Silicon Graphics
Computer Systems

Asynchronous Transfer Mode:

Implementations for 1994

by:
John Talbott
Silicon Graphics Inc.

March 16, 1994 5:06 pm

SUMMARY: The most exciting networking concept of the early 1990's is Asynchronous Transfer Mode or ATM. ATM promises to deliver much higher bandwidth over wider areas with lower latency than ever before available, but in the fervor to bring higher speeds to networks, many of the implementation details of ATM have yet to be determined. In this paper, we discuss the current state of the ATM standard and suggest appropriate uses for the technology today.

Asynchronous Transfer Mode: Implementations for 1994

Silicon Graphics, Inc.

Mountain View, California

Executive Summary

Asynchronous Transfer Mode (ATM) is emerging quickly as a likely "network of choice" for replacing today's point-to-point leased lines and today's X.25 networks, and for use in emerging multimedia delivery systems to the home. Some proponents also argue for replacing current LAN technology (Ethernet, Token Ring, and FDDI) with ATM but the reasons for not doing so today are quite clear.

ATM comes from a heritage of telephone company developed networking protocols designed to move significant amounts of information over long distances with little delay. ATM reflects the current state of the network in that it is designed without the error control characteristics of X.25 and Frame Relay. ATM also transfers huge amounts of data very quickly, since it works best over the high speed, low error rate SONET/SDH networks deployed by the telephone companies today.

The emerging markets of interactive television and wide area multimedia will also benefit from the growth of ATM. As a data network, ATM has unique characteristics of guaranteed low latency to insure that the network delivers voice, video, and data delivery in a predictable manner. This enables the server to play movies over the network and to host interactive gaming, shopping, and other services.

On the other hand, ATM has not yet matured so that it will handle traditional data traffic well in local area networks. Advanced concepts such as multicasting of IP packets have no standard mechanisms. Flow control does not exist and congestion management is rudimentary. Since local area networks respond very poorly to data loss and these deficiencies of ATM make the probability of data loss on normal network tasks too high to be acceptable, customers should avoid installing ATM as a "fast LAN" replacement. Instead, customers who need more speed for traditional traffic should look to FDDI and HiPPI.

ATM has several strengths. The protocol does not define a data link, so ATM will scale from today's 45 and 155 Mb/s speeds to tomorrow's gigabit networks. The deterministic latency characteristics make ATM ideal for high speed voice and video traffic. Using switching allows an ATM user to access many endpoints for the cost of a single line and thus save money on wide area network facilities. If applied where the user can take advantage of its strengths, ATM will provide a major step forward in networking capabilities.

Section 1 A Brief History of High Bandwidth Networks

If one were to read the trade press, one might imagine that Asynchronous Transfer Mode (ATM) will obsolete all other protocols and that starting next year, the industry will not have to worry about any other type of networking. This promise sounds great. Unfortunately, like all protocols, ATM has a great deal of promise associated with it, but will take a few years to realize its potential. What ATM does promise is a high speed network with several very compelling features.

Application
Presentation
Session
Transport
Network
Data Link
Physical

Table 1: The Open Systems Interconnection Model

Asynchronous Transfer Mode defines a way of implementing a small portion of the complete protocol stack required in any computer to computer communications. If one thinks in terms of the OSI reference model (Table 1), ATM occurs in, but does not completely comprise, the third layer (Network). This layer has the responsibility for taking messages from the network independent sublayer (the route to IP, CONS, and other network providers), segmenting those messages appropriately for the chosen data link layer, mapping them to the frames of the data link layer, and then delivering them correctly to that data link transport. In the case of ATM, this means that packets delivered to ATM are segmented into a fixed size cell, addressed according to a circuit identifier, and passed to a circuit for delivery.

From Table 1, one can see that ATM must interface to a data link and to a network. As with X.25, ATM may be used without upper layer protocols for certain types of traffic but because these popular protocol layers (TCP and IP, for instance) provide useful functions, they will continue to exist in the ATM world. Additionally, ATM will be defined for use over many different data links. This means ATM provides only a small part of the mechanism for computer to computer communication. If this suggests to you that maybe the whole thing is about to get a lot more complex, you are either perceptive or experienced. As we begin to understand ATM, we need to remember that it is a part of a much larger whole and that ATM is clearly defined to work within the worldwide communications framework of today and of tomorrow.

1.1 First Generation Digital Networks

1.1.1 Why a new Data Link?

The network that has existed before is a hierarchical digital network based on the standard T-1/E-1[1] carrier. These signals multiplex together to form higher bandwidth carriers like T-3/E-3. Service on the backbone networks of the major phone companies runs at rates in excess of 2 gigabits per second but the design of the network is such that these high data rate connections are really lots of T-1 speed connections multiplexed together. It is possible now to buy service at faster speeds than T-1 and in the case of Switched Multimegabit Data Service (SMDS) that service does not have to be dedicated to a point-to-point link, but these carriers are still designed as dedicated bandwidth channels.

1.1.2 Optimized Data Channels

Voice, by nature, functions well on a synchronous network. The phone company digitizes the incoming analog signal (8 bits @ 8 khz = 64 kbps) and reserves a slot on the T-1/E-1 carrier for the duration of the call. When a user speaks across this connection, the digital signal carrying that voice information will be placed in a slot on an allocated channel and carried to the other end of the channel. When the speaker pauses and no data (other than silence) is being sent, the network continues to hold

1. The T-1 carrier is the base level carrier in the Americas and parts of Asia. It consists of 24 voice channels, each digitized at a rate of 8kb/s with an 8 bit sample. The resulting 64kb/s data rate when multiplied by 24 is 1.536 Mb/s. The T-1 rate of 1.544 Mb/s includes additional signalling. In Europe, the standard carrier, often called E-1, is 30 channels plus signalling resulting in a rate 2.048 Mb/s.

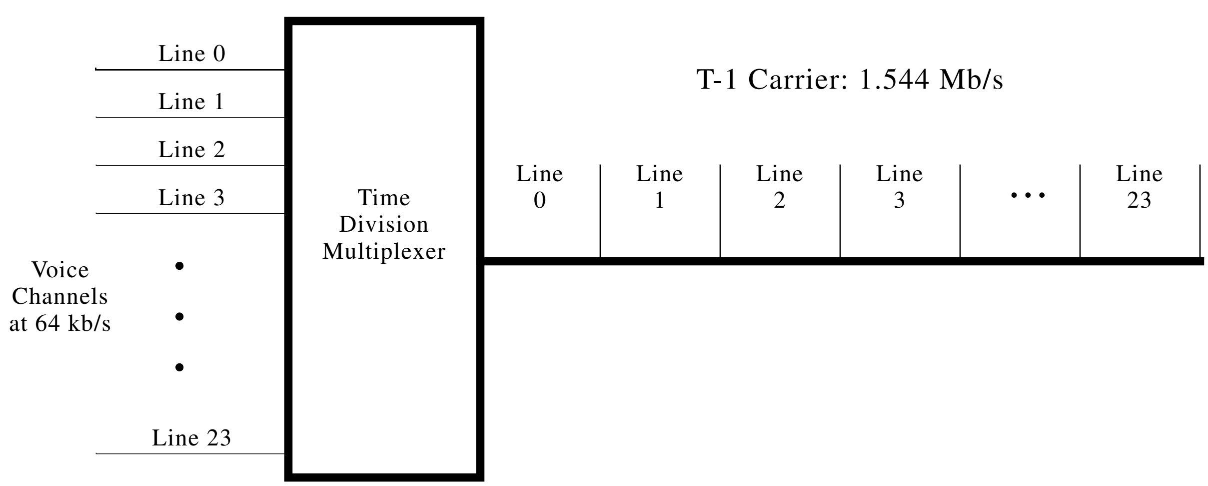

Figure 1. Multiplexing of 64 kb/s voice channels onto a T-1 Carrier

that slot open until the connection closes. When the connection terminates, that slot becomes available for another user. We will refer to this scheme as time-division multiplexing (TDM) because each piece of data has a slot on the channel and that slot recurs at synchronous intervals. The traffic on such a network is called "constant bit rate" because it is sent at a constant rate.

Instead of dedicated bandwidth, data traffic benefits from packet transmission where the network transmits packets on a first come, first served basis (for the purpose of this argument). We call this type of traffic "bursty" because transmissions (like a file transfer) occur at random intervals. In effect, the station puts bursts of traffic on the network.. Shared networks, like Ethernet work well for bursty traffic because every user shares the bandwidth leaving the whole of the bandwidth available for use by any user. However, because the network does not reserve any bandwidth for a particular connection, every piece of data must wait until the network is available, making it impossible to guarantee the latency across the network.

1.1.3 Enter Multimedia

The conflict between packet-oriented traffic and TDM traffic becomes moot when the assumptions about data traffic change. We can only assume that data traffic is bursty when the data consists primarily of files and text. When the computer generated information requires timely, synchronous delivery (audio, video, visual simulation, real-time graphics, and so on) the network must take on the characteristics of a voice network in that it must deliver information to the screen (or speakers) without pauses, glitches, or time-outs. At the same time, the network must not drop a single piece of information, because file transfers and textual information need accuracy.

1.1.4 Differentiating Between True Multimedia Traffic and Multimedia Files

It is important to note at this point that traffic which appears to be multimedia traffic may not be what it appears. Multimedia programs often transfer data directly across the network as a file and store it on a local disk or read it directly into RAM memory. In this case, the traffic comes as quickly as possible and carries none of the characteristics of constant bit rate traffic, so should be considered traditional bursty traffic.

On the other hand, many users now "watch" video over the network without storing the video data (beyond memory required for a couple of seconds of buffering) in the local (client) machine. This means that the data must be delivered at a reasonably predictable rate and that the session begins with the arrival of the first frame of video. Video conferencing works this way and, of course, so does regular voice traffic. Video servers which allow the user to play the video over the network rather than transfer the whole file before playing work this way as well.

1.2 High Bandwidth, Synchronous Networks

Because of the need to move large amounts of data at very high rates, the telephony community developed a standard for transporting data over fiber backbones called Synchronous Optical Network (SONET) in North America or Synchronous Digital Hierarchy (SDH) in other parts of the world.[1] The purpose of this network is to provide a wide pipe which can pass any type of data, achieve good bandwidth utilization, and provide integrated constant bit rate and bursty traffic services. To achieve this, the SONET/SDH requires upper layer protocols to handle routing, signalling, and switching. Collectively, this network is often called the "Broadband Integrated Services Digital Network" or BISDN. ATM has been chosen as the switching protocol for asynchronous traffic on this network.

The SONET/SDH define something called a Synchronous Payload Envelope (SPE), the basic unit of transport on the network. Upper layer protocols are mapped into an SPE and transported along a SONET path to their destination. ATM cells are mapped into a specific SPE; other protocols are mapped into different SPE formats. ATM, therefore, can be said to run over SONET, but may not be the only thing running on a particular SONET link.

This implies that the SONET and SDH stand alone as a network. For those who are interested in pursuing the topic further, please refer to one of the many high quality text books available today on high speed telephony. From this point on, we will exclusively discuss ATM.

Section 2 What is Asynchronous Transfer Mode?

2.1 Cell Relay and Switching

Packet switching models use circuit identifiers to route packets through the network. Voice networks switch data through the network by reserving space on facilities throughout the network then releasing that space when the connection finishes. In each case, the concept of switching allows the users to share the network.

This may sound simpler than it should because the telephone companies have provided high quality switching systems for Plain Old Telephone Service (POTS - a very common acronym) for many years. However, much of wide area networking has been done with dedicated leased lines instead of packet switching like X.25. This creates the problem of underutilization of the network. Recent improvements in the quality of the network have created Frame Relay and SMDS because they offer the advantages of switching packets (or relaying frames). Frame Relay and SMDS as well as ATM can be considered improvements over

1. SONET and SDH are not exactly the same, but the differences are not important to this paper. In the implementation Silicon Graphics has chosen for ATM, the optical carrier chipset handles both SONET and SDH.

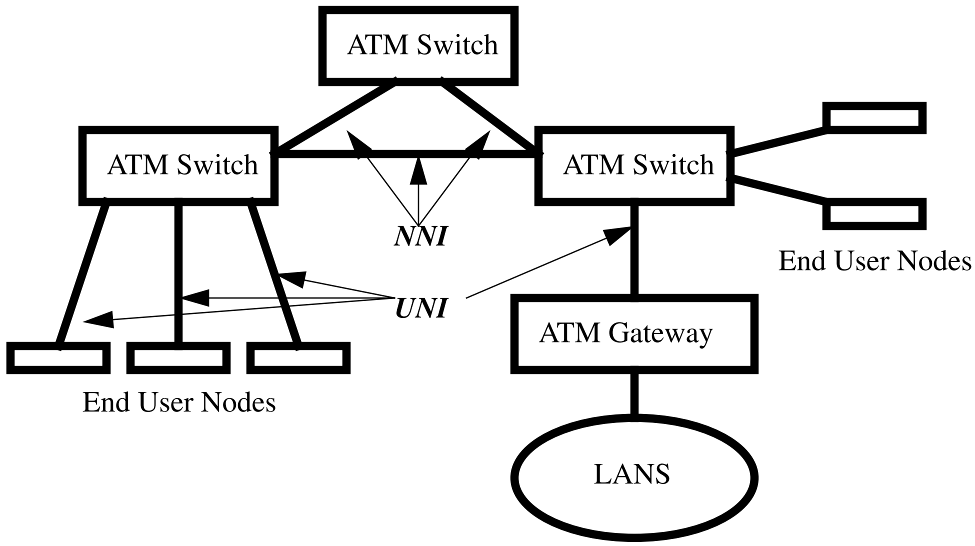

Figure 2. Typical ATM Network Implementation

X.25 on the packet switching model, but Frame Relay and SMDS do not include any accommodations for multimedia. Those who need guaranteed delivery times must use leased lines.

ATM bridges this gap. At its core, ATM provides services for switching data (called cells in ATM) across high speed backbone networks. Users must ask the network for the amount of bandwidth that they need and the network (in general) won't allow over-subscription. Because the network is switched, the path through the network is known. For these two reasons, the ATM system can guarantee delivery of cells with known delay through the network at a known bandwidth.

We can now define ATM as the switching protocol used to establish paths and deliver data across high bandwidth networks. Furthermore, ATM consists of hosts which generate traffic and switches which move traffic through the network. The interface to the network switch from the host is called the User-Network Interface (UNI) and the interface between network switches is called the Network-Network Interface (NNI).

2.1.1 The Construction of the ATM Cell

Figures 3 and 4 show the bit layout for the 53 bytes of the ATM cell transferred from the host to the first network switch. The definitions of each field in the header are not important to this paper and are available in any good reference on ATM. What is important here is that the ATM cell is very small and is fixed length. AAL data is encapsulated within the data portion of the cell so that the larger packets can be reconstructed. The only identifiers read during the network transport are those in the Virtual Path Identifier (VPI) and Virtual Circuit Identifier (VCI) fields. The VPI indicates the destination of the next hop and the

Header 5 Octets	Information Field -- 48 Octets

Figure 3. The ATM Cell

General Flow Control 4 bits	Virtual Path Identifier 8 bits	Virtual Circuit Identifier 16 bits	Payload Type Identifier 3 bits	Cell Loss Priority 1 bit	Header Error Control 8 bits

Figure 4. The Header of the ATM Cell on the Host to Network Connection

VCI indicates the end-to-end circuit designation. This means that ATM does not check and does not care what the upper layer protocol might be. The AAL does, of course, care what lives above it but that is only for the purposes of interfacing correctly to the next layer.

2.2 *Physical Network Options*

ATM can run over any medium on which the proper mappings exist. Since the people who originally designed ATM (mostly at telephony research organizations) intended it for high bandwidth synchronous networks, most people implement it around these specifications. However, installed base considerations and economics have caused a variety of physical media to emerge. We will discuss the most popular in Section 3.

2.3 *ATM Adaptation Layer*

Since ATM defines a data link layer for network communication, we must now define the upper layer interface. This interface, called the ATM Adaptation Layer (AAL) must connect to several different types of traffic. Systems handle traffic most efficiently when processing large sized frames (64 kB) rather than tiny ATM cells, so it is desireable to hand large frames to the AAL. The AAL must be able to handle the segmentation of this data into ATM 53 byte cells. Multimedia traffic requires direct connection to the ATM layer since no upper layer protocol is

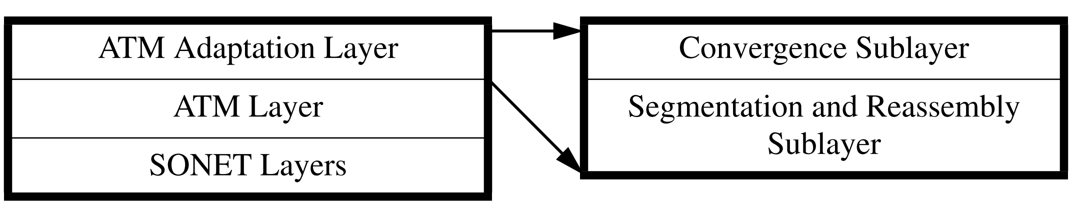

Figure 5. ATM Layers and the ATM Adaptation Sublayers necessary for the transportation of video and audio, but data and, some formats of multimedia specify a transport protocol, so allowances must be available to service these protocols.

The ATM Adaptation Layer (AAL) divides into two parts: the Segmentation and Reassembly (SAR) portion and the Convergence Sublayer (CS). The SAR handles the assembly and segmentation of data packets into the 53 byte ATM cells. A frame delivered to this layer may be quite large, depending on the upper layer protocol, so the SAR handles the task of breaking up and reassembling these data frames. Usually, the SAR is implemented in hardware.

The CS handles the functions of message identification and synchronization. These functions are necessary to the mapping (converging) of multiple upper layer protocols onto one data link network.

5 different functions have been defined by the ATM Forum for the AAL and they are referred to by their number: AAL1-AAL5.

AAL 1	Adaptation for Constant Bit Rate Services
AAL 2	Adaptation for Variable Bit Rate Services
AAL 3/4	Adaptation for Data Services
AAL 5	Adaptation for Data Services

Table 4: ATM Adaptation Layers

AAL 3/4 has already been discarded by the ATM Forum for computer traffic because it has high overhead. AAL 1 and AAL 2 are useful for those choosing to use direct connections to the ATM layer without using the upper layer protocols such as TCP. AAL5 has been defined by the ATM Forum and adopted by the International Telecommunications Union[1] (ITU) specifically for the use of data communications and is the layer SGI will expect to use for IP and all other traffic which depends on upper layer protocols. The Forum is considering other versions of the AAL to handle emerging traffic types.

2.4 *Virtual Circuits*

Two popular varieties of networks exist today. Datagram networks, like IP over Ethernet, do not determine a route through the network based on a connection established at the beginning of the session. Instead, each router in the network must decide where to next send any packet it receives. Therefore, each packet carries an address (or several addresses) indicating its destination.

1. The ITU is the governing body for ATM and many telecommunications standards. The familiar acronym, CCITT, was a committee of the ITU which has ceased to exist.

Virtual circuit networks, like ATM, determine a path prior to the transmission of any data. Each router or switch in the network maintains knowledge of the circuit and keeps that knowledge until the circuit is disconnected. The cell does not need to have a fixed identifier and, in fact, only needs to carry information which will identify it at its next destination. The destination can then change the identifier to identify the data unit to the next destination.

ATM cells have two identification fields, the Virtual Circuit Identifier (VCI) and Virtual Path Identier (VPI). Together, they form the identifier necessary for the ATM switch to forward the cell to the next destination. The process of establishing the knowledge throughout the network is called "set-up" and the mechanism is signalling. One can think of this as a "data phone call" and the signalling as being the call set up.

2.5 Signalling

Signalling protocols tend to be fairly sophisticated in complex networks like the phone network even though the user only sees certain "signals". For the phone, these signals are the touch tones pressed at call setup. The levels of signalling that go on behind the scene in order to properly route calls are very sophisticated, but are easy to control if you have an organization like the major network providers to define and implement the protocols. Private networks based on ATM will require industry standards to function, if the network will become heterogeneous.

2.5.1 Permanent Virtual Circuits

In the case of ATM, standardized signalling throughout the network has not yet been implemented. The compensate for this, permanent virtual circuits (PVC's), allow the network manager to manually allocate a given circuit to a connection between two hosts. The network manager creates a table of circuit identifiers and destinations in the host and in the switch. Based on this table, the sending host attaches a circuit identifier to the cell, the switch matches the identifier to a route, and the switch forwards the cell to its destination. In this process, the switches and hosts have pre-allocated the bandwidth for the circuit, so predicting the response is easy. However, when no traffic uses the circuit, the bandwidth is wasted. While these type of circuits are commonly used today, they will likely disappear when more sophisticated signalling protocols such as the ATM Forum Signalling Protocol become ubiquitous.

While PVC's exist in virtually all ATM equipment available today, each vendor uses a different management scheme for setting up and tearing down connections. This does not render the switches non-inter-operable, but it does create significant headaches for those who have to maintain the system. Nonetheless, creating PVC's is a good way to gain ATM functionality today.

2.5.2 Switched Virtual Circuits

Unlike PVC's, Switched Virtual Circuit (SVC) protocols allow the network elements to set up and tear down virtual circuits at the inception of each communication. Using SVC's, the network utilization increases because because network only reserves bandwidth for circuits while those circuits are in use. When the hosts finish using the circuit, they remove it, and the circuit and bandwidth become available for other users.

The newly adopted ATM Forum Signalling Protocol provides an industry standard mechanism for managing these circuits. However, as this protocol stands today, only the UNI is standardized. Switches tend to use proprietary signalling mechanisms for the NNI, so SVC's may not be available between switches from different vendors. Today, switches which incorporate the ATM Forum Signalling Protocol typically will allow SVC's through one switch or through several switches from that same vendor.

2.5.3 LAN Emulation

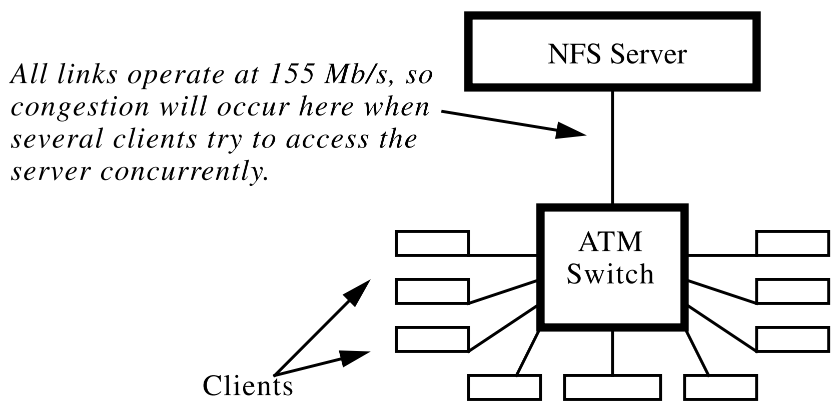

Figure 6. Congestion in a Local Area Network

The protocol used for SVC's works very nicely for large file transfers and other types of connections which involve long holding times on the circuits. Unfortunately, a great deal of local area network traffic does not obey this model. NFS, for instance, operates on a connectionless model and the majority of NFS packets are very small. In addition, NFS uses a client-server model in which a single server will serve many clients, so the possibility of congesting the link that runs from the server to the switch rises dramatically.

The ATM Forum has not adopted a standard method for solving this problem today. The subcommittee in charge of LAN emulation is reviewing several proposals and expects to define a standard sometime in early 1995. In the meantime, several vendors have created their own mechanism which works with their switches and host adapters. Though proprietary, these standards do allow at least limited functionality in LAN environments. Network managers need to plan for the differences between vendor schemes and avoid trying to run differing LAN emulation protocols between switches and host adapters.

2.6 Constant Bit Rate Traffic

The one thing that ATM does is to deliver constant bit rate traffic over the network. As mentioned before, multimedia traffic places a high demand on the network to deliver information with low latency. Because the host has a direct physical connection to the host, each host receives a given amount of bandwidth to the local switch. At the initiation of a virtual circuit, the initiator can reserve a certain amount of bandwidth through the entire network. This allows the host to send constant bit rate traffic over a channel without any queuing latency.

2.7 Rate Control

Popular data networks have no defined method for guaranteeing bandwidth to processes which require delivery of information in a manner suitable for multimedia. TDM networks, like the phone network, offer very specific guaranteed data rates which are not very flexible.

ATM implements a concept called rate control. When an application opens a session, the application asks the network for a given amount of bandwidth. That bandwidth is granted for the duration of the connection, and the network expects the transmitting agent to limit its transmissions to the agreed data rate on that virtual circuit. If the originating host trys to send more data then allowed, the first switch should discard that data.

Hopefully, this method will prevent switches from granting more bandwidth than the physical connection to the network allows. Some implementations of ATM today do not include rate control. In some cases, it is also possible to oversubscribe the switch. In either case, the probability of overloading a data link between a switch and a host is non-zero, so, without flow control, data can be lost.

Section 3 Current ATM Implementations

3.1 Installed Telephony Carriers

The first implementations available from public network vendors will naturally incorporate the data pipes already in place. Several companies have announced public ATM service on DS-3 circuits at 45 Mb/s. To accomplish this, the sender encapsulates the ATM cells in the in the DS-3 protocol adding extra overhead, but providing a functional link. The maximum throughput one might expect on this line is 36 Mb/s.

Carriers are beginning to announce implementation of ATM at T-1/E-1 speeds as well. This would require encapsulating ATM cells in the DS-1 protocol and would also result in high overhead, but they conform to existing deployments of wide area networking facilities. These slower speed connections also exist within public network tariff structures, installed switching and carrier systems will handle the protocols, and the pricing structure is well-known.

SONET/SDH rates of 155 Mb/s and above are common in the hierarchy of the telephone company backbone, but do not exist at the customer point of access today. In a few special cases (including several in which Silicon Graphics is participating), carriers have extended their network to customers for ATM testing. When the pricing and the fiber optic infrastructure exist to support wide deployment of these rates, they will be available.[1]

3.2 LAN Carriers

Local and private network users do not face the same constraints as the public network, so users are more free to experiment. Because multimode fiber used in FDDI was available in the early 1990's, the TAXI chipset[2] became the most popular carrier for ATM. In 1994, however, most vendors will switch to the SONET/SDH rates as those chipsets become cheaper and more readily available.

The most popular speed for local installations is STM-1 or OC-3c at 155 Mb/s. Local networks typically use multimode fiber[3] or copper Unshielded Twisted Pair (UTP). No specification exists today for STS-3c[4] on UTP, but the ATM Forum will move toward one this year, hopefully lowering the cost of the network installation. The standard will inevitably use category 5 wiring. Vendors are implementing this today.

Several slower speed implementations have been adopted by some vendors so that customers could continue to use category 2 UTP. These include 51 Mb/s, a defined SONET/SDH speed, and 25 Mb/s, a speed for which IBM has argued. Neither of these speeds offers much improvement over FDDI and Ethernet, though they do provide switched networks over installed wiring. Some customers will adopt them.

1. Several countries and several public network providers have aggressively deployed high speed network in hopes of capturing a competitive advantage in the coming years. In Singapore, the government is currently wiring the entire city for fiber communications. Japan and Germany have moved strongly toward digital services based on ISDN, and are currently installing high bandwidth networks. AT&T, SPRINT, and MCI all chose to completely rewire their backbone networks for fiber in the late 1980's and have since completed those projects.

2. TAXI has been deployed at two data rates: 100 Mb/s and 140 Mb/s. TAXI uses an encoding scheme known as 4B/5B and is often identified by this acronym.

3. Again, this is not an important distinction. Single mode fiber is more expensive than multimode and will transport a given signal with less loss over a given distance. This means that long haul carriers need to install fewer repeaters on their single mode fiber backbones. In local networks, the distance between nodes rarely spans the 2 km limit of multimode fiber, so the increased cost of single mode fiber is not useful.

4. Naturally, copper cabling would not use "Optical Carrier 3". STS-3 defines the logical carrier used at 155Mb/s whether on fiber or on copper.

The following table lists the currently defined data rates for the networks which carry ATM and their names. The last entry, OC-12c, has not yet been heavily adopted in computer systems because of cost considerations. The current market is pushing towards OC-3 as the standard for this year. OC-12 will emerge when the market grows large enough to support the infrastructure required.

Existing	Proprietary	SDH	SONET	Data Rate (Mb/s)	Medium	Area
T-1				1.544	UTP	Americas, Asia
E-1				2.048	UTP	Europe, Australia
	IBM			25		IBM, some PC
E-3				34		Europe, World
T-3, DS-3				45	Coax	World
			STS-1	51	Fiber, UTP	World
	TAXI 4B/5B			100	Fiber	World
		STM-1	STS-3c	155	Fiber, UTP	World

Table 5: Available Carriers for ATM

3.3 Functions of the Host Adapter

The Host Adapter card design for ATM becomes very important because of the interaction required between the host and the adapter. By definition, ATM cells consist of 5 header bytes and 48 data bytes. At a rate of 155 Mb/s, the network will send somewhere in excess of 350,000 cells to the host adapter every second. If the host adapter cannot handle this traffic, the achievable data rate will drop. In addition, if the host adapter depends too heavily on a host which cannot handle this data rate, the achievable rate will drop again.

At one point, it was clear that the entire protocol needed to run on a "smart" card so that the host processor spent little effort on input/output (I/O). Today’s fast processors have become powerful enough to saturate Ethernet with little effect on the processor, so it has become cheaper to run the protocols in the host and to use host memory and simple host adapters. When FDDI came along, the amount of processor power required to service an interface increased and so vendors added more power to their network interface cards (NIC’s).

In the case of ATM today, the amount of processor required to service the large number of cells suggests that much of the processing should occur on the board. This requires a significant investment in a host adapter processor and on memory dedicated to the process. Some boards available today use DMA techniques to use host memory instead of dedicated memory, but this puts a significant amount of traffic on the bus and a load on the processor. Boards using this design perform quite poorly, but are available at a lower cost.

3.4 IP Encapsulation

The ATM Forum has proceeded quickly to adopt several standards for ATM, but the exact use of the AAL for encapsulation of IP (and other network traffic) has not been determined. RFC 1483 suggests two methods for protocol encapsulation. The first, LLC encapsulation, uses the familiar LLC mappings attached to the protocol data unit. The LLC mapping allows the AAL to distinguish between various upper layer protocols and therefore direct the completed packet to the right place.

The second method, VC Based Multiplexing, allows the system to open a separate virtual circuit for each upper layer protocol. IP packets would then arrive on one virtual circuit, CONS would arrive on another, and so on. VC Based Multiplexing becomes unwieldy if the number of different protocols serviced is quite high leading to a large number of virtual circuits, but this is not likely to be a technical problem. For a discussion of the two methods, please see RFC 1483.[1]

However, in discussing LAN Emulation, we noted that many vendors have incorporated proprietary interfaces in lieu of usable standards. It is at this point, upper layer protocol encapsulation, that the implementations become proprietary. Because of this, IP inter-operability over ATM is not guaranteed between different manufacturers. Silicon Graphics stresses the need to adhere to standards, and will therefore implement VC Based Multiplexing. Other vendors of host adapters and switches have not chosen this path.

3.5 Flow Control

Another area where much contention exists with ATM is in the area of flow control and specifically link-to-link flow control. No flow control has been specified for the link because super high bandwidth networks like ATM would not require it. Unfortunately, in the ten years since ATM was proposed, this has not proved to be true. The ATM Forum members have suggested various methods for LAN emulation and flow control, but none has yet been adopted.[2]

---

1. Heinanen, Juha. RFC 1483: Multiprotocol Encapsulation over ATM Adaptation Layer 5. The texts of the RFC's are available on many Internet Network Hosts typically in a directory called "rfc". The common ftp site is "ds.internic.net".

Section 4 The Silicon Graphics Implementation

Now that we have described ATM, we describe how Silicon Graphics will choose to implement this technology.

4.1 The Industry View

If one listens to those promoting ATM in telephone and networking companies, one will find fervent supporters of ATM proclaiming the good news that the reign of ATM is at hand.

> As customers evolve toward broadband networking, ATM can be used to "reinvent the LAN," especially for users of demanding client/server applications. ATM LANs make it possible to have true "plug-and-play" LAN administration, where switches learn the identity of workstations that have been added to the network or moved within it...
>
> ATM not only offers higher speeds than most LANs today, but it can also move to even higher levels of performance over the next few years without requiring a change in network architecture.[1]

Industry pundits show a little more restraint:

> Craig Partridge, a senior scientist at Bolt Beranek and Newman (BBN) and author of Gigabit Networking (Addison-Wesley, 1994), says that although enterprise ATM switches from Fore Systems and Lightstream (a joint venture of Ungermann-Bass and BBN) can deliver their rated 155 Mbps speeds. some high-end 622 Mbps switches have failed from a lack of adequate flow control. ATM switches are really designed for steady, not bursty, traffic. But data communications is bursty, so they drop bits all over the floor.[2]

Some analysts are even more cautious. In his presentation on video servers, Al Lill of the Gartner Group proclaims that ATM will not be a viable LAN architecture until 1997-98, but sees its importance down the line:

> While the core of most segment and virtual LAN switches is proprietary today, we anticipate that ATM will become the dominant form of switching, due to the scalability and extremely high performance that will be required for LAN hubs to support hundreds of users at full wire speed transmission rates...
>
> But ATM to the desktop will not happen until 1997-98.[3]

---

2. In his book on ATM, Asynchronous Transfer Mode, Solution for Broadband ISDN, Martin de Prycker gives this subject a completely cursory overview (Section 3.2.5). By simply stating that "In principle, no flow control will be applied to information streams at the ATM layer inside the public BISDN network," de Prycker simply leaves the discussion to those who design network interface modules.

1. McQuillan, John R. Ph.D., "Why ATM?", Business Communications Review, February 1993, Page ATM 1.

2. Reinhardt, Andy, "Building the Data Highway", Byte Magazine, March 1994, Page 66.

4.2 Moving Forward Cautiously

Silicon Graphics agrees to a certain extent with the less enthusiastic of the crowd. This does not mean that Silicon Graphics thinks that ATM will not prove viable. However, following several years of experience with high speed, high performance networking, Silicon Graphics understands that the issues encountered in wholesale replacement by a new networking technology are not simple. In fact, we are very excited to explore higher bandwidth networking possibilities, but we want to be careful not to make promises which the current state of the technology cannot support.

We also realize that customers of Silicon Graphics need more network performance than those who purchase other equipment. IRIS workstations and Challenge servers achieve higher bandwidth and provide more throughput to service the strong needs of graphics, supercomputer, and multimedia users. Therefore, we will begin to supply solutions today for target markets.

It should be noted here that the reservations expressed by Mr. Partridge above are very real. Flow control issues, signalling, and lack of LAN emulation are real problems for ATM today:

> A thorough specification for congestion and flow control is essential to ensure that traffic from user nodes does not exceed the capacity of ATM networks. If traffic flow from user nodes is not regulated, the memory buffers of those switches can overflow, leading to data losses.[1]

Warning! Data Loss May Occur

Local Area Networks function very poorly when a high probability of data loss exists. NFS traffic in particular shows massive performance degradation when data loss occurs. What this means is that early adopters will still experience start-up pain. Those who do choose to adopt ATM should look for switches with very large buffers. More importantly, those customers with bursty data needs should continue to use FDDI and other well understood LAN technologies to meet those needs.

4.3 Standards and Functionality

Since the market is still young, many vendors exist with different types of solutions. Fore Systems has been particularly aggressive in providing solutions for the Silicon Graphics Indigo™ family. Recently, they have expanded their product offerings to include VME host adapters for CHALLENGE™ and Onyx™. These hosts adapters provide the features listed in Table 6.

The initial rollouts of ATM from all vendors include the basic functionality of SONET/SDH or TAXI interface and permanent virtual circuits. These features, while useful for testing, do not provide a full featured implementation. Switch

---

3. Gartner Group Presentation: Video Servers. March 3, 1994, Pan Pacific Hotel, San Francisco.

1. Saunders, Stephen. "ATM Forum Ponders Congestion Control Options", Data Communications Magazine, March 1994, Page 55.

vendors like Fore and Digital Equipment Corporation (and many others) have created their own LAN emulation to provide basic local area network services to the users. This creates problems for inter-operability in heterogeneous networks in addition to the complexity already associated with an emerging protocol.

Silicon Graphics takes the position today that we should support standards as they become available. Therefore, we will implement to the SONET/SDH OC-3c/STM-1 for the first release using multimode fiber. The only standard switching protocol, the ATM Forum Switching Protocol, will become available on our system in late 1994.

RFC 1483 suggests two implementations of IP encapsulation and both are important. At this point, it is impossible to know exactly which method to support. Further implementations of ATM will address protocol enhancements to support flow control and IP. As ATM becomes more readily available, Silicon Graphics will support the ATM Forum standards.

4.4 The Fore Implementation

Fore Systems has chosen to pursue ATM aggressively, selling switches and host adapters for local area networks. Fore provides GIO-32, EISA, and VME bus implementations thus covering the Silicon Graphics product line. These host

Feature	Silicon Graphics	Fore Systems	Newbridge
Host Adapter for Indigo Family		X	Indy
Host Adapter for CHALLENGE and Onyx	X	X	
VC Based Multiplexing	X		X
LAN Emulation		X	X
ATM Forum Signalling Protocol		Planned	Planned
Flow Control			
OC-3c/STM-1	X	Planned	X
TAXI		X	
On Board RISC Processor with Local Memory	X		Processor
On Board Rate Control Queues	X		
Availability	N/A	Now	Sep. 1994

Table 6: Features of Silicon Graphics and Fore Implementations of ATM

adapters and switches will function together as a local area network because Fore has chosen to develop and support their own methods for handling IP encapsulation. Fore has not implemented rate control. Fore's systems have been increasing in their performance over time and provide a reasonable alternative for the customer looking for ATM in the local area. Fore's trade name for their complete network emulation package is SPANS.

No one should expect, however, that systems using Fore's adapter cards will inter-operate with systems from all vendors. Limited inter-operability, based on standards, exists for PVC traffic today, but full LAN emulation requires the user to have a network of Fore switches and host adapters only.

4.5 The Newbridge Implementation

Newbridge Microsystems has also developed a GIO implementation for the Silicon Graphics Indy and will begin shipping sometime around September of 1994. Like Fore, Newbridge has developed a complete system with adapters for many different types of workstation. Newbridge provides a complete package called "VIVID" for VIdeo, Voice, and Image Data. Newbridge has also chosen to integrate their offering into existing networks by providing the "Ridge" product, an FDDI to ATM UNI, or a 12 port Ethernet switch with an ATM UNI. These products help the user in transition to ATM by not taking away the traditional local area network.

4.6 The Orlando Implementation

Silicon Graphics has taken part in the most aggressive use of ATM today in Time-Warner's Orlando Full Service Network project scheduled for deployment this summer. This interactive television trial takes networking into the customer's home, a difficult environment for computer networking. For the first time, users will have access to a high bandwidth, bidirectional (though not symmetrical), asynchronous network.

Using this network, Time-Warner will deliver high quality interactive services to 4000 homes using Silicon Graphics technology. The network in this implementation will deliver ATM to each home on a PVC network with OC-3 and DS-3 rate channels. The return channel will use IP over a TDM channel returning back to the ATM network and running IP over ATM back to the host. The network will include both fiber optic cable and coaxial cable so ATM will have to function on different physical carriers as well as at different data rates. In addition, the delivery of movies will show ATM delivering multimedia traffic while interactive data transfer will show ATM carrying critical, reliable data. This network will demonstrate several of the unique features of ATM as well as its ability to exist in a standard network.

Silicon Graphics and Time-Warner chose to use ATM as the primary broadband protocol because it is an emerging wide area network standard which will provide the services desired. Other interactive television implementors are trying to use channelized DS-3, DS-1, X.25, and Asymmetric Digital Subscriber Line (ADSL) for delivery and response, but these choices reflect the dependence these vendors place on Video-on-Demand as the key application. Video-on-Demand will work in a synchronous network because one can deliver a compressed movie on slots in a synchronous, time-division multiplexed channel. ATM, with its ability to use various amounts of network bandwidth, allows efficient use of network bandwidth and the ability to program and use services in a natural and flexible manner.

It is important to note that Silicon Graphics still considers this technology to be prototype and has no immediate plans to release this ATM technology as a general purpose product.

Section 5 Conclusion

Asynchronous Transfer Mode promises a new generation of networking. Optimized multimedia transmission, high bandwidth, and switched services all bring significantly more power to the network than today's networks can provide. ATM can realize this promise, and will over time. Several key standards still need definition and implementation, but some of the qualities of ATM are available today.

In 1994, Silicon Graphics recognizes that early adopters are beginning to deploy ATM for networks beyond a typical test bed. These adopters have met with some success in the deployment of wide area networks using permanent virtual circuits. The resulting network looks very much like a network with dedicated pipes interconnecting LANs across wide areas. In collaborative engineering design across continents, where engineers must share data in real-time, ATM has allowed the companies involved to dramatically shorten engineering schedules by allowing network users to share ideas directly. In medical imaging, ATM has given doctors access to libraries of remotely stored data, cutting the need to bring patients to the hospital for consultation. In supercomputing, ATM has provided access to shared supercomputer technology allowing precious resources to be spread amongst a much wider audience.

Each of these examples shows that ATM technology has delivered a portion of the promise. Typically, these early implementations have barely begun to access the power of the protocol. None of them yet use switching, nor have they generated a wide enough customer base to test advanced congestion management techniques. This will happen this year, however, and ATM will move to the mainstream of wide area networking. The future is more hazy when it comes to the local area network market, but the driving force multimedia will create places heavy demands on current technology. In wide area networks, such as interactive television delivered to the home, ATM will work well. It is still the case, however, that local area networks, on which the majority of the traffic is bursty, should be built on known technologies: FDDI and Ethernet.