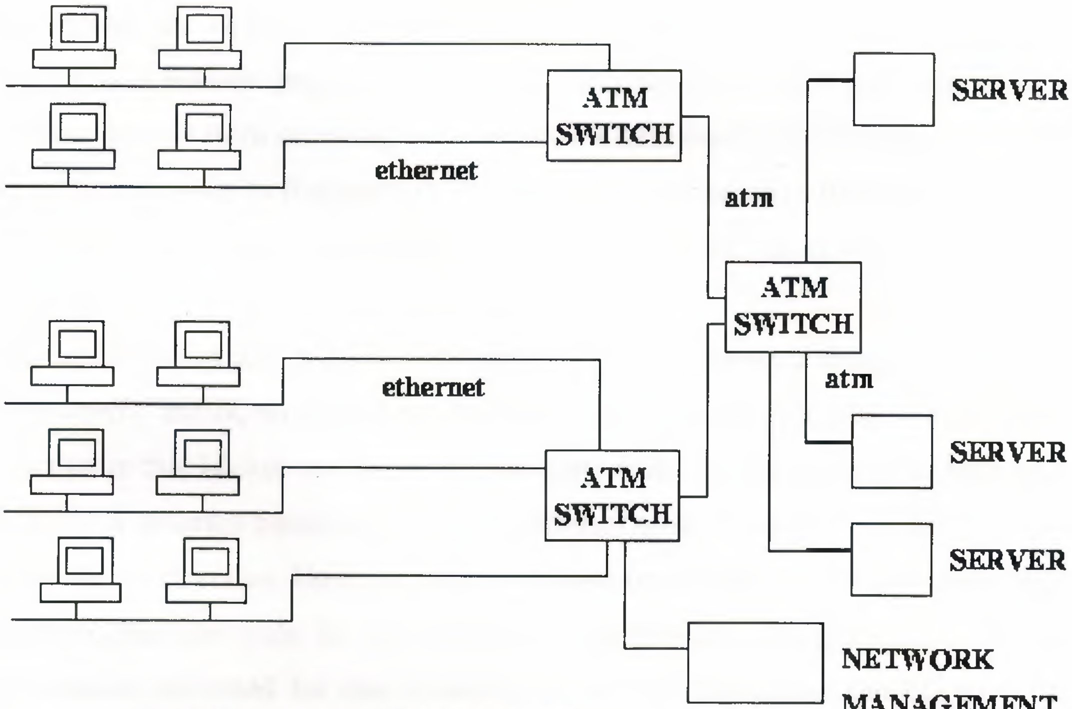

"VIRTUAL LAN" ATM SOLUTION

What is this acronym ATM?

ATM stands for (no not automated teller machines) "Asynchronous Transfer Mode". It is primarily driven by telecommunications companies and is a proposed telecommunications standard for Broadband ISDN.

Motivation for ATM

In order to understand what ATM is all about, a brief introduction to STM is in order. ATM is the complement of STM which stands for "Synchronous Transfer Mode". STM is used by telecommunication backbone networks to transfer packetized voice and data across long distances. It is a circuit switched networking mechanism, where a connection is established between two end points before data transfer commences, and torn down when the two end points are done. Thus the end points allocate and reserve the connection bandwidth for the entire duration, even when they may not actually be transmitting the data. The way data is transported across an STM network is to divide the bandwidth of the STM links (familiar to most people as T1 and T3 links) into a fundamental unit of transmission called time-slots or buckets. These buckets are organized into a train containing a fixed number of buckets and are labeled from 1 to N. The train repeats periodically every T timeperiod, with the buckets in the train always in the same position with the same label. There can be up to M different trains labeled from 1 to M, all repeating with the time period T, and all arriving within the time period T. The parameters N, T, and M are determined by standards committees, and are different for Europe and America. For the trivia enthusiasts, the timeperiod T is a historic legacy of the classic Nyquist sampling criteria for information recovery. It is derived from sampling the traditional 4 KHz bandwidth of analog voice signals over phone lines at twice its frequency or 8 KHz, which translates to a timeperiod of 125 usec. This is the most fundamental unit in almost all of telecommunications today, and is likely to remain with us for a long time.

On a given STM link, a connection between two end points is assigned a fixed bucket number between 1 and N, on a fixed train between 1 and M, and data from that connection is always carried in that bucket number on the assigned train. If there are intermediate nodes, it is possible that a different bucket number on a different train is assigned on each STM link in the route for that connection. However, there is always one known bucket reserved a priori on each link throughout the route. In other words, once a time-slot is assigned to a connection, it generally remains allocated for that connections sole use throughout the life time of that connection.

To better understand this, imagine the same train arriving at a station every T timeperiod. Then if a connection has any data to transmit, it drops its data into its assigned bucket(time-slot) and the train departs. And if the connection does not have any data to transmit, that bucket in that train goes empty. No passengers waiting in line can get on that empty bucket. If there are a large number of trains, and a large number of total buckets are going empty most of the time (although during rush hours the trains may get quite full), this is a significant wastage of bandwidth, and limits the number of connections that can be supported simultaneously. Furthermore, the number of connections can never exceed the total number of buckets on all th2 different trains (N*M). And this is the raison-d'etre for ATM.

Advent of ATM

The telecommunications companies are investigating fiber optic cross country and cross oceanic links with Gigabit/sec speeds, and would like to carry in an integrated way, both real time traffic such as voice and hi-res video which can tolerate some loss but not delay, as well as non real time traffic such as computer data and file transfer which may tolerate some delay but not loss. The problem with carrying these different characteristics of traffic on the same medium in an integrated fashion is that the peak bandwidth requirement of these traffic sources may be quite high as in high-res full motion video, but the duration for which the data is actually transmitted may be quite small. In other words, the data comes in bursts and must be transmitted at the peak rate of the burst, but the average arrival time between bursts may be quite large and randomly distributed. For such bursty connections, it would be a considerable waste of bandwidth to reserve them a bucket at their peak bandwidth rate for all times, when on the average only 1 in 10 bucket may actually carry the data. It would be nice if that bucket could be reused for another pending connection. And thus using STM mode of transfer becomes inefficient as the peak bandwidth of the link, peak transfer rate of the traffic, and overall burstiness of the traffic expressed as a ratio of peak/average, all go up. In the judgement of the industry pundits, this is definitely the indicated trend for multimedia integrated telecommunications and data communications demands of global economies in the late 90's and early 21st century.

Hence ATM is conceived. It was independently proposed by Bellcore, the research arm of AT&T in the US, and several giant telecommunications companies in Europe, which is why there may be two possible standards in the future. The main idea here was to say, instead of always identifying a connection by the bucket number, just carry the connection identifier along with the data in any bucket, and keep the size of the bucket small so that if any one bucket got dropped enroute due to congestion, not too much data would get lost, and in some cases could easily be recovered. And this sounded very much like packet switching, so they called it "Fast packet switching with short fixed length packets." And the fixed size of the packets arose out of hidden motivation from the telecommunications companies to sustain the same transmitted voice quality as in STM networks, but in the presence of some lost packets on ATM networks.

Thus two end points in an ATM network are associated with each other via an identifier called the "Virtual Circuit Identifier" (VCI label) instead of by a time-slot or bucket number as in a STM network. The VCI is carried in the header portion of the fast packet. The fast packet itself is carried in the same type of bucket as before, but there is no label or designation for the bucket anymore. The terms fast packet, cell, and bucket are used interchangeably in ATM literature and refer to the same thing.

Statistical Multiplexing

Fast packet switching is attempting to solve the unused bucket problem of STM by statistically multiplexing several connections on the same link based on their traffic characteristics. In other words, if a large number of connections are very bursty (i.e. their peak/average ratio is 10:1 or higher), then all of them may be assigned to the same link in the hope that statistically they will not all burst at the same time. And if some of them do burst simultaneously, that that there is sufficient elasticity that the burst can be buffered up and put in subsequently available free buckets. This is called statistical multiplexing, and it allows the sum of the peak bandwidth requirement of all connections on a link to even exceed the aggregate available bandwidth of the link under certain conditions of discipline. This was impossible on an STM network, and it is the main distinction of an ATM network.

The ATM discipline and future challenges

The disciplines conditions under which statistical multiplexing can work efficiently in an ATM network are an active area of research and experimentation in both academia and industry. It has also been a prodigious source of technical publications and considerable speculations. Telecommunications companies in the US, Europe, and Japan as well as several research organizations and standards committees are actively investigating how BEST to do statistical multiplexing in such a way that the link bandwidth in an ATM network is utilized efficiently, and the quality of service requirements of delay and loss for different types of real time and non real time as well as bursty and continuous traffics are also satisfied during periods of congestion. The reason why this problem is so challenging is that if peak bandwidth requirement of every connection is allocated to it, then ATM just degenerates into STM and no statistical advantage is gained from the anticipated bursty nature of many of the future broadband integrated traffic profiles.

Thus the past few years publications in "IEEE Journal of Selected Areas in Communications" and the "IEEE Network and Communications Magazines" are filled with topics of resource allocation in broadband networks, policing metering and shaping misbehaving traffic and congestion avoidance and control in ATM networks, and last but not least, multitudinous mathematical models and classifications speculating what the broadband integrated traffic of the future might actually look like, and how it might be managed effectively in a statistics based nondeterministic traffic transportation system such as an ATM network. The more adventurous readers desirous of learning more about ATM networks are encouraged to seek out these and the standards committees' publications.

Fortunately however, these are problems that the service providers and ATM vendors like the telecommunications companies have to solve, and not the users. The users basically get access to the ATM network through well defined and well controlled interfaces called "User Network Interface" (UNI), and basically pump data into the network based on certain agreed upon requirements that they specify to the network at connection setup time. The network will then try to ensure that the connection stays within those requirements and that the quality of service parameters for that connection remain satisfied for the entire duration of the connection.

Who are the standards bodies investigating ATM?

In the US, ATM is being supported and investigated by T1S1 subcommittee (ANSI sponsored). In Europe, it is being supported and investigated by ETSI. There are minor differences between the two proposed standards, but may converge into one common standard, unless telecommunications companies in Europe and America insist on having two standards so that they can have the pleasure of supporting both to inter-operate. The differences however are minor and do not impact the concepts discussed here. The international standards organization CCITT has also dedicated a study group XVIII to Broadband ISDN with the objective of merging differences and coming up with a single global worldwide standard for user interfaces to Broadband networks. No conclusions yet.

Types of User Network Interfaces (UNI) for ATM

It is envisioned that the ATM network service providers may offer several types of interfaces to their networks. One interface that is likely to be popular with companies that build routers and bridges for local area networks is a Frame based interface. One or more of the IEEE 802.X or FDDI frames may be supported at the UNI, with frame to ATM cell conversion and reassembly being done inside the UNI at the source and destination end points respectively. Thus a gateway host on a local area network might directly connect its ethernet, token ring, fddi, or other LAN/MAN interface to the UNI, and thus bridge two widely separated LANs with an ATM backbone network. This will preserve the existing investment in these standards and equipments, and enable a gradual transition of the ATM networks into the market place.

An alternate interface likely to be more popular in the longer runs, and for which the concept of Broadband-ISDN really makes sense, is direct interface at the UNI with standard ATM cells. Such a streaming interface can hook subscriber telecom, datacom, and computer equipment directly to the network, and allow orders of magnitude greater performance and bandwidth utilization for integrated multimedia traffic of the future. Thus it is by no accident that the IEEE 802.6 packet for the MAC layer of the Metropolitan Area Network (MAN) DQDB protocol (Distributed Queue Dual Bus) looks very much like an ATM cell.

It is quite likely that companies may crop up (if they have not already done so) to design ATM multiplexers for interface to the UNI of a larger ATM backbone network. Especially if the CCITT succeeds in standardizing an interface definition for UNI, it will be an additional boon to this market. The multiplexers with multiple taps on the user side can connect to one fat ATM pipe at the network side. Such a multiplexer would hide the details of ATM network interface from the user, and provide simple, easy to use, low cost ATM cell taps to hook the user equipment into.

Companies with investment in existing STM networks such as T1 and T3 backbones, are likely to want a direct T3 interface to the UNI, thus allowing them to slowly integrate the newer ATM technology into their existing one. Thus it is possible to see a flurry of small startups in the future rushing to make large T3 multiplexers for connecting several T3 pipes into one large ATM pipe at the UNI.

Typically, an ATM network will require a network management agent or proxy to be running at every UNI which can communicate and exchange administrative messages with the user attachments at the UNI for connection setup, tear down, and flow control of the payload using some standard signalling protocol. A direct user attachment at the UNI is likely to cost more and be more complex, then a user attachment to something which in turns interfaces to the UNI.

What does an ATM packet look like

An ATM cell or packet as specified by T1S1 sub-committee is 53 bytes.5 bytes comprise the header, and 48 bytes are payload. The header and payload are specified as follows:

```
<------------------- 5 bytes ------------------->|<---------- 48 bytes --------->|
-----------------------------------------------------------------------------------------
| VCI Label | control | header checksum | optional adaptation | payload |
|  24 bits    | 8 bits  |        8 bits            |     layer 4 bytes     | 44 or 48 |
-----------------------------------------------------------------------------------------
```

The 48 bytes of payload may optionally contain a 4 byte ATM adaptation layer and 44 bytes of actual data, or all 48 bytes may be data, based on a bit in the control field of the header. This enables fragmentation and reassembly of cells into larger packets at the source and destination respectively. (Since the header definition may still be in flux, it is possible that presense or absence of an adaptation layer information may not be explicitly indicated with a bit in the header, but rather implicitly derived from the VCI label). The control field may also contain a bit to specify whether this is a flow control cell or an ordinary cell, an advisory bit to indicate whether this cell is dropable in the face of congestion in the network or not, etc.

The ETSI definition of an ATM cell is similar, 53 bytes cell size, 5 byte header, 48 bytes data. However the difference is in number of bits for the VCI field, number of bits in the header checksum, and semantics and number of some of the control bits.

For a more detailed specification of the ATM header, see the appropriate standards committees' documents.

Connections on an ATM network

As in STM networks, where a datum may undergo a time-slot-interchange between two intermediate nodes in a route, the VCI label in an ATM cell may also undergo a VCI label interchange at intermediate nodes in the route. Otherwise, the connections in the ATM network look remarkably similar to STM networks.

An Example:

Assume an ATM network with nodes in NY, ATLANTA, DALLAS, and SF. Say that Chuck while vacationing in NY decides to play Aviator with his buddies in Mtn view who are still grinding away on MPsniff. Also assume that we have ATM cell interfaces at UNI's in both NY and SF. This is what can happen: Chuck's portable $3K laptop makes a connection request to the UNI in NY. After an exchange of connection parameters between his laptop and the UNI (such as destination, traffic type, peak and average bandwidth requirement, delay and cell loss requirement, how much money he has got left to spend, etc}, the UNI forwards the request to the network. The software running on the network computes a route based on the cost function specified by Chuck, and figures out which links on each leg of the route can best support the requested quality of service and bandwidth. Then it sends a connection setup request to all the nodes in the path enroute to the destination node in SF.

Let's say that the route selected was NY--AT--DA--SF. Each of the four nodes might pick an unused VCI label on their respective nodes and reserve it for the connection in the connection lookup tables inside their respective switches. Say, NY picks VC1. It will send it to AT. AT in turn picks VC2, associates it with VC1 in its connection table, and forwards VC2 to DA. DA picks VC3 and associates it with VC2 in its connection tables and forwards VC3 to SF. SF picks VC4 and associates it with VC3 in its connection tables, and pings the addressed UNI to see if it would accept this connection request. Fortunately, the UNI finds Chuck's buddies and returns affirmative. So SF hands the UNI and Chuck's friends VC4 as a connection identifier for this connection. SF then acks back to DA. DA acks back to AT and sends it VC3. AT puts VC3 in its connection tables to identify the path going in the reverse direction, and acks to NY sending it VC2. NY associates VC2 in its connection tables with VC1, and acks the originating UNI with VC1. The UNI hands chuck's laptop VC1 and connection is established.

Chuck identifies the connection with VCI label VC1, and his buddies identify the connection with VCI label VC4. The labels get translated at each node to the next outgoing label like so:

NY AT DA SF

Chuck -> VC1 -> VC2 -> VC3 -> VC4 -> buddies

Chuck <- VC1 <- VC2 <- VC3 <- VC4 <- buddies

Other scenarios are also possible and would depend on a vendor's implementation of the ATM network.

When Chuck has had enough playing Aviator and wants to get back to some serious scuba diving off the Atlantic coast, the connection is torn down, and the VCI labels are resued for other connections.

What Assumptions can a user attachment make for a VCI label?

As is probably obvious from the above example, none. The VCI labels are owned by network nodes, and get randomized quite quickly as connections come and go. A VCI label is handed to a user attachment only as an opaque cookie, and not much can be assumed about its spatial distribution other than quite random.

It may be possible to have certain reserved VCI labels similar in concept to "well known port definitions of UDP and TCP", as identifiers for special well known services that may be provided by the network. However very little can be assumed about the dynamically assigned VCI labels for most user related connections.

A service provider is unlikely to accede to any special request by any one service requester to allocate it a chunk of VCI labels, unless the network itself is owned by the service requester. Furthermore, the address space of the VCI labels is limited to 24 bits and only designed to identify the connections between two points on a single link. The address space would disappear rather quickly if customers started to requisition portions of the VCI label for their own semantics.

If there is a specific need to assume semantics for the VCI label outside of the ATM network, i.e. require it to be within a certain range on the user attachments at the UNI, it is probably best to provide a lookup table in hardware inside the user attachments which can map the pretty much randomized VCI label assigned by the network to n bits of a new label to which the user attachment can assign its own semantics to its silicon's content.

What Protocol layer is ATM?

As is probably evident by now, ATM is designed for switching short fixed length packets in hardware over Gigabit/sec links across very large distances. Thus its place in the protocol stack concept is somewhere around the data link layer. However it does not cleanly fit in to the abstract layered model, because within the ATM network itself, end-to-end connection, flow control, and routing are all done at the ATM cell level. So there are a few aspects of traditional higher layer functions present in it. In the OSI reference model, it would be considered layer 2 (where layer 1 is the physical layer and layer 2 is the datalink layer in the internet protocol stack). But it is not very important to assign a clean layer name to ATM, so long as it is recognized that it is a hardware implemented packet switched protocol using 53 byte fixed length packets.

What is perhaps more relevant is how will all this interact with current TCP/IP and IP networks in general, and with applications which want to talk ATM directly in particular. A convenient model for an ATM interface is to consider it another communications port in the system. Thus from a system software point of view, it can be treated like any other data link layer port. Thus for instance, in IP networks connected via gateways to ATM backbones, the model would be no different then it presently is for a virtual circuit connection carried over an STM link except that an IP packet over an ATM network would get fragmented into cells at the transmitting UNI, and reassembled into the IP packet at the destination UNI. Thus a typical protocol stack might look like this:

Data

TCP

IP

ATM Adaptation Layer

ATM Datalink layer

Physical Layer (SONET STS-3c STS-12 STS-48)

Thus, just like an ethernet port on a host is assigned an IP address, the ATM port may also be assigned an IP address. Thus the IP software in a router decides which port to send a packet to base on the IP address, and hands the packet to the port. The port then does the right thing with it. For an ethernet port, the ethernet header is tacked on and the Frame transmitted in ethernet style. Similarly, for an ATM port, the IP datagram is fragmented into cells for which an ATM adaptation layer is specified in the standards. The fragmentation and reassembly is done in hardware on the sending and receiving sides. A VCI label acquired via an initial one time connection establishment phase is placed in the header of each cell, and the cells are drained down the fat ATM datalink layer pipe. On the receiving side, the cells are reassembled in hardware using the ATM adaptation layer, and the original IP packet is reformulated and handed to the receiving host on the UNI. The adaptation layer is not a separate header, but is actually carried in the payload section of the ATM cell as discussed earlier.

For direct interface to an ATM cell stream from an application, new interfaces have to be designed in the software that can provide the application with nice and fast mechanisms for connection establishment, data transfer, keep alive, tear down, and even application level flow control. In this case the software processing steps may look like this:

Application Streaming Data

OS interface to application

ATM virtual circuit management/signalling

Driver interface to ATM

ATM

----

where the ATM virtual circuit management represents software which understands the ATM header specifics, sets up and tears down connections, does demultiplexing of the payload to appropriate connections, and responds to whatever standard signalling protocol is employed by the ATM interface at the UNI for connection management.

The Physical Layer

The physical layer specification is not explicitly a part of the ATM definition, but is being considered by the same subcommittees. T1S1 has standardized on SONET as the preferred physical layer, and the STS classifications refer to the speeds of the SONET link. STS-3c is 155.5 Mbit/sec. STS-12 is 622 Mbit/sec, and STS-48 is 2.4 Gbit/sec. The SONET physical layer specifications chalk out a world wide digital telecommunications network hierarchy which is internationally known as the Synchronous Digital Hierarchy (SDH). It standardizes transmission around the bit rate of 51.84 Mbit/sec which is also called STS-1, and multiples of this bit rates comprise higher bit rate streams. Thus STS-3 is 3 times STS-1, STS-12 is 12 times STS-1, and so on. STS-3c is of particular interest as this is the lowest bit rate expected to carry the ATM traffic, and is also referred to as STM-1 (Synchronous Transport Module-Level 1). The term SONET stands for Synchronous Optical Network and is the US terminology for SDH (since they had to differ in something). So much for the acronym soup.

The SDH specifies how payload data is framed and transported synchronously across fiber optic transmission links without requiring all the links and nodes to have the same synchronized clock for data transmission and recovery (i.e. both the clock frequency and phase are allowed to have variations, or be plesiochronous). The intention being that products from multiple vendors across geographical and administrative boundaries should be able to plug and play in a standard way and the Broadband ISDN network be a true international network. And guess what the fundamental clock frequency is around which the SDH or SONET framing is done? You guessed it, 8 KHz or 125 usec.

However all of this sits below the ATM layer and the ATM cells are transported across the physical layer as opaque payload, also called the SONET payload or the Synchronous Payload Envelope (SPE). The physical layer is independent of the payload type, and can just as easily carry STM cells as ATM cells. Refer to the standards documents for more details.

Flow control in ATM

Unlike the reactive end to end flow control mechanisms of TCP in internetworking, the gigabits/sec capacity of ATM network generates a different set of requirements for flow control. If flow control was left on end to end feedback, then by the time the flow control message was received at the source, the source would have already transmitted over several Mbytes of data into the ATM pipe exacerbating the congestion. And by the time the source reacted to the flow control message, the congestion condition might have disappeared altogether unnecessarily quenching the source. The time constant of end to end feedback in ATM networks (actually feedback_delay * link_bandwidth product) may be so large that solely relying on the user attachments to keep up with the dynamic network is impractical. The congestion conditions in ATM networks are expected to be extremely dynamic requiring fast hardware mechanisms for relaxing the network to steady state, and necessitating the network itself to be actively involved in quickly achieving this steady state. Thus a simplistic approach of end to end closed loop reactive control to congestion conditions is not considered sufficient for ATM networks.

The present consensus among the researchers in this field is to use a holistic approach to flow control. They recommend employing a collection of flow control schemes along with proper resource allocation and dimensioning of the networks to altogether try and avoid congestion, to try and detect congestion build up early by closely monitoring the internal queues inside the ATM switches and reacting gradually as the queues reach different high watermarks, and to try and control the injection of the connection data into the network at the UNI such that its rate of injection is modulated and metered there first before having to go to the user attachement for a more drastic source quenching. The concept is to exercise flow control in hardware very quickly, gradually, and in anticipation rather than in desperation. Rate based schemes which inject a controlled amount of data at a specified rate that is agreed upon at connection setup time, and automatically modulate the rate based on the past history of the connection as well as the present congestion state of the network have seen much press lately. The network state may be communicated to the UNI by the network very quickly by generating a flow control cell whenever a cell is to be dropped on some node due to congestion (i.e. the queues are getting full). The UNI may then police the connection by changing its injection rate, or notify the user attachment for source quenching depending on the severity level of the congestion condition.

The major challenge during flow control is to try and only affect those connection streams which are responsible for causing the congestion, and not affect other streams which are well behaved. And at the same time, allow a connection stream to utilize as much bandwidth as it needs if there is no congestion. This topic is still an area of active research, experimentation, and prolific publications including several PhD theses.

Does an ATM network provide inorder delivery?

Yes. An ATM cell may encounter congestion and suffer variable delay due to bufferring within the ATM switches, and may even be dropped either due to congestion control or due to header checksum error. However an ATM connection always obeys causality, the cells in a connection (i.e. cells with the same VCI label) arrive inorder at the destination. This is so because there is no store and forwarding in the network, cells travel over a single virtual circuit path, the ATM switches do not switch the cells in the same VCI out of order, and no retransmissions is done at any point in the ATM network.

Connectionless services are also supported on ATM networks, but these are implemented as a higher layer service layered over the ATM datalink layer. Thus cells in a connectionless service may arrive out-of-order because there might be multiple VCIs over multiple paths setup to deliver the connectionless datagrams and cells may arrive over different paths in different order. Thus the fragmentation reassembly engine which implements the connectionless datagrams, and which is layered on top of the basic connection oriented service of the ATM layer, must carry sequence numbers in the adaptation layer in each cell and correct any reordering of the cells at reassembly time. This is what the IEEE 802.6 protocol for MAN does to support its connectionless service class.

Does an ATM network provide reliable delivery?

No. There is no end-to-end reliable delivery service at the ATM layer. The ATM layer does not do any retransmissions and there are no end-to-end acknowledgements for what has been received. Reliable delivery service can be implemented as a layer on top of the basic connection oriented ATM layer, where acknowledgement of received data and retransmission of missing data can be done for connections requiring reliable delivery. Thus a TCP type transport layer protocol (layer 4 in the OSI model) layered on top of the ATM layer is required for guaranteed delivery.

Performance of an ATM interface

Unlike STM networks, ATM networks must rely on considerable user supplied information for the traffic profile in order to provide the connection with the desired service quality. There are some sources of traffic which are easier to describe than others, and herein lies the cost/performance challenge for best bandwidth utilization in an ATM interface.

An ATM network can support many types of services. Connection oriented as well as connection less. It can support services which may fall in any of the four categories (loss sensitive, delay sensitive), (loss insensitive, delay sensitive), (loss sensitive, delay insensitive), and (loss insensitive, delay insensitive). It can further reserve and allocate a fixed bandwidth for a connection carrying a continuous bit stream for isochronous traffic (repeating in time such as 8khz voice samples), allocate a bandwidth range for a variable bit stream for plesiochronous traffic (variable frequency such as interactive compressed video), as well as allocate no specific amount of bandwidth and rely on statistical sharing among bursty sources. It may also provide multiple priorities in any of the above categories. The services can span the entire gamut from interactive such as telephony and on-line data retrieval, to distributed such as video and stereo Hi-Fi broadcasts and multicasts for conferencing and database updates.

Thus the performance that one might get from ones ATM connection is very much dependent on the parameters that are specified at connection setup time. Just because the link bandwidth may be an STS-12, does not necessarily imply that the end to end payload bandwidth that the ATM interface can sustain will also be STS-12. It will in fact be considerably lower based on connection setup parameters and the quality of service request, and whether bandwidth was reserved or statistically multiplexed, and the load on the ATM network.

Typically, the ATM network may not permit 100% loading of any link bandwidth, and in fact user available bandwidth may not be allowed to exceed more than 80% of the peak bandwidth of the link. The UNI may start policing and/or denying new connection requests on the link if utilization exceeds this amount. Add the approx 10% overhead of the 5 byte header in the 53 byte cell, and the max sustainable payload throughput on an ATM cell stream interface may peak at 72% of the peak link bandwidth. And this does not include any adaptation layer overhead if present, signalling overhead or physical layer overheads of SDH SONET framing and inter-cell spacing gaps.

And of course, application to application bandwidth may be even less, unless the software datapath from the interface driver through the OS to the application (and vice versa) is very carefully optimized. It would hardly be received very well if the end-to-end throughput from application to application would turn out to be no better for an ATM port than for an ethernet or fddi port due to software overheads.

How many cells might be realistically received or transmitted at a sustained rate on an ATM cell interface in a processor? Hard to say for sure as there is no existence proof as yet.

However, what can be stated is that the transmitter and receiver performance is independent of each other. The transmitter side is constrained by the flow control of the simultaneous connection streams by pacing the injection rate according to the respective negotiated class of service and bandwidth requirements. The receiver side is constrained by asynchronous reception of cells at a variable rate, and with bufferring capacity for a large number of simultaneous connections each of which can be receiving data simultaneously. And if an adaptation layer is used, then the reassembly of these cells into a higher layer protocol data unit (PDU) must also be done in hardware by the receiver side. Thus a lot of thought is required in designing an ATM interface to a host system, poor design of which can cripple the system performance.

When can I have my own connection to an ATM network?

The Broadband ISDN with ATM is an enabling technology. It will enable new kinds of applications and new types of usage which are only in people's imagination today. It is a complete overhaul of the copper based low bandwidth telecommunications technology that has existed until now, and represents a massive investment both in researchand development, as well as deployment and integration. The software investment required to make the ATM network work is tremendous, and many of the algorithms and theories about how to manage the ATM network are still in their infancy and mostly on paper. Considerable work is also required in developing new network management paradigms and protocols to effectively control and manage the vasts quantities of bandwidths and services that the revolution in communication technology is promising to offer.

At the present time, there are no commercially available ATM networks in the US (to my knowledge), though there are several ATM prototype switches and experiments in existence. The earliest anticipated roll out of commercial ATM switches is expected no sooner than 1995 time frame. And full fledged deployment of ATM networks with "COST

EFFECTIVE" multi-media integrated services to end-users is still a lot farther away, probably closer to the end of this decade. But it's coming...hang on.

Conclusion

ATM is a scalable, flexible and "future-proof" technology that allows for transportation of various forms of information, including data, voice, video or multiples of them. Global standards are being introduced to ensure compatibility. It is primarily for these reasons that ATM was chosen as the method for the implementation of broadband services.

References

The following is a list of references and articles for further reading.

- Martin de Prycker, "Asynchronous Transfer Mode: Solution for Broadband ISDN", Ellis Horwood Limited, 1991.
- Rainer Haendel, Manfred N. Huber, "Integrated Broadband Networks: An introduction to ATM-Based Networks", Addison-Wesley Publishers Ltd., 1991.
- IEEE Communications Magazine, September 1991, Vol. 29, No. 9.
- The ATM Forum, "ATM User-Network Specification", Version 3.0, September 10, 1993.
- Jean-Yves Le Boudec, "The Asynchronous Transfer Mode: a tutorial", Computer Networks and ISDN Systems 24(1992), pp 279-309.
- Craig Partridge, "Gigabit Networking", Addison-Wesley Publishers Ltd., 1994.
- L.G. Cuthbert, J.C. Sapanel, "ATM: The broadband telecommunications solution", London Institution of Electronic Engineers, 1993.
- S. Rao, "Internetworking in broadband networks: heterogeneity and global networking", Oxford IOS, 1993.