ATM at the Desktop

A Technology White Paper

Author, Martin Taylor

April 1995

Executive Summary

This white paper is aimed at network planners and managers who are looking at the deployment of ATM networking technology in their LANs to support evolving business needs that demand new kinds of networked applications. The paper describes the characteristics of ATM networking technology in outline, discusses the application of ATM to support networked multimedia and bandwidth-intensive applications at the desktop, and describes the networking hardware and software elements that provide for practical implementation of ATM at the desktop within an existing LAN environment. Finally, the paper looks at the wider implications of ATM for LAN backbone and wide area networking, and discusses alternative LAN technologies for supporting multimedia applications at the desktop.

About the Author

Martin Taylor is Vice President, Network Architecture for Madge Networks, Inc. He is responsible for the strategic planning of the company's products, including LAN switching hubs, ATM adapter cards and LAN-ATM access products. Martin joined Madge Networks in 1991, and formerly held the position of Director of Product Marketing. Prior to joining Madge, Martin held business development positions in Local Area Networks and in fiber optic cabling systems with GPT Ltd, the UK's leading supplier of telecommunications equipment. Martin has a total of 9 years experience in the communications industry, and a further 6 years experience in MIS. He can be reached on email at mtaylor@madge.com.

Contents

A Better Way to Network

This paper is about the deployment of ATM networking at the desktop. It addresses a wide range of issues, from the applications that benefit from ATM as a desktop technology to the practicalities of implementing ATM at the desktop in an existing LAN environment. It does *not* set out to provide a comprehensive tutorial on ATM. There have been plenty of other materials published in which the reader can find detailed information on all aspects of the workings of ATM. However, a brief description of ATM networking is provided in Appendix A to assist readers unfamiliar with the technology.

ATM will be important at the desktop for two reasons. First, ATM offers some unique features which have great value in certain kinds of networked applications - notably those that combine data communication with video and voice, often referred to as multimedia applications.

Secondly, and far more importantly, ATM is simply a better way to network. ATM network users get dedicated bandwidth, which means no more performance degradation under heavy load. ATM networks are constructed with a single building block, the ATM switch, eliminating the need for separate routing and hubbing functions, and providing far greater capacity per dollar invested. And ATM networks provide the scalability, and the ability to support all imaginable future applications, which are crucially lacking in the LAN technologies that we depend upon today.

Application Needs for Desktop ATM

ATM offers such compelling advantages that it is likely to overshadow all other networking technologies at some time in the future. But in the short term, ATM may be costlier than current networks, and, until the required standards are all in place, may face some interoperability issues. For these reasons, immediate interest in ATM is likely to focus on those applications that derive particular advantage from the technology.

We will look at two classes of desktop applications that benefit very specifically from ATM's capabilities: networked multimedia, and applications requiring very fast file transfer, such as high resolution image processing. There are other classes of desktop applications that benefit from ATM, such as distributed supercomputing, but these are of interest only to a highly specialized technical computing audience.

Delay-sensitive Traffic - Networked Voice and Video

Network users are showing growing interest in a variety of applications that require the network to deliver video, voice and data to the desktop. These applications include delivery of multimedia training materials over the network, dissemination of news and other information by means of live or stored video feeds, and desktop-to-desktop conferencing with voice and video, together with data and application sharing.

Latency and Delay Variation in Networks

To date, few networked multimedia applications have been implemented on a large scale, and one of the key reasons for this lies in the shortcomings of current LAN and internetworking technology. Today's LAN internetworks cannot provide either the bandwidth or the delivery of voice and video streams with low delay that is essential for acceptable quality of presentation.

Both delay and delay variation in LAN internetworks arise from the transmission of data in large packets and the use of store-and-forward devices such as bridges and routers to link network segments together. At the originating station, a delay is incurred while a stream of data is accumulated to form a complete packet. This delay is then exacerbated at each subsequent hop through the network as each bridge or router has to receive the entire packet into memory before commencing transmission on the next link. Delay variation occurs because packets have to queue for transmission on busy links. Since all packets are treated with the same priority, time-sensitive packets containing video data can easily be stuck in a queue behind batch file transfer packets.

Users who have experimented with networked multimedia typically report that applications can be run successfully over a single LAN segment provided that the number of concurrent users is strictly limited, but that the applications do not run satisfactorily over bridges or routers.

ATM Quality of Service

ATM provides a radically better solution to networked multimedia than LAN internetworks. ATM's cell-based transmissions are inherently less subject to store-and-forward delay, and queuing delay is addressed by the provision of cell prioritization, which allows time-sensitive traffic to jump the queue. And on top of this, ATM permits a station to request the allocation of bandwidth on a link for time-sensitive traffic, thereby guaranteeing that there is sufficient capacity for the required quality of service.

ATM's bandwidth allocation mechanisms are implemented by classifying traffic as Constant Bit Rate (CBR), Variable Bit Rate (VBR) or Available Bit Rate (ABR). A station wishing to send or receive video or voice traffic via a CBR or VBR connection must request the bandwidth required when placing the call. If the network is unable to find the requested bandwidth, the call fails with a "busy tone". Otherwise the bandwidth required is reserved for this channel. Data traffic is normally sent over an ABR connection, and uses as much bandwidth as is instantaneously available in the network. If stations attempt to send traffic at a greater rate than the available bandwidth, then flow control mechanisms are brought into play and the sending station(s) are requested to slow down transmissions.

Video Compression

Raw digitized broadcast quality video requires over 100 Mbps for transmission, so it is normal to apply compression to video before sending it over a digital network.

A variety of video compression schemes are available, with different characteristics that suit different degrees of compression, and different applications such as point-to-point real time or stored video playback.

Compression Scheme	Intended For	Bit Rate	Compression Processing Requirement	De-compress Processing Requirement	Comments
H.261	Wide area video conferencing	64 kbps - 2 Mbps	High	High	Designed for synchronous leased lines and ISDN
MPEG-2	Video distribution	1.5 Mbps - 15 Mbps	Medium-High	Low	Provides VCR and broadcast quality
Motion JPEG	General purpose video	1.5 Mbps - 15 Mbps	Medium	Medium	Symmetric processing reqmt suits video conferencing
Indeo	LAN-based video playback	1.2 Mbps	High	Low	Proprietary Intel scheme based on i750 processor

For video applications in a business context, typical desktop bandwidth needs will range from 128 kbps for video conference sessions that are making use of ISDN connections in the WAN up to perhaps 4 Mbps for VCR quality video playback from a local video server. Thus multimedia applications do not need very high speed connections to the desktop, they simply need connections that deliver network traffic with low delay and jitter.

Voice over ATM

The ability to handle voice transmission over ATM to the desktop is particularly interesting because it has the potential to displace the traditional telephone handset and PBX network. In the short term, deployment of desktop ATM for voice communications is likely to be limited to specific applications where voice and data integration is important, for example Automated Call Distribution in customer service centers, and Interactive Voice Response systems. Longer term, the PC may become the voice terminal for general purpose desktops.

In this scenario, PCs equipped for sound and connected to a desktop ATM network would work together to provide distributed PBX functionality. Call processing and messaging functions would be provided in software running on low-cost commodity server platforms. The PC's graphical user interface would provide point-and-click control over messaging and call processing functions, which for many users will be far easier to handle than phone keypad-activated controls. Voice, email, fax and video messaging would all be integrated as a single PC-based application. And all person-to-person communications across the network, from a simple point-to-point telephone call to a multi-point videoconference with application sharing, would make use of the same network infrastructure and the same control semantics.

In PBX networks, a voice call is handled as a synchronous 64 kbps traffic stream. Such a traffic stream can be carried in an ATM network as a stream of ATM cells. Each cell holds 48 bytes of information, equivalent to 384 bits, so we need to send 167 cells per second to provide 64 kbps bandwidth.

One important issue faced by voice over ATM is *packetization delay*. Because we have to accumulate 48 bytes of voice data before we can send a fully-loaded cell, the traffic stream is subjected to a delay equivalent to the time it takes to accumulate this amount of data. With a 64 kbps voice stream, the packetization delay is about 6 milliseconds. When we take this delay into account along with ATM cell switching delays and receive buffering delays, the total round trip delay may approach or exceed the CCITT recommended maximum values for acceptable speech quality. In this case additional processing may be required at each end station to provide echo cancelling.

The amount of bandwidth needed for voice is tiny in comparison to the bandwidth provided in the desktop ATM connection and in the ATM switches. We can take advantage of this to provide “hi-fi” quality voice communications over the ATM network by using a faster sampling rate to digitize the voice. Telephone quality voice has a frequency response up to 3.5 kHz. If we sample voice at 16 kHz (double the standard rate) then we can provide an analogue frequency response up to 7 kHz over a 128 kbps ATM connection, which offers a subtantial improvement in the perceived quality of the communication. At this higher speed, packetization delay is also reduced.

Voice communication systems based on ATM networking will need gateways to external networks such as the public telephone system. These gateways will be implemented in specialized kinds of communications servers. Additional “handset servers” will be needed to connect conventional handsets into the ATM network locally.

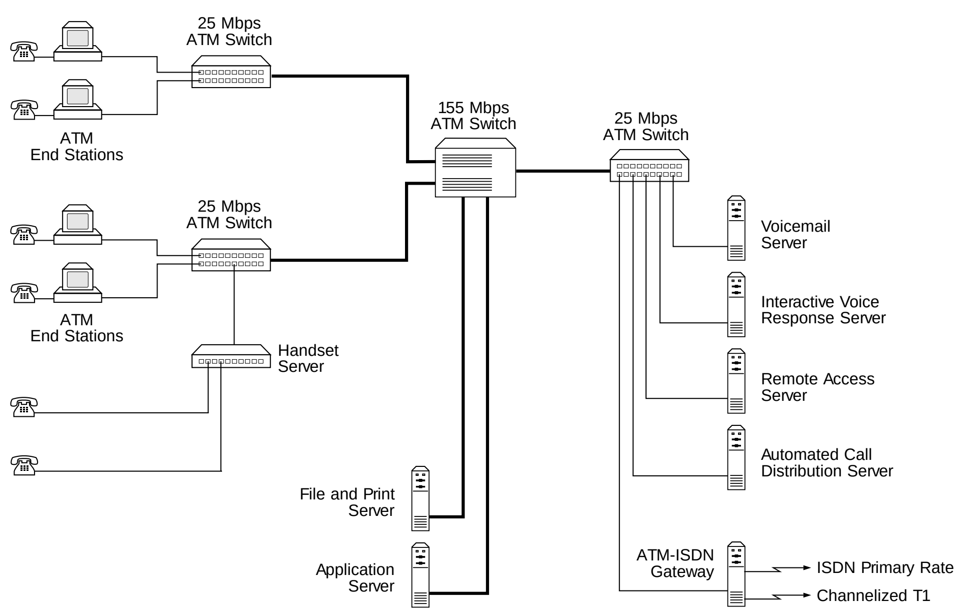

Figure 1: ATM LAN as distributed PBX with server-based voice applications

Bandwidth-intensive Applications

Certain kinds of applications require very large volumes of data to be moved on demand to and from the desktop. In this case the concern is not with the consistency of timing of data delivery, rather it is raw speed.

Applications based on high resolution image retrieval and processing are most often cited in this context. These include medical applications (X-ray images) and applications in publishing, advertising and graphic arts. Other applications that can benefit from very high speed bulk data movement include software development, CAD/CAM and engineering calculations.

ATM provides a range of link speeds that can be deployed at the desktop to meet the needs of data-hungry applications. Available speeds include 25, 51, 100, 155 and 622 Mbps.

The ability of end stations to exploit the very high speeds offered by ATM at the top end will depend on the nature of the application and the processing power available at both the client workstation and the server. High-end PCs are capable of handling memory-to-memory data transfers to and from servers at 80-90 Mbps or more, but any data transfer involving extensive reading or writing to disk storage is likely to run far slower than this. In reality, there are few PC-based applications that will benefit from connections providing much more than 25 Mbps of dedicated bandwidth.

ATM Technology for the Desktop

To support ATM networking cost-effectively at the desktop, three key elements of technology have to be in place.

- An ATM transmission scheme compatible with standard copper-based desktop cabling installations between the wiring closet and the desktop.

- A software environment for end stations that supports the use of current applications over the ATM network.

- An ATM network infrastructure which supports connectivity and switching between ATM end stations.

The standards available for ATM transmission to the desktop and support for current applications over ATM networks are described below, together with a discussion of the elements needed to build an ATM LAN infrastructure.

ATM Transmission Technologies for the Desktop

The key requirement for practical deployment of ATM at the desktop is compatibility with current standards for structured copper cabling schemes. The three transmission technologies described will all operate on UTP cable of appropriate quality, up to 100 meters from the wiring closet. All the schemes are also capable of runnning over STP cables such as IBM Type 1.

155 Mbps over Category 5

The ATM Forum standard for 155 Mbps ATM over Category 5 UTP was agreed in the summer of 1994. The scheme is based on the SONET OC-3 standard for ATM over optical fiber, which was originally developed to handle ATM on telco network trunk links. This has been adapted to run over copper, using a simple NRZI encoding scheme with very precise balancing of the transmit pairs to bring radiated emission within the allowable limits permitted by FCC regulations.

51 Mbps over Category 3

The decision to standardize 51 Mbps ATM over Category 3 UTP was made by the ATM Forum in November 1993. Like the 155 Mbps scheme, this standard uses SONET framing, the OC-1 rate giving 51 Mbps.

Transmitting at 51 Mbps over Category 3 cable while remaining within FCC limits for radiation represents a substantial technical challenge. The problem has been solved by means of a complex encoding called CAP-16, which is similar to the kinds of phase/amplitude modulation schemes used in high performance modems.

Unfortunately, it has turned out to be more difficult than expected to implement the CAP-16 encoding in silicon. The chips are still not shipping, and costs are somewhat higher than predicted. Uncertainty over the future of this technology has held most vendors back from announcing plans to support the scheme, and there are no 51 Mbps products on the horizon as yet.

25 Mbps over Category 3

The 155 Mbps and 51 Mbps standards for ATM over copper cabling have been derived by adapting telco-based transmission schemes that use SONET framing. As a result, both standards include a certain amount of "baggage" carried over from their telco origins that has little or no value on a desktop link.

An alliance has been formed between a number of vendors who have taken an alternative "bottom up" approach to the design of a desktop transmission scheme for ATM, resulting in a very simple system which delivers 25.6 Mbps ATM over Category 3 UTP cabling. This scheme is based on existing transmission technology which is widely used in Token Ring networks, and avoids the need for complex line coding logic and SONET frame processing. As a result, it exhibits costs which are a fraction of those associated with either of the two higher speed schemes for copper cabling.

In February 1995, the 25 Mbps ATM specification being promoted by the Desktop ATM25 Alliance was adopted by the ATM Forum as the baseline text for a new standard.

Choice of Speed

Some observers have questioned the value of an ATM connection to the desktop running at what is considered to be (in the ATM world) such a low speed. The usefulness of 25 Mbps ATM has to be judged against a background where the vast majority of desktop users today are connected by shared media networks at 4, 10 or 16 Mbps, on which typically 40 to 80 users are sharing access to a single half-duplex channel at these speeds. ATM at 25 Mbps gives each user a dedicated full-duplex 25 Mbps connection, which represents a huge step forward in terms of capacity per user. Furthermore, a 25 Mbps connection at the desktop will support several concurrent sessions of compressed video, and is therefore more than adequate to meet the needs of networked multimedia applications.

Only a few specialized applications, such as high resolution image processing, will benefit from access to higher speed connections such as 155 Mbps. ATM networks allow different link speeds to be mixed, so that each user can benefit from access to a link at an appropriate speed for their specific needs. The needs of the vast majority of users will be successfully met by 25 Mbps ATM links.

Supporting Desktop Applications over ATM

Having brought a physical ATM connection to the desktop at an appropriate speed, we need now to consider how desktop applications can be supported over the ATM network connection.

The principal issue to be addressed concerns the underlying differences between the operation of current LAN technologies and ATM networks. Current LAN applications are unable to operate directly over an ATM network because of these differences, so a technique known as LAN Emulation over ATM is used to provide a familiar interface to application software in the end station.

LAN Emulation over ATM

Today's LAN applications are designed to run over a shared media network such as Ethernet or Token Ring. These applications are supported by standard software interfaces to network adapter cards such as Novell's ODI or Microsoft's NDIS interface.

Software interfaces such as ODI and NDIS provide network transport services to the application software that runs in both servers and client workstations, and these services reflect the underlying nature of the LAN technology. Today's Ethernet and Token Ring LANs provide connectionless "best effort" delivery of variable length data packets which are addressed by a unique 6-byte value that identifies the physical interface to which the packet is to be sent. These LAN technologies also support broadcast and multicast packet delivery, whereby a packet with a special kind of address will be delivered to all of, or a subset of, the stations on the LAN.

The upper layers of the communications software in LAN end stations, including the client "shell" and the server operating system, have evolved entirely around this underlying functionality. For example, end stations in an IPX network rely on "find nearest server" broadcast requests to locate resources on the LAN.

The operation of an ATM network is entirely different. Connections between end stations have to be set up before any data can be transferred. Data is sent in fixed length cells rather than variable length packets, and is addressed by a virtual channel identifier rather than by physical destination identity. Finally, although ATM networks do support point-to-multipoint connections, they do not inherently provide an equivalent to the broadcast and multicast services available on LANs.

LAN Emulation over ATM provides a solution to the problem of running current LAN application software over an ATM network. This technique defines a "LAN Emulation Client" process to run in each end station, which adapts the connection-oriented cell transport service provided by ATM to the connectionless packet transport service demanded by the applications.

The LAN Emulation Client (LEC) process handles the transmission and receipt of LAN packets from the ATM network, as follows:

- When requested to transmit a packet, the LEC process resolves the packet's LAN destination addresses to an ATM station address. If no virtual channel connection is currently set up to this address, it signals to the ATM switch to request a connection. Once the connection is in place, the LAN packet is segmented into ATM cells and transmitted on the appropriate virtual channel connection.

- When requested to transmit a broadcast packet, the LEC segments the packet into cells and transmits these cells to a special station in the ATM network called the Broadcast and Unknown Server (BUS). The BUS maintains a list of stations that belong to this ATM “LAN” and is responsible for sending the cells on to all the stations on the list, so that they all receive the broadcast packet.

- On receiving a stream of ATM cells from the network, the LEC process re-assembles the incoming cells to re-create the original LAN packet, and passes this to the application software.

The LEC process communicates with the end station’s application software through one of the standard network software interfaces such as ODI or NDIS. In effect, the LEC emulates a standard ODI or NDIS network adapter card driver, fooling the end station software into thinking that it is talking to a standard Ethernet or Token Ring adapter. In practice, the LEC process will be built into the driver software that comes with the ATM adapter, so it will be a simple matter to install an ATM adapter complete with LAN Emulation to operate in a standard server or client software environment.

LAN Emulation over ATM is currently in the process of standardization by the ATM Forum. The standard describes the address resolution protocols, the operation of the Broadcast and Unknown server, and the packet segmentation format so that multi-vendor interoperability can be achieved in an emulated LAN environment. A detailed explanation of the ATM Forum’s LAN Emulation scheme can be found in “LAN Emulation over ATM — a Technology White Paper” published by Madge Networks, November 1994.

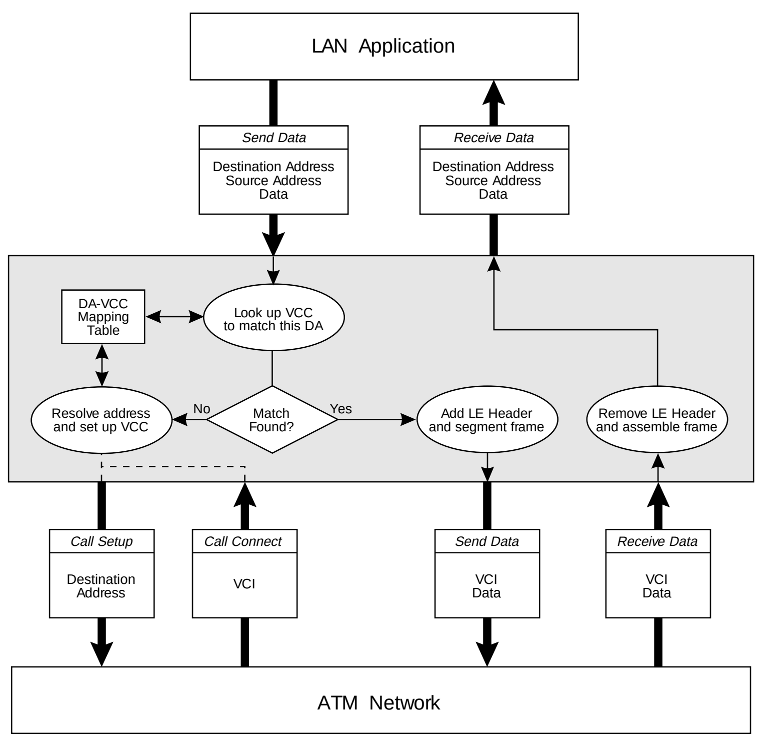

Figure 2:
LAN Emulation Client process adapting current LAN applications to ATM networking

Native ATM Support

While LAN Emulation over ATM provides an effective solution to the problem of operating current LAN applications over an ATM network, it is often seen as an interim step which bridges the gap until applications become available that are specifically designed to run over ATM.

Interest in application support for "native ATM" is motivated by two observations. First, that LAN Emulation adds considerable complexity to the network, requiring additional processing in end stations and additional processes in the network for address resolution and broadcast forwarding. And secondly, that LAN Emulation does not permit applications to take advantage of some of the unique features of ATM, including the ability to specify quality of service, and to reserve bandwidth across the network for constant bit rate traffic such as video.

Adapting communications processes, such as the interaction between clients and file servers, to operate directly over ATM requires a number of complex issues to be addressed at various levels in the protocol stack. Currrent transport and network layer protocols, such as TCP/IP, are not necessarily well tuned to the underlying characteristics of ATM networks. Ideally, new kinds of transport protocols are needed. And upper layer processes that depend on broadcasts need to be re-designed around the point-to-point nature of ATM communications.

Some networking software vendors have recognized this, and have announced plans to enhance the protocols to provide improved ATM support. For example, Novell has announced that it is working on Connection-Oriented IPX (CO-IPX) to support NetWare over ATM networks. But CO-IPX is still some way from reality, and until it becomes available, NetWare users will look to LAN Emulation over ATM to support their applications.

Video applications over ATM represent a far more pressing need for native ATM support than traditional LAN applications like file and print services. A video application that is designed to run over an Ethernet or Token Ring LAN can obviously be supported by LAN Emulation over ATM, but some of the specific advantages of ATM for handling video traffic will be lost.

Video over Native ATM

There are two inherent capabilities of ATM networking that are very valuable for handling video traffic.

First is the ability of ATM to distinguish between time-sensitive traffic like video and other types of data transfer. Connections in an ATM network can be classified as Constant Bit Rate (CBR), Variable Bit Rate (VBR) or Available Bit Rate (ABR). When setting up a CBR or VBR connection, bandwidth is reserved on all the links along the route to make sure that capacity is always available to carry the traffic. Traffic on ABR connections makes use of whatever bandwidth is instantaneously available that is not being used by CBR and VBR traffic, which always takes precedence when a link becomes congested. Video is clearly best transmitted on CBR or VBR connections, while LAN Emulation makes use of ABR connections.

And secondly, ATM switches provide efficient hardware-based support for point-to-multipoint virtual channel connections which are ideal for supporting video broadcasting and multicasting over the network. Using point-to-multipoint connections allows a number of live video feeds to be connected to the network, and stations on the network can elect to join one or more point-to-multipoint groups to "tune into" specific video feeds.

In a LAN Emulation environment, all broadcasts and multicasts must be handled by the software-based Broadcast and Unknown Server. This is necessary because LAN Emulation requires support for any-to-any broadcasting rather than one-to-many. Video broadcasting via a BUS would be very inefficient, and the BUS itself would represent a potentially serious bottleneck.

In order to take advantage of ATM's capabilities for guaranteeing the quality of video transmissions using CBR or VBR connections, and for handling video distribution by switched point-to-multipoint virtual channel connections, we need a protocol stack that supports video directly over ATM.

Work is taking place in the ATM Forum to address the standardization of video transport over ATM. However, because the Forum is addressing a variety of different needs, including distribution of video material to residential premises, and because the technical issues are complex, this work is expected to continue for a considerable period before substantive progress is made towards a standard or set of standards.

The first implementations of video over native ATM are likely to result from co-operative efforts between applications and networking vendors. Experience gained from these kinds of proprietary implementations is likely to accelerate progress towards standards.

Infrastructure to Support Desktop ATM

ATM networks consist only of ATM-connected end stations and ATM switches. The infrastructure needed to support ATM at the desktop therefore consists simply of interconnected ATM switches. These switches are likely to vary in their physical configuration according to their precise role.

ATM Switches and Concentrators

Structured cabling schemes are generally based on cable runs from desktops to wiring closets (floor cabling), and additional cable runs from wiring closets to some central location (trunk cabling). There are generally far fewer cable runs in the trunks than on the floor, and so there has to be some kind of concentrator function in the wiring closet to enable many desktops to share the trunk cabling connections to the central resources.

In an Ethernet or Token Ring LAN, this concentrator function is performed by the wiring hub. In an ATM network, an ATM switch is needed.

Most ATM switches currently on the market are symmetrical in their design, meaning that the switch architecture is optimized for balanced traffic flows on each port. A switch which has 16 x 155 Mbps ports will generally be able to handle close to wire speed on each port, for an aggregate capacity of 2.4 Gbps. This kind of switch is suited for a core switching function in a network, where the traffic flows across the switch are typically any-to-any.

Most practical LAN configurations, though, involve communication from many clients to a much smaller number of servers. The implication of this is that an ATM switch which is located in the wiring closet to connect a number of desktops into the network will actually be required to handle many-to-few communications across its ports. Many clients will be communicating with few servers; these servers may be locally connected to the switch, or linked by other switches across the trunk cabling.

The ideal ATM switch for the wiring closet therefore has an asymmetric design, optimized to concentrate traffic from many ports to few.

Asymmetric switches for wiring closets may be true switches, which allow traffic from any port to be switched to any other port, or they may be pure concentrators, which connect all the traffic from desktop ports to trunk ports. With a pure concentrator, traffic which needs to move from one local desktop to another would have to be transmitted over the trunk link to a central switch, and then transmitted back over the trunk to the concentrator. Concentrators may provide a more economical solution than true switches in wiring closet situations where the majority of traffic flows follow a hierarchical route from desktop to center.

Capacity Provision for Desktop ATM

ATM networks provide dedicated bandwidth on each link to the desktop. In other words, end stations do not share the bandwidth on the wire that links them to the ATM switch with any other stations. However, when multiple ATM end stations are communicating with the same server then the bandwidth on that link *is* being shared. Likewise, bandwidth is shared on a trunk link between ATM switches when multiple end stations access common resources across the trunk.

With any network that provides many-to-few communications, we can only preserve dedicated bandwidth end-to-end by provisioning sufficient capacity on all links to ensure that no link is ever overloaded. Applying this logically, a server with a 155 Mbps ATM connection could only support 6 users with 25 Mbps desktop links. If the server had a 622 Mbps connection, then it could support 24 users with 25 Mbps links, or 4 users with 155 Mbps links.

Clearly this is absurd. In today’s LANs, a server connected to a 16 Mbps Token Ring backbone can support many tens or even hundreds of users — and in this situation, the server is sharing the 16 Mbps backbone with other servers. Admittedly, we may be looking for improved performance — but not to the extent where we have a 622 Mbps server connection supporting only 24 users.

Real applications in practice are not going to demand data transfer at full ATM wire speed for continuous periods. In an ATM network supporting a mix of transaction processing, client/server, file transfer and multimedia applications, there is going to be a mix of continuous traffic to support networked video and voice, and bursty traffic from the data applications. The continuous traffic is likely to occupy only a fraction of the wire speed on any link, while the bursty traffic will peak at or near wire speed, for very brief periods.

When we combine multiple traffic flows that combine continuous and bursty traffic, we can rely on the principal of statistical multiplexing to allow us to deploy links that are apparently over-committed, but which actually deliver very acceptable performance in practice. For example, we could use a wiring closet switch to connect 12 users with 25 Mbps ATM desktop links to a single 155 Mbps connection to a central server farm.

The 155 Mbps link is apparently 2:1 over-committed, but in practice this configuration would support a continuous 4 Mbps VCR quality video stream to each user, with over 100 Mbps of capacity still available to handle bursty data traffic. Because it is not likely that all the users will be sending or receiving video all the time, and because 100 Mbps is likely to provide more than enough aggregate data handling capacity for this number of users, we could probably stretch safely to a 4:1 over-commitment, and support 24 users at 25 Mbps each on a single 155 Mbps downlink.

The extent to which we can over-commit capacity on those links in the ATM network that are shared depends very much on the nature of the applications and the intensity of their use. It is difficult to give general guidelines, but the fact that the vast majority of users today are successfully sharing 4, 10 or 16 Mbps access with dozens of others is an indication that, in most cases, a high degree of over-commitment can comfortably be tolerated.

The value of over-committing capacity in an ATM network is that we can economize on costly high-capacity switch ports, and keep the overall system cost down. If we find that certain trunk links are, in practice, becoming overloaded, it is a simple matter to increase link speed or add additional links to deal with this.

Alternatives to ATM at the Desktop

Shared-media Ethernet and Token Ring LANs do a fine job of servicing the generic data applications that are used on the vast majority of PCs today. There are many who argue that Ethernet and Token Ring are more than adequate for most desktop access purposes, and that there is little practical benefit in moving to ATM at the desktop.

At this point, it is worth re-iterating some of the reasons why ATM is a better way to network:

- ATM provides a fully scalable network solution, offering a range of link speeds that address all conceivable network requirements. Thus ATM is the most future-proof of all networking technologies.
- ATM reduces network administration and management burden by simplifying moves and changes and eliminating network congestion issues.
- ATM provides efficient support for a whole new class of applications that combine data with time-sensitive voice and video traffic. Organizations that successfully exploit these kinds of applications have the potential to make major improvements in business effectiveness.

Other LAN technologies continue to make advances to meet the ATM challenge. Some of these advances are compatible with the installed base of hubs and adapter cards, while others require the renewal of large parts of the LAN infrastructure. None of them match ATM in terms of all-round capability.

Switched LANs

LAN switching provides an efficient means of interconnecting multiple LAN segments of the same technology, with good performance. Using LAN switching, LAN segments can be reduced in size so that the available bandwidth is shared by fewer stations. Today, most LAN switches are being used to implement collapsed backbones, interconnecting workgroup and server segments with 50 or more stations per segment. But typical switch costs per port may be low enough to justify reducing shared segment size to just a handful of stations, or even to provide a dedicated private "LAN per port" connection to a single station. This provides some of the dedicated bandwidth characteristics of ATM at the desktop, but with the advantage that it makes use of an existing LAN adapter in the station.

Switched Ethernet

Ethernet switching offers a logical enhancement for existing Ethernet installations. Switch costs have fallen rapidly, and Ethernet switching holds out the promise of dedicated 10 Mbps per desktop at well under $500 per connection, based on the use of existing adapter cards.

Although Ethernet switching may be the cheapest way to bring dedicated bandwidth to the desktop, it does not compare very favorably with ATM in terms of capabilities. Standard Ethernet offers only half-duplex transmission, and a single bi-directional videoconference session could eat up all the available bandwidth on the link. Full-duplex Ethernet is being developed, but is not compatible with the installed base of adapter cards. Ethernet also provides no facility for specifying Quality of Service, for prioritizing delay-sensitive traffic, or for reserving bandwidth on an end-to-end connection for video or voice communications. And finally, connecting switched Ethernet LANs together across the wide area requires the use of routers, which will introduce unacceptable delays to voice and video traffic when heavily loaded.

Ethernet switching can be used as a means of connecting existing Ethernet stations to an ATM network. In this case, the Ethernet switch is optimized to provide multi-port Ethernet to ATM conversion, usually based on LAN Emulation. The idea is to bring the power of ATM to the wiring closet, and make use of existing Ethernet adapter cards in PCs to provide the last 100 meters connection to the desktop.

Ethernet Enhancements for Multimedia

The use of Ethernet as the access connection from the desktop to the ATM network results in the loss of many of ATM's most important advantages. One vendor is attempting to improve this situation by implementing enhancements to Ethernet switching which provide for prioritization of different types of Ethernet traffic.

This scheme, proposed by 3Com in November 1994, is known as Priority Access Control Enabled (PACE). Details released so far have been sketchy, but the scheme appears to rely on the assignment of two MAC addresses to the adapter card for two different classes of traffic. The Ethernet switch prioritizes traffic on the link to the desktop according to the MAC address in each packet. The scheme is said to operate with existing adapter cards.

Many questions about PACE remain unanswered. For example, how do applications interact with the adapter card to specify which priority (and hence which MAC address) should be used? How is the second MAC address for each adapter card assigned and administered? How many other vendors will support PACE?

Switched Token Ring

As in Ethernet installations, Token Ring users are embracing the concept of LAN switching for collapsed backbones. Token Ring users currently have less incentive to reduce workgroup segment size, however, since the higher wire speed and deterministic access method of Token Ring ensures much less likelihood of segment overload.

Token Ring switching can be deployed to deliver dedicated 16 Mbps connections to each desktop. Dedicated Token Ring links provided by switches can also provide full duplex connections, by means of a proposed enhancement to the 802.5 standard for Token Ring. Some standard Token Ring adapters can be upgraded through software to support full duplex Token Ring.

Currently, the cost per port of Token Ring switching is somewhat higher than for Ethernet switching, owing to the greater complexity of the technology. Costs will come down, but the cost of dedicated Token Ring may well compare unfavorably with that of 25 Mbps ATM.

Token Ring Priority

Unlike Ethernet, Token Ring does embody packet prioritization mechanisms that enhance its suitability for transporting multimedia traffic.

In a Token Ring network, access to the ring for data transmission is normally provided in strict ring sequence by the circulation of the token. Where the priority mechanism is invoked, a station may gain privileged access to the token ahead of all other stations with lower priority data.

If we are sitting at station with a high priority packet to transmit, and there are several stations upstream of us waiting for the token so that they can transmit, then we need a way of notifying them that we have a higher priority so that we can transmit before them. In this situation, we look for the first data packet that comes past us from an upstream station, and in the header of this packet we set a special 3-bit sequence called the priority reservation bits. This packet continues on around the ring back to the station that originated it, which strips it off the ring. When this station finds the priority reservation bits in the header of the packet, it issues a new token with the requested priority level. When downstream stations examine this token, they see that the priority level is higher than their data, so they pass the token on without transmitting their data until it arrives at our station, where we are waiting to send the high priority packet. Once we have transmitted this packet, we re-issue the token with normal priority.

Tests carried out by IBM and others have shown that the use of this priority mechanism can ensure that multimedia traffic is handled with well-controlled delay even at traffic loadings exceeding 75% of network bandwidth and in the presence of bursty data traffic.

Packet prioitization can provide real advantage for supporting multimedia applications locally in existing Token Ring environment, but Token Ring still suffers from the aggregate bandwidth limitations of shared media, and lack of support for multiple priorities in the bridges and routers that are necessary to link rings together.

When Token Ring is used as the desktop network in conjunction with an ATM backbone, it should be possible to build a mapping between Token Ring packet priority and ATM Quality of Service in the switch devices that link Token Rings with the ATM backbone. This could provide a way of supporting end-to-end prioritization of multimedia traffic in a mixed Token Ring and ATM environment.

100 Mbps LANs

In general, the 100 Mbps shared media LANs provide a potential alternative to ATM for desktops that require very high speeds for bursty data traffic, but they suffer from the same kinds of limitations as Token Ring and Ethernet regarding time-sensitive multimedia traffic. As a result, they do not match the potential offered by ATM for a revolutionary step forward in network service.

FDDI/CDDI

FDDI and its implementation over copper cabling, known as CDDI, can be considered as 100 Mbps versions of Token Ring. The token-passing ring topology of FDDI and CDDI provides a deterministic access mechanism which allows the full 100 Mbps bandwidth to be exploited, and a prioity mechanism similar to that of Token Ring is supported.

FDDI and CDDI represent the only mature and widely available 100 Mbps LAN technologies today, but high costs have so far deterred the widespread adoption of this technology at the desktop.

Fast Ethernet

Fast Ethernet is a 100 Mbps implementation of the CSMA/CD acccess method used by standard 10 Mbps Ethernet. It lacks packet prioritization. As with Ethernet, the collision detection aspects of the access protocol limit the practical utilization of the available bandwidth to about 40%.

Fast Ethernet is becoming available from a number of vendors at costs which are substantially lower than those for FDDI, and is likely to prove attractive as an upgrade to Ethernet workgroups that have requirements for high speed bursty data movement.

100VG-AnyLAN

100VG-AnyLAN is a 100 Mbps shared media LAN technology that implements an entirely new kind of access protocol, called Demand Priority Access. For ease of integration in both Ethernet and Token Ring environments, 100VG-AnyLAN supports both Ethernet and Token Ring frame formats, and it also supports multiple levels of packet priority. Costs are comparable to those for Fast Ethernet.

The 100VG-AnyLAN concept, though technically elegant, has so far failed to attract widespread support from established vendors of hubs, routers and adapter cards. In the absence of multivendor support, potential users of this technology are likely to face severe difficulties in integrating it with an existing LAN environment.

Desktop ATM in the Wider Context

We have discussed ATM desktop connectivity at length, together with the immediate ATM network infrastructure needed to support practical applications over the ATM network.

In most situations, the desktop ATM network is not going to exist in isolation. It may need to be connected to LANs and WANs based on other technologies, or it may be part of a larger ATM network that extends across private or public wide area links. In this section we look at desktop ATM in the context of such mixed networks.

ATM LAN Backbones

The ability of ATM to handle very large volumes of traffic with low delay makes it well-suited for LAN backbone applications, linking numbers of workgroup Ethernet or Token Ring segments with centralized high performance servers. In this situation, the desktop ATM networks and the ATM LAN backbone can share a common infrastructure. In networks where some desktops are connected by “legacy LANs” — Ethernet and Token Ring — and some desktops are connected by ATM, it is vitally important that all can share access to the same common resources, whether these resources be attached by ATM links or by Ethernet or Token Ring.

In order to provide for interoperability between Ethernet, Token Ring and ATM desktops, we need to ensure that the means for interconnecting legacy LANs with the ATM backbone provides for standardized translation between legacy LAN protocols and ATM. LAN Emulation over ATM provides the answer to this.

Interworking with Legacy LANs

The simplest method for providing interconnection between Ethernet or Token Ring LANs and an ATM backbone is to use LAN switching with LAN Emulation over ATM.

LAN switching is a technique for interconnecting LAN segments of the same technology, ie multiple Ethernets or multiple Token Rings. So how can we use this technique to interconnect with an ATM backbone? Simple — because LAN Emulation allows us to make the ATM backbone look exactly like an Ethernet or Token Ring segment. We can treat the ATM network just like a physical Ethernet or Token Ring, and switch data packets directly into it.

Once a Token Ring packet has been switched onto an ATM network via LAN Emulation, it looks just like other packets that have originated at desktops with direct ATM connections — provided, of course, that these desktops are also running Token Ring LAN Emulation. With a consistent format for Token Ring packets on the ATM network, whether they originated at Token Ring desktops or ATM desktops, we have complete interoperability between the two kinds of desktop access. This means that both Token Ring and ATM desktops can gain direct access to servers which may be located either on the ATM backbone or on another Token Ring.

We have an exactly equivalent situation with Ethernet. Physical Ethernet segments connected to an ATM backbone via an Ethernet switch with LAN Emulation over ATM can interoperate with ATM desktops and servers running Ethernet LAN Emulation.

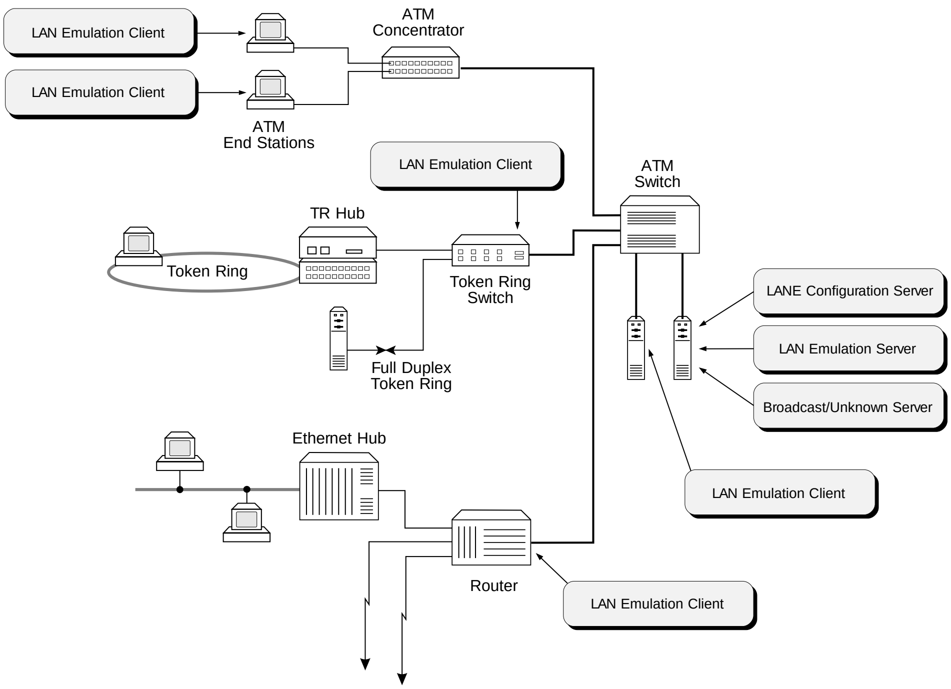

Figure 3: Switches and Routers with LAN Emulation in a mixed ATM LAN environment

In a mixed environment, where both Ethernet and Token Ring segments are switched onto the ATM backbone, there will be two kinds of packet formats present on the ATM network. In this situation, what the ATM network will *not* do is provide protocol conversion from Ethernet to Token Ring or vice versa. However, if we have a server attached directly to the ATM backbone, then the ATM adapter in the server can provide both Token Ring and Ethernet LAN Emulation at the same time. Incoming cells are sorted out according to their frame type, and the adapter card appears to the network operating system as two virtual adapters, Ethernet and Token Ring. This allows for common ATM-attached servers to be accessed by Ethernet, Token Ring *and* ATM desktops without the need for any kind of protocol conversion.

The Role of Routers

Routers have become established as the key building blocks of large networks that comprise interconnected LANs on multiple sites. Routers determine optimum paths through a complex WAN, re-route around failed links, control the propagation of broadcasts and provide local-to-wide area protocol translation — and all these capabilities are key requirements for effective wide area internetworking.

Within a LAN, though, the case for routers is less clear. Nevertheless, routers are widely used today to interconnect multiple local LAN segments. There are probably two main reasons for this.

First, the desire to deal with congestion on overloaded backbones has led to widespread acceptance of collapsed backbone architectures. In the absence of suitable LAN switches on which to base a collapsed backbone, large routers are often chosen to perform this task.

The second reason arises from limitations in the IP addressing scheme. Many organizations have to work with multiple Class C addresses, which limits the number of hosts per subnet to 255. Since the typical size of a LAN segment may be 80-100 stations, it is common for each LAN segment to be designated as a single subnet. In this case, each LAN segment is connected to its own router port.

LAN Emulation allows an ATM network to function as a single very large LAN segment. With LAN switching providing a connection between Ethernet and Token Ring segments and the ATM network, the combined LAN can operate as a single large flat subnet.

The advantage of basing the LAN on a single flat subnet is that any station can communicate directly with any other without needing to pass through a router — eliminating the performance problems that arise from router throughput limitations and latency. The flat subnet model works very well for all protocols apart from IP, where limitations in IP addressing may force the adoption of subnetting within the switched LAN. In practice, the problems can be overcome in one of three ways.

- Adopt a private IP addressing scheme based on Class A addressing, which allows plenty of address space to build large switched LANs that can be addressed as a single subnet. For access to the public Internet, an address translation gateway is required.

- Make use of the "Variable Length Subnet Mask" feature provided in many routers to allow large subnets that accommodate switched LANs to be configured on large sites, while smaller subnets with fewer host addresses can be deployed on smaller sites.

- Configure the router to recognize two or more subnet numbers on the ports that connect it to the switched LAN, so that stations belonging to multiple subnets can be distributed throughout the switched LAN.

Large flat subnets are often thought to be undesirable because of the risk of "broadcast storms". In practice, the effective control of broadcasts can be accomplished within LAN switches and in the broadcast forwarding function of the LAN Emulation scheme, with the aid of manually or dynamically configured broadcast filtering schemes.

By removing the router from the core of the LAN infrastructure and re-positioning it as primarily a means to connect wide-area links and to provide Token Ring to Ethernet translation, the performance demands placed on the router are much reduced and a lower capacity (and hence lower cost) router can be employed.

ATM in the WAN

One of the great advantages of ATM over traditional LAN technologies is that ATM can be deployed in both the LAN and the WAN. Conventional LAN internetworks comprise LAN segments, leased lines, and connections to public packet and circuit-switching networks, all linked by routers. Each different kind of link operates a different protocol, requiring complex and performance-sapping translation when moving between LAN and WAN. And the processing power required to carry out this protocol translation in the router represents a major element of cost.

ATM cells originating at the desktop can be switched directly onto wide area ATM links, without any kind of protocol translation. ATM switches can provide multiple different kinds of physical level interfaces to drive different kinds of wide area link.

Private ATM WANs

In the short term, most wide area ATM networking is likely be based on private ATM networks. These will comprise ATM switches with interfaces to leased lines such as T1 (1.55 Mbps), T3 (45 Mbps) or E1 (2 Mbps). ATM Forum standards exist for each of these kinds of ATM WAN interface.

Just as in an ATM LAN comprising multiple switches, a private ATM WAN makes use of the Private Network-to-Network Interface (P-NNI) to communicate between switches to determine the optimum route through the network for any particular Virtual Channel Connection.

Public Cell Switching Services

Longer term, many organizations will find it makes economic sense to make use of public cell-switching services rather than building their own private ATM WAN based on leased lines. In this case, the local ATM network at each site is connected by a high-speed line to the nearest ATM switch belonging to the service provider. The public network can be regarded as an ATM "cloud" which carries ATM cells between the customer's various access points.

The expected lower costs of using public cell switching are based on economies of scale in the hardware installed by the service provider, and in more efficient use of the bandwidth on long distance links.

Connecting to public cell-switching services requires additional interface standards for the private-to-public switch interface. Work on these standards is under way in the ATM Forum. For public cell-switching services to become widely available, a standard is also required for the interface between switches belonging to different carriers. The complexity inherent in this interface is likely to mean that metropolitan area services provided by single carriers will be available some time before long distance services.

Conclusion

ATM at the desktop provides an ideal infrastructure for the support of networked multimedia applications, including desktop conferencing, interactive training and the retrieval of stored video data objects. Desktop ATM also represents a far more powerful LAN solution than either Ethernet or Token Ring for all general purpose LAN applications.

LAN Emulation over ATM provides the key technology element that supports a graceful migration strategy from Ethernet and Token Ring to ATM at the desktop, since it addresses the requirements to support existing LAN applications at desktops equipped with ATM and to interconnect existing desktop and backbone LANs with ATM.

Where ATM workgroups are being deployed to meet specific application needs, an ATM backbone will serve the dual purpose of interconnecting these workgroups, and greatly improving the capacity and capability of the existing LAN interconnections.

The rapid market acceptance of LAN switching as a means of upgrading today's Ethernet and Token Ring environments carries a powerful message about the value of switched networks. One day, switching will be pervasive in our networks, and ATM, as the purest embodiment of switching, will be right there on our desktops.

Appendix A

Brief Introduction to ATM

ATM is a new communications technology that combines elements of current techniques for both voice and data networking to provide a universal solution for carrying all types of communications traffic in both the local and wide area.

Networking Voice and Data Today

Data internetworks today are based on packet switching techniques. Packets contain a variable amount of data, up to several thousand bytes, and contain a header with addressing information that network elements use to direct the packet to its final destination. Most data internetworks are “connectionless”, which means that there is no concept of a channel between the originator and the recipient of a data packet; each packet header contains all the information necessary for the network to deliver the packet to its destination.

Voice and video networks today are based on time division multiplexing techniques, where high speed links are used to carry multiple channels of voice by dividing the available bandwidth up into fixed time slots, each of which is occupied by one voice channel. Voice communications are invariably “connection-oriented”, which means that a connection has to be set up between the originator and the recipient before any information can be transferred. In other words, you have to dial the other party.

ATM - Asynchronous Transfer Mode

ATM combines elements of both traditional voice and data communications techniques to provide a solution that meets the needs of both. ATM is based on packets, but instead of variable length packets that may be several thousand bytes long, it uses short, fixed length data units called “cells”. Each cell carries 48 bytes of information. Unlike the synchronous time slots of a traditional time-division multiplexed voice network, ATM cells have no fixed timing relationship - hence the term “Asynchronous Transfer Mode”.

ATM is a connection-oriented technology, which means that an end station must set up a connection to a destination before it can transfer information. This means that cell headers do not have to contain the full address of the destination, only the identification of the connection. This information can be stored in a cell header only 5 bytes in length, which is much shorter that conventional data packet headers. Details of each connection, known as a “virtual channel”, are held at every intermediate switch along the path of the connection through the network. Switches use this information to determine how each cell should be forwarded to reach its eventual destination.

ATM Networks

ATM networks consist only of end stations and switches. In this sense they resemble voice networks far more than data internetworks, which consist of many different elements including wiring concentrators, repeaters, local and remote bridges, switches, and routers.

In an ATM network, every end station has a dedicated full-duplex connection to an ATM switch. ATM switches are connected to both end stations and to other ATM switches. Both end station links and inter-switch links carry ATM cells, but operate slightly different signalling protocols, described by User-to-Network Interface (UNI) and Network-to-Network Interface (NNI) specifications.

To communicate across an ATM network, an end station must first signal a connection request to the switch to which it is attached, specifying the address of the end station to which it wants to talk. This is exactly analogous to dialling a phone call. The ATM switches co-operate to locate the called station, determine the best route, and set up the virtual channel connection between the two stations. The switch nearest to the originating station notifies it when the channel has been set up, and supplies a virtual channel connection number that identifies this connection. The station places this channel number in each cell header that it transmits to the relevant destination, and the switches in the network interpret the channel number in each cell header to determine on which port to forward the cells.

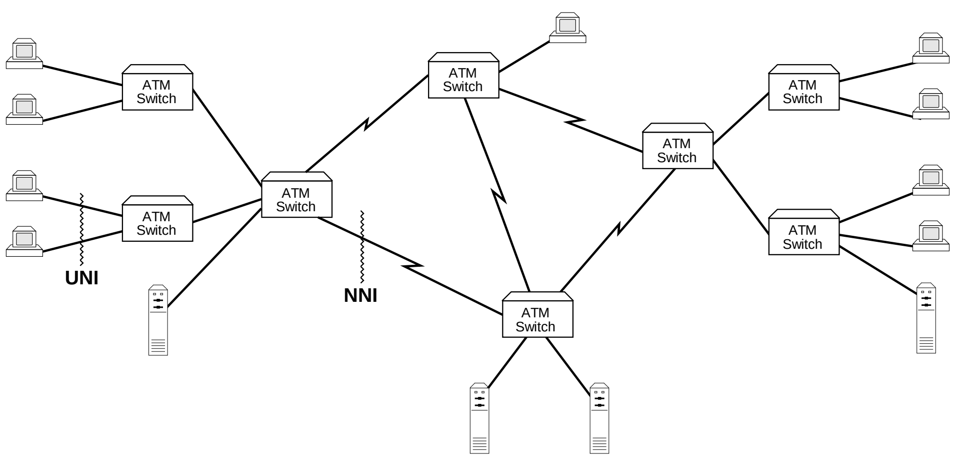

Figure A-1: ATM network consisting of ATM switches and end stations

Benefits of ATM

The value of ATM lies in its ability cost-effectively to handle data, video and voice communications at a variety of speeds in both the local and wide area, with equal facility.

ATM's unique capabilities stem from three key features: the fine "granularity" of the cell-based approach to data transport, the processing efficiency of connection-oriented cell forwarding, and the point-to-point nature of all communications links.

Carrying all traffic, voice and data alike, in small fixed-length cells allows different types of traffic to share common links without introducing the variable delays that arise with large data packets. In conventional networks, such packets tend to hog bandwidth on low speed links. Large data packets also result in excessive end-to-end transmission delays, as each packet is subject to store-and-forward latency at every intermediate bridge or router.

Fixed-length cells addressed by connection information lend themselves to very efficient processing by switches. ATM switches can be built from dedicated silicon that carries out address translation and data forwarding entirely in hardware, achieving data handling cost efficiencies that are orders of magnitude greater than software-intensive routers.

Point-to-point links eliminate the shared medium access protocols used in today's LAN protocols, with their inherent limits on the span of an individual LAN segment. This has two benefits: ATM links can operate over very long distances if needed, and at a variety of different speeds. For this reason, ATM can be implemented successfully in both local and wide area networks.

Amsterdam

Atlanta

Bangkok

Berlin

Boston

Brussels

Cape Town

Chicago

Cologne

Copenhagen

Dallas

Denver

Detroit

Frankfurt

Hong Kong

Jakarta

Johannesburg

Kuala Lumpur

London

Los Angeles

Madrid

Minneapolis

Munich

Nashville

New York

Oslo

Paris

Philadelphia

San Jose

Seattle

Singapore

Stockholm

Sydney

Tokyo

Toronto

Vancouver

Washington DC

Madge Europe
Madge Networks Ltd
Loudwater
High Wycombe
Bucks HP10 9QZ
England
Tel: +44 1628 858000
Fax: +44 1628 858011

Madge Americas
Madge Networks Inc
2310 North First Street
San Jose
Calif. 95131-1011
United States
Tel: +1 408 955 0700
Fax: +1 408 955 0970

Madge Asia
Madge International Ltd
64-01 Central Plaza
18 Harbour Road
Wanchai
Hong Kong
Tel: +852 2593 9888
Fax: +852 2519 8022

Madge Japan
Madge Japan KK
Believe Mita
43-16 Shiba 3-chome
Minato-ku, Tokyo 105
Japan
Tel: +81 3 5232 3281
Fax: +81 3 5232 3208

Madge Networks reserves the right to change specifications without notice.

980-096-01 4/95