211-TP-004-001

Asynchronous Transfer Mode (ATM): A Comprehensive Overview

Technical Paper

Technical paper—ECS Project only — Not for Distribution

March 1995

Prepared Under Contract NAS5-60000

RESPONSIBLE ENGINEER

Chuck Pross /s/ 3/7/95

Chuck Pross, Sr. Systems Engineer Date
EOSDIS Core System Project

SUBMITTED BY

Peter G. O'Neill /s/ 3/7/95

Pete O'Neill, SI&P Office Manager Date
EOSDIS Core System Project

Hughes Applied Information Systems
Landover, Maryland

Abstract

This document is intended to provide a comprehensive overview of the Asynchronous Transfer Mode (ATM) development status and position in the marketplace.

Keywords: asynchronous, transfer, mode, bandwidth, networks, connectivity, interface, LAN, WAN, development.

This page intentionally left blank.

Contents

Abstract

1. Introduction

2. Background and History

3. ATM Characteristics

4. How Will ATM Evolve?

5. Status of Standards and Availability of Products

6. ATM as a WAN

7. ATM as a LAN

8. Sample ATM Product Plans

9. Summary

Appendix A: Fore Systems Plans

Appendix B: References

1. Introduction

1.1 Purpose

This document is intended to provide a comprehensive overview of the Asynchronous Transfer Mode (ATM) development status and position in the marketplace. ATM is one of the fastest growing and evolving technologies. It is therefore important that the reader realize that while the information presented here was current at the time it was collected from its original source, portions of it may be out-of-date by the time it has been published and read.

The material in this report comes from a variety of sources: periodicals, vendor publications, and vendor presentations (Appendix B). Some of the information may appear contradictory--that is because the technology is very recent, and consensus has not yet been reached either by standards groups, vendors, or users.

Much of the material in this report consists of edited verbatim copies of material from referenced sources. The material presented was cross-checked with unreferenced material from other sources, for consistency, and as a "sanity check."

1.2 Organization

This paper is organized as follows:

Section 1: Describes the document purpose and organization.

Section 2: Provides the background and history of the subject.

Section 3: Describes ATM characteristics and qualities.

Section 4: Examines the current and future evolution of ATM.

Section 5: Provides the status of ATM standards and product availability.

Section 6: Examines ATM as a WAN.

Section 7: Examines ATM as a LAN.

Section 8: Provides sample ATM product plans.

Section 9: Provides a document summary.

At the end of this document are two appendices: Appendix A provides two vendor’s strategies and products; Appendix A lists references used in this document.

Questions regarding technical information contained within this Paper should be addressed to the following ECS contacts:

– Chuck Pross, Sr. Systems Engineer, (301)925-0716, chuck@eos.hitc.com

Questions concerning distribution or control of this document should be addressed to:

Data Management Office
The ECS Project Office
Hughes Applied Information Systems
1616 McCormick Dr.
Landover, MD 20785

2. Background and History

Over the last decade, the role of LANs has changed considerably. Initially conceived as a data highway with seemingly infinite bandwidth, the LAN was installed as a pipe within a facility into which each user could tap, much like a water main into which homes tap for water.

The size of the pipe was huge, typically 10 megabits per second (Mbps) compared to the data rates supported by attached devices, typically 9.6 kilobits per second (Kbps). In the early days of LANs, it was inconceivable that the pipe would run out of capacity.

Within a few years, however, the numbers and speed of attached devices taxed LANs to the point of requiring network segmentation. Typically, the network was divided into logical workgroups, which were segregated from the network backbone by a bridge. In most cases, the backbone was no faster than the segments it connected; however, the perceived performance of the network increased because the bridges filtered traffic that was local to each workgroup. In the past few years, LANs have been placed within a box, known as a concentrator, with a logical network backbone running on the concentrator's backplane.

Despite this evolution, LANs continue to run out of bandwidth, and the problem is getting more acute because of the power and proliferation of computing devices attached to LANs. The issues driving LAN growth boil down to: the economics of client-server computing versus terminal-host computing, the persuasiveness of the computing technology (PCs and workstations), the desire to share data throughout the organization, and the sheer power of the workstations.

The power of workstations (in millions of instructions per second) has risen by an order of magnitude in the past three years and is likely going to rise by another in the next three years. The power of the workstations is outstripping the capabilities of the networks to which they are attached. This problem is particularly severe in scientific and engineering applications.

A number of tactical solutions have emerged to address the LAN bandwidth problem: so-called Etherstretch solutions in the Ethernet world. One solution is the advent of the multisegment concentrator. Rather than placing a single LAN segment on the backplane of a concentrator, the leading vendors today are placing many segments on the backplane.

Another tactical solution is the multiport media switch, such as the Ethernet switch. An Ethernet switch provides dynamic circuit switching between Ethernet segments. Such approaches are good short-term solutions, but in the long run they are architecturally doomed compared to the inherent power of an ATM-based LAN.

A third solution is the deployment and implementation of higher speed backbones, such as FDDI backbones connecting lower speed Ethernet or token-ring segments.

The fundamental problem with all these techniques is that they continue to deliver <u>fixed amounts of bandwidth</u> to networks supporting more powerful and larger quantities of workstations. Ethernet, Token Ring, and FDDI all have a fixed amount of bandwidth that must be shared among all users on the segment. As users are added, the bandwidth available to each declines proportionally.

While all these approaches, or a combination of them, may stem the tide temporarily, they fail to offer an innovative long-term architectural solution to the LAN bandwidth problem. ATM-based campus networks, by contrast, offer a radical architectural departure from the past and promise to offer a permanent solution to the LAN bandwidth problem.

<u>The ATM Solution: Network Scaleable Bandwidth</u>

ATM is a type of packet-switching technology. The topology of ATM-based LANs is a two-tier star network. At the network center is an ATM cell relay switch. A high speed user interface branches out to the attached LANs. This interface has been defined by the ATM Forum, a technical consortium of companies charged with hammering out the user equipment-to-network interface. Typically, a router or bridge attaches to an ATM switch over this interface. Routers provide a migration path from the existing LAN infrastructure to ATM. For a solution such as this to be effective, users will have to use high-speed second-generation routers, capable of routing upwards of 100,000 packets per second. In the longer term, workstations will connect directly over the user interface.

Adaptive Corp. has coined the term, the unlimited area network, to promote the salability of ATM in LAN applications. ATM provides a dedicated amount of bandwidth between each user (or community of users) and the ATM switch. The key to ATM is plug-in bandwidth: additional high speed connections to the switch can be added without degrading the performance of existing switch connections. As additional users are added to the network, they can connect to a dedicated ATM port, either directly or through a router or bridge.

2.1 Managing Moves and Changes

The restructuring of businesses today (their expansion or contraction)has placed a large onus on network administrators. A major problem reported by these administrators is managing moves or changes. For example, splitting a workgroup on a TCP/IP-based LAN causes these problems:

- The administrator needs to create a new IP subnetwork.
- Each node in the relocated department needs to be assigned a new IP address.
- The existence of the new subnetwork needs to be configured into the network's routers.
- Security and access between the old group and the relocated group need to be resolved.
- Other departments may need to be restricted from accessing the new subnetwork.

These problems are not unique to TCP/IP. They exist in some form for all LAN-based networking protocols.

The solution: Virtual circuits between LANs. In an ATM network, LANs are attached using a concept called a virtual circuit--a logical pipe between two LANs. Devices attached to the LANs are unaware of the ATM switch between them. A department can be split into two geographical locations through an ATM switch without reconfiguring the network addresses on each segment. ATM allows the logical network to be separated from the physical network.

The granularity of moves and changes can be brought down to the individual device. A single device can be attached to the ATM switch over a virtual circuit going to one or more LANs[1]. By separating the physical locations of devices from their network addresses, the LAN--administrator can create a logical network with little resemblance to the physical.

2.2 Backbone Congestion

LAN backbones were created to relieve congestion within workgroups. When workgroup congestion rises to the point of noticeable network delay, a LAN can be reconfigured by installing bridges. As the size of LANs increases, the backbones themselves become congested, requiring the installation of a second level of backbones. As the layer of backbones increases, there is a diminishing return due to each frame's need to traverse a large number of bridges.

With a backbone-based network, the backbone and bridges act collectively as the switching fabric between individual LAN segments. Typically, backbones offer the same bit rate as the segments they attach; for instance, Ethernet segments are usually connected by an Ethernet backbone. In this case, the speed of the backbone often becomes the limiting performance factor. In other cases, the backbone is faster than the segments. Many token-ring networks use 4Mbps segments and 16Mbps backbones. A few LANs use FDDI backbones. At 100Mbps, FDDI is about one order of magnitude faster than the segments it is connecting.

Today, some larger LANs are using the concept of a collapsed backbone; the collapsed backbone is a number of workgroup segments with a router as the hub. Collapsed backbone networks suffer from two problems: the speed of the router and the speed of the media. For these network designs to be effective, the hubbing router must be capable of supporting full bandwidth routing--routing of all ports concurrently at media speed. Presently, only second-generation routers are capable of routing at media speeds on all ports.

The second major limitation of collapsed backbone topologies is the speed to the medium itself. Even with a high-speed router, the network is limited to 4Mbps, 10Mbps, or 16Mbps, and more importantly, the media is fixed in bandwidth. ATM addresses this problem in a nonintrusive way by eliminating the backbone. This is done by consolidating all the traffic through an ATM switch. The ATM switch becomes the switching fabric between the LANs. An ATM-based LAN offers (and this is significantly different than backbone networks) a switching fabric capable of supporting bit rates two orders of magnitude faster than the segments it is connecting. A well-designed ATM-based LAN will provide a large growth path, and hence a long economic life.

---

[1] (In the short term, it is unlikely that an individual device will be connected this way because of the costs involved in dedicating an ATM port switch to an individual device and because most devices today do not need 100 Mbps of dedicated bandwidth from a switch. They can share the ATM connection with a group of devices through a bridge or router)

2.3 Infrastructure Diversity

An important issue facing wide area network (WAN) architects today is the diversity of infrastructure. In Fortune 500 corporations, it is common to have one WAN for voice, another for SNA traffic, and possibly a third for LAN traffic. These separate networks have grown largely from the understanding that each network user community has different needs. The advent of video-conferencing has created another set of needs, bandwidth-on-demand, and its own subindustry of inverse multiplexer vendors.

When faced with the diverse (and diverging) set of requirements, network architects have had two fundamental choices in the past: create duplicate networks or consolidate onto a single infrastructure. When networks with diverse needs are consolidated, the fit for any one of them is poor.

For instance, one prevalent application today is the sharing of voice and LANs over a T-1 WAN. In this example, the two needs share the same media but are not highly integrated. The T-1 bandwidth is allocated semi-statically, yet the needs of the LAN are often bursty in nature. One instant the LAN cannot get enough bandwidth from the T-1 link; the next instant the link is idle.

Because of its versatility, one promise of ATM is to create a single infrastructure to handle all these needs and let network planners consolidate these disparate WANs into a single WAN.

Part and parcel of solving the infrastructure problem is tightly integrating LANs and WANs with the secondary benefit of consolidated network management. Again, because of its versatility, ATM will integrate multiple applications transparently.

The ATM solution: Consolidate the infrastructure. Solving the disparities of infrastructure is a network design problem in disguise. As a basic transport technology, ATM has been designed to handle voice, video, LAN, data, and fax--all within a single switching fabric. It is standards-based and architecturally flexible enough to handle this diverse set of requirements. Most importantly, ATM is a new standard and as such has had the benefit of prior industry experience in solving complex technical problems.

What makes ATM a viable solution? The small cell size of 53 octets, for instance (an octet is an 8-bit byte), makes the standard usable for both voice and data. The small variability of network delay makes ATM a viable solution for video and voice applications. Additionally, ATM provides nonblocking access, allowing it to be used for video applications. ATM will support permanent and switched virtual circuits, allowing it to connect the LAN and data infrastructure. Because ATM is a fixed-length cell technology, ATM switches can switch millions of cells per second, letting the ATM network carry all the traffic.

Finally, because ATM is a consolidating technology, it benefits from the law of large numbers. It is statistically unlikely that all users will demand network services at any particular instant; the network appears to be available instantly to all users at all times.

3. ATM Characteristics

3.1 Cell-Switching Technology

ATM s a cell-switching technology capable of concurrently carrying voice, video, data, and facsimile. It has been designed to provide a service for any intermixed combination of traffic with relative ease. With current leased-line technology, such as a T-1 network, data and voice can share the link, but they are separated into different channels within the T-1 stream. ATM, by contrast, intermixes traffic freely, based on the dynamic demands of the network users. LAN traffic, for instance, tends to be bursty in nature. During a LAN burst, the ATM network can dynamically provide additional bandwidth to carry the load. When the burst ends, the bandwidth is available to other users.

Although it is possible for all users to demand a large amount of bandwidth simultaneously, in large networks it is unlikely that this would actually happen. <u>However, ATM does not currently support congestion control mechanisms to handle this event.</u>

All ATM traffic resides in a fixed-length 53-octet (an octet is an 8-bit byte) cell. Fixed-length cells were chosen to facilitate very high-speed switching (millions of cells per second), Forty-eight octets of the cell carry a payload (that is, user data). Five octets are for overhead. A 48-octet payload was chosen as a compromise between the needs of voice, which prefers small cells with short latency, and data, which prefers a better ratio of user data to overhead and therefore a larger cell.

ATM's key qualities include its isochronous nature (ideal for voice), which has low network delay, and more importantly, low variability of network delay: all cells arrive within a predictably fixed time. For data, ATM was designed to support nonblocking transmission through the concept of virtual tributary streams: the network guarantees that each user will have a minimum bit rate service level. The user can exceed this level with the understanding that excess frames may be deleted during periods of network congestion.

4. How Will ATM Evolve?

4.1 From LAN to MAN to WAN

While it is not clear how any new technology will evolve, it is likely that ATM initially will be used in campus LAN environments. Limited network bandwidth as well as moves and changes are two urgent problems facing LAN administrators today. With no foreseeable slowdown in the power of LAN-based computing systems, the need for ATM in campus environments is immediate and real.

ATM will quickly migrate from the LAN, out of the metropolitan area network (MAN), and finally to the WAN as ever-increasing concentric circles. Within the LAN revolution, however, several phases will occur. The first of these is the migration phase. ATM will be used in conjunction with high-speed routers and bridges to provide LAN connectivity. Most first-generation routers cannot adequately support the speeds required to achieve effective throughput on ATM. Their basic problem is a mismatch between routing speed and ATM technology; their meager performance is ill suited for connection to the ATM infrastructure.

Now however, a new generation of routers is available. These so-called second-generation routers, multiprocessor RISC-based machines supporting very high-speed routing, are well positioned to connect existing LANs to the ATM infrastructure. The second phase will be the direct connection phase. As the cost to connecting directly to an ATM switch declines (both on the workstation and ATM switch ends), workstation vendors will offer direct ATM interfaces. Initially, these interfaces will be offered only on high-end workstations. In time, however, the technology will commoditize, and many workstations will provide direct connections to the switch. Regardless of the speed with which this happens, however, administrators will always need routers to connect the existing LAN infrastructure to the switch --in the same way that hundreds of thousands of terminals of yesteryear attach to the LAN infrastructure with terminal servers.

From LANs to MANs. From LANs, the ATM infrastructure will grow outward to the MAN. Fortune 100 companies will grow SONET-based private networks to connect their sites within a MAN. Over time, many of the bypass and local carriers will pick up on this cue and offer ATM services in major metropolitan areas, just as they offered T-1 services following a wave of private network installations during the 1980s.

ATM over WANs. Finally, ATM service is available over some WANs. Such application of ATM technology requires a massive investment in the WAN infrastructure. However, in time, all local exchange carriers and interexchange carriers will be forced to migrate to ATM to remain competitive.

<u>The most imminent application of ATM is in LANs</u>. Although commonly thought of as a wide area technology, ATM will change the face of LANs within the next three years. The

technology--and some of the products now available to support it--offer a compelling argument as to why ATM can and will play a role in next-generation LAN architectures.

ATM will find its first home in the LAN, primarily due to the urgency of the problems facing LAN administrators. The growth of the LAN, coupled with the explosion of LAN-attached computing power, has created a spiraling requirement for LAN bandwidth. Because of their static bandwidth, the current media-based solutions--Ethernet, Token Ring, and FDDI--offer only a temporary reprieve from the tidal wave of demand.

Additionally, ATM can solve the problem of managing continual moves and changes. Because ATM separates the physical and logical networks, physical location can be completely unrelated to a station's' network address.

ATM also solves backbone congestion by collapsing the backbone into a single, ultrahigh-speed switching fabric. Unlike a backbone, which is typically no more than ten times the speed of the LANs it is connecting, an ATM switch is 100 times faster than the attached LANs.

Finally, ATM solves the network design problem. ATM is a single, well-conceived architecture, based on a wealth of experience accumulated in previous-generation networks.

5. Status of Standards and Availability of Products

Public carriers do not question whether ATM switching offers a number of benefits in wide area networks. Using ATM as a switching backbone will allow them to provide economical high-speed LAN interconnect, data, voice, video, and multimedia services to their customers. The long-haul companies, such as AT&T, MCI, Sprint, and WilTel, have led the charge, and the Regional Bell Operating Companies (RBOCs) will offer ATM as well. Competitive access providers (CAPs), such as MFS Datanet, are also making such services available to customers throughout the United States.

And wide area networking isn't the only arena in which ATM can offer significant advantages. Major LAN equipment vendors are also investing heavily in ATM, anticipating that a huge ATM LAN market will develop. Given the high level of interest among the user community at this early stage, the vendors' assumptions are probably correct.

However, ATM is not a networking panacea. ATM standards, the source of so many potential benefits, represent a compromise: ATM isn't the ideal solution for data, voice, or video. Existing LAN applications can't fully exploit ATM technology, and ATM-specific applications don't exist yet. Although some ATM switch vendors provide proprietary APIs for their products, there is no ATM standard for APIs.

Moreover, several key ATM specifications are still under development and need a lot of work. And the critical issue of cost-effective migration from current LAN technology to ATM remains an open question.

5.1 Using ATM Switches

Still, design and development of ATM products is going forward at an amazing pace, and many users are actively planning their ATM networks.

ATM switches fall into two major categories of use: private enterprise switches and public WAN backbone switches. The later category is the domain of very large WAN switches designed for public carriers' central offices and serving as nodes in nationwide and eventually global public wide-area backbones.

Somewhat smaller versions of these switches will also be deployed by public carriers to provide access points to the ATM backbone. These access points will support other transmission technologies and services such as T-1, frame relay, and Switched Multimegabit Data Service (SMDS). These products will have limited availability over the next two years and will primarily be used in the carriers' ATM trials and early WAN service offerings.

Within the broad category of private ATM switches are three somewhat different applications of the technology: workgroup, campus or enterprise backbone, and WAN access. ATM workgroup switches are designed to be used in conjunction with ATM adapter cards in end-station devices such as desktop computers and workstations. These switches can also be linked to form a campus ATM backbone. Eventually, they will also be used as WAN access switches, linking private campus networks to public carrier ATM services.

Fore Systems' (Pittsburgh) ASX100, SynOptics Communications' (Santa Clara, CA) LattisCell ATM switch, Network Equipment Technologies' (NET, Redwood City, CA) ATMX, and Newbridge Networks' (Herndon, VA) Vivid ATM Workgroup Switch are examples of this type of product.

5.1.1 Workgroup ATM

ATM workgroups, with ATM to the desktop, will see limited deployment over the next two to three years. Other technology already in use will continue to be adequate for supporting most desktop applications, and most companies will not be able to justify the expense of delivering ATM bandwidth to the desktop.

So is anyone really using ATM to the desktop yet? A small number of installations are. Many tests are underway, but production applications are few and far between. Fore Systems, the current leading ATM switch vendor, boasts more than 500 switches installed at more than 300 customer sites. Many of these switches are being used in trials of one kind or another, but some of the sites are using Fore's ASX-100 switches to support workgroup networking applications on a production basis.

Applications for switches from Fore Systems, NET, Newbridge, and others include:

- linking clusters of high-performance workstations processing in parallel to form a less-expensive replacement for supercomputers;
- scientific visualization or the visual modeling of complex structures and processes; and
- workstation-based video conferencing.

Applications such as these that require high-bandwidth links between workstations are not common yet. However, with ATM technology now available, engineers are beginning to imagine previously undreamed-of possibilities, and more high-end applications will be developed over the next few years. This development, in turn, will increase demand for ATM to the desktop.

5.1.2 Enterprise ATM

Desktop applications will drive demand for ATM in the future, but for now, bandwidth on the backbone is the biggest selling point. People are looking to ATM to relieve congestion on campus network backbones and to make those backbones easier to reconfigure and manage. Vendors are addressing these issues by incorporating ATM ports into intelligent hubs and routers, allowing these devices to use ATM's high bandwidth to increase backbone throughput. Other functionality that can be provided by ATM on the backbone, such as virtual networking, will also be very attractive to network managers.

Several vendors are building frame-to-ATM cell-conversion modules, called segmentation and reassembly (SAR) devices, into their hubs and routers, transforming them into gateways between legacy LANs and campus ATM backbones. A SAR device takes variable-length Ethernet, Token Ring, or Fiber Distributed Data Interface (FDDI) frames and segments them so they can be placed into the 48-byte "payload" portion of fixed-length ATM cells.

The SAR also reassembles frames from pieces carried in cells coming into its ATM port. Most vendors call the SAR devices they incorporate into their products by a trade name, such as SynOptics' EtherCell, or simply as an ATM interface on their hub or router product.

Using a SAR device in a hub or router doesn't turn it into a switch, however. An ATM switch by definition has more than one ATM port (usually many). Hubs and routers with ATM interfaces are designed to connect into ATM switches that will route their traffic through the ATM backbone.

Putting ATM interfaces on intelligent hubs and routers gives users one of the keys to migrating from existing shared-media LAN technologies to ATM in manageable stages. However, putting a Data exchange Interface (DXI) on a router or making a new router with a full frame-to-cell conversion module solves only the physical-layer connection problem. It doesn't touch the more complex problems of multivendor interoperability or the integration of ATM and legacy LANs. In spite of intensive work by the ATM Forum (Mountain View, CA), a consortium of equipment manufacturers, service providers, researchers, and users, many of these problems remain largely unresolved.

Direct Bearing:

The Forum's technical committees publish prestandard specifications, based whenever possible on existing international standards. The Forum's specifications, while not approved international standards, are complete enough for vendors to implement in ATM products.

As with most communications technologies, ATM is undergoing a development evolution that began at the Physical-layer protocols and is proceeding upward to the higher layers. The Forum's Physical Layer Working Group has one specification under its belt and is working on others. In the interim, vendors are building products using existing interfaces.

The Forum is focusing on three areas: standardizing as many of the interswitch signaling and connection-management functions as possible, which will eventually make multivendor switch interoperability possible; defining how Layer 3 protocol routing used in legacy networks (for example, IP or IPX routing) will interact with ATM's native routing capabilities; and detailing how ATM switches will be managed in a network that includes legacy LAN equipment.

Connection Management:

Connection management refers to a group of functions, including call setup and tear down, call routing, address resolution, and management of existing circuits. Eventually, a number of standards will specify how these functions should operate, but most are still in development.

Some specifications have been finalized to the point that vendors are designing products that implement them. Where complete specifications are not yet available, vendors are using their own proprietary connection management protocols, which means limited interoperability.

The placement of connection-management software is also an important factor in switch design. Some vendors are choosing to place connection-management services in a centralized server somewhere in the network, while others have it distributed to all the switches in the network. Both approaches have their benefits.

Centralized connection management results in lower switch cost because the switch doesn't need an additional processor to run the software and because this design makes upgrading to new versions easier. Communication with the centralized control server may be in-band, using part of the network bandwidth, or out-of-band, using a separate connection between server and switches. Out-of-band communication causes potential connection setup delay and requires more connections. It is also a potential single point of failure.

Distributed connection management can result in faster connection setup and tear down because decisions are made locally. Distributed control also eliminates a single point of failure. Some people argue that a distributed-control architecture scales better than one with centralized control. On the downside, distributing these services to the switches increases cost by requiring additional processing power in each switch.

Signaling:

Signaling is one of the key connection-management functions. The Signaling Working Group of the ATM Forum's Technical Committee recently completed version 3.0 of the user-to-network (UNI) recommendation. UNI 3.0 specifies signaling for SVC and PVC call setup and tear down; UNI 2.0 dealt only with PVC signaling.

The 3.0 specification defines addressing conventions for end stations in both private and public network nodes. On private ATM networks, addresses will be modeled after the Network Service Access Point (NSAP) format defined in Layer 3 of the OSI protocol stack. It is similar to an IP address but can be longer.

The NSAP format will yield globally unique 48-bit addresses that can function as Layer 3 addresses in a pure ATM network but can also be used as Layer 2 addresses by Layer 3 LAN protocols such as IP and IPX. Another advantage of this format is that an infrastructure and procedure to administer them worldwide is already in place.

With UNI 3.0, functional specifications exist for both PVC and SVC signaling. Using these protocols as specified, switches from different vendors can exchange addresses and share call setup and tear-down information (assuming both vendors build products to the same version of the specification).

UNI 3.0 also specifies using SNMP for local connection-management services across the UNI, including the definition of an ATM UNI Management Information Base (MIB). The Interim Local Management Interface (ILMI), as it is called, will use Simple Network Management Protocol (SNMP) messages to communicate local signaling management information across the UNI.

UNI 3.0 is a big step forward for standardizing key elements of ATM signaling. However, several issues in the area of signaling remain unresolved, and much more work needs to be done before users can purchase and use ATM switches as easily as they buy Ethernet or Token Ring hubs.

With UNI 3.0, two different vendors could set up and tear down circuits between their switches, but they wouldn't be able to fully negotiate quality of service for the connection, for example. This means that even if vendors support UNI 3.0, they will likely use their own proprietary protocols to cover elements of signaling and connection management not addressed by the specification.

According to Jim Grace, chair of the UNI Working Group and an employee of Ungermann-Bass (Santa Clara, CA), the group plans work on several fronts during the year. Their work includes issuing a UNI 3.1 specification that will bring current recommendations in line with the International Telecommunications Union-Telecommunications Standards Sector (ITU-TSS, formerly the CCITT) Q.2931 signaling specification.

This update means that ATM devices using UNI 3.0 and those using UNI 3.1 signaling will not interoperate. The working group will also develop what Grace terms "Phase 2" signaling, which will add several enhancement features to UNI signaling, such as the ability to have multiple connections per call (good for multimedia applications), more provisions for negotiating bandwidth and quality of service (QOS), and the ability to set up virtual paths through the network (the current spec defines only virtual channel setup).

The UNI Working Group is working in conjunction with the Private Node-to-Node Interface (P-NNI) Working Group on SVC signaling across the NNI. Such a standard will allow signaling interoperability between switches from different vendors. Any specification for interswitch communication will require two components: a signaling element and a routing element. According to Mike Goguen, chair of the PNNI Working Group and an employee of SynOptics, the signaling element is relatively straightforward. The routing element, which defines how routing will be accomplished in an ATM internetwork, is a far more difficult task.

The NNI working group is aggressively targeting 1995 as the year by which to complete a comprehensive Phase 1 NNI specification. The group is taking the broad view of signaling in large internetworks that may include both private and public switches in multiple routing domains.

However, a subgroup of members is also working on a draft "Phase 0" specification that will allow users to at least establish switch-to-switch signaling and routing between products from different vendors in a single domain. This specification will, Goguen concedes, use static routing tables that are far from ideal but will at least allow basic multivendor switch interoperability.

LAN Emulation:

LAN emulation for linking Ethernet, Token Ring, FDDI, and other LANs into and through an ATM network is another unresolved issue. The LAN Emulation Working Group has been working on a specification for some time, however, and is close to completion.

The committee took major steps toward completing the LAN Emulation specification which, according to Keith McCloghrie, committee chair and an employee of Cisco (San Jose, CA), should be finalized by late this year and published in early 1995.

LAN emulation is essentially the process of converting Ethernet, Token Ring, or FDDI MAC addresses into ATM addresses. But, since it is key to integrating ATM into existing LANs, a standard specification for LAN emulation is vital to ATM's acceptance by users.

5.2 Management of ATM Devices

Management of ATM devices is another critical element of ATM networking still under development. The ATM Forum specified SNMP as the management protocol for ATM.

The Internet Engineering Task Force (IETF), one of the standards bodies the Forum works with, is approving an ATM MIB, also known as the AToM MIB. This MIB includes information for monitoring and configuring both ATM switches and end stations and will allow standard SNMP management stations to monitor and configure ATM devices.

Vendors that incorporate this MIB into their switches will ensure that the switches are manageable by SNMP management stations. Users who plan to purchase ATM switches and use SNMP network management should then be able to use standard SNMP-based management stations, such as OpenView, SunNet Manager, and NetView 6000.

The larger issue currently looming in the network management arena is that the ATM Forum and IETF have standardized on SNMP. However, the ITU and the American National Standards Institute (ANSI) are working to adapt the Telecommunications Management Network architecture which is based on the Common Management Information Protocol (CMIP).

Effective management of an enterprise ATM network., which will almost always comprise both private campus and public wide-area components, will require this difference to be resolved. The current thinking is that SNMP will be used in private ATM networks and CMIP will be used to manage public networks, which means the two management protocols will have to interoperate.

We have a long way to go before we can feel as comfortable with ATM as we do with Ethernet and Token Ring. Legacy networking technologies will continue to exist like the city streets and access roads that stay in place as new superhighways are built to link them. But as ATM technology evolves, it will allow us to go places which we could never get to on the old roads.

Speeds:

ATM can be run at a variety of speeds over shielded and unshielded twisted pair (STP and UTP) and fiber optic cable. The Physical Layer Working Group of the ATM Forum's Technical Committee is currently working on definitions for specific speeds and physical interfaces.

The working group has recently completed a specification for 155.52 Mbps over Category 5 UTP. This speed is equivalent to the Synchronous Optical Network (SONET) STS-3c rate and uses SONET-like framing

The group is also near completion of a specification for 51.84Mbps (the SONET STS-1 rate) over Category 3 UTP at runs of up to 100 meters. That specification will use SONET-like framing. The key issue here is meeting various government agencies' emissions requirements.

In the meantime, most vendors are supplying their switches with a hodgepodge of already standard speed and interface options. Among those most commonly supported are 155Mbps over single or multimode fiber UTP using SONET framing, 45Mbps over coax with DS-3 framing, and 100Mbps TAXI (an interface that uses Advanced Micro Devices' TAXI FDDI physical layer chip) on fiber.

Switch vendors such as Fore Systems, NET, Ungermann-Bass, and Newbridge are looking to support speeds and interfaces that can be used to link end stations to switches as well as to link switches together in ATM LAN backbones. The standard high-speed interfaces they currently offer were originally designed for use in wide-area networking and don't lend themselves well to desktop connectivity. For example, most use fiber or coax, neither of which is ideal for supporting desktop applications. Only Fore Systems currently supports a twisted-pair interface.

Vendors such as Alcatel Data Networks (Reston, VA), Cascade Communications (Westford, MA), and Stratacom (San Jose, CA) make switches primarily to connect private ATM networks to public networks or to serve as carrier-owned public WAN access switches. The interfaces they provide, such as standard DS3 and E3, support existing transmission standards in wide-area networking and are better suited to their task.

In the long run, ATM switch vendors should standardize on twisted-pair interfaces for desktop connectivity and most local interswitch links. ATM switches used to connect local networks to ATM WANs will support standard wide-area speeds and interfaces provided by the carriers. In most cases these will follow SONET/SDH specifications for fiber.

6. ATM as a WAN

A lot of network savvy IS managers are hoping that they'll find the best magic wand in ATM. It's working out well for LANs, say early users, and now it's becoming available as a wide area service. But don't be fooled: wide-area ATM is a quantum leap more complex, and its immediate benefits are correspondingly more elusive than its local area predecessor.

Four carriers have announced interstate ATM services, and some are already in operation. The biggest, AT&T in Basking Ridge, N.J., will start general service by the end of this year. San Jose-based MFS Datanet Inc.; Sprint Corp. of Westwood, Kans.; and WilTel Inc. of Tulsa, Okla., are also in the ATM game. MCI Communications Corp. will be before the end of 1994, says Paul Weichselbaum, vice president of data marketing at the Washington, D.C.-based company.

Yet it's Hughes Aircraft Co. becoming one of the first commercial wide area ATM users. Like many of the other firms that will be conducting trials, Hughes has a vision of a single network for the enterprise, a unifying mechanism capable of making strategic changes in corporations.

Hughes has been testing ATM switches for use as local area backbone networks and is beginning a trial of a four-node WAN that has a 45-megabit-per-second access rate. For the trial, Hughes will use the WAN services of San Francisco-based Pacific Bell Co. (which will connect three California offices), as well as the services of Sprint (which will handle communications to Hughes Information Technology Co. in Reston, Va.) The trials examine such technical questions as how different applications work on ATM's variable and constant rate services, how workstations can set up sessions across the network and how well routers work on ATM networks. If the trial is successful, Hughes will be expanding its ATM network' to Denver and perhaps to some other U.S. cities this year, and to international sites in the future.

That's the vision, but early services are likely to be inflexible, far more of a headache to deal with than LANs. What's more, most users and consultants are far from optimistic that it will be cost effective in the near term. For the next 12 months, there will be limited use of ATM in wide area networks. ATM is an emerging technology, not one to bet the store on right now.

To some extent, confusion and inflexibility is inevitable with any early telecommunications service. But the differences between users' experience with local ATM and what they are likely to get with ATM WANs are quite stark. For example, unlike ATM LANs, for which there is only one category of service, users of ATM WANs will eventually need to ponder the differences among four Classes of service (A, B, C and D although only A and C are currently available[2]).

Furthermore, users of ATM LANs don't have to worry about the differences among committed information rates, burst rates, constant bit rates or variable bit rates. These details vary from

---

2 The specifics of these classes are beyond the scope of this paper.

carrier to carrier, and users not only have to ponder the pricing of all of these options, they also have to model their expected traffic in order to use services without wasting lots of money.

Cutting down on this planning problem, to some extent, is one of the most recently announced ATM services, WilTel's High Speed LAN Interconnection Service; unveiled last month. This service offers such nice features as the ability to reallocate connections and bandwidth quickly through an on-line management service. WilTel also allows circuits to be assigned asymmetrically, meaning that you can send large amounts of data in one direction without having to pay for an unused equivalent amount of bandwidth going in the returning direction.

Still another convenience is that WilTel allows its services to be accessed through inverse multiplexers, which it leases to organizations. These devices allow data to be sent at rates between T-1 and T-3 without having to pay for an expensive T-3 local access circuit. WilTel also offers a service specially tailored for connecting IBM mainframes to high-speed devices across IBM's standard 4-megabyte per second (32Mbps) channel connections, which in turn ride over ATM.

However, WilTel's two ATM-based services are so geared for particular types of data that they may not be very good for video or voice. Christine Heckart, manager of broadband services for the carrier, says that all services are the equivalent of variable rate Class C services. And Class C, most carriers insist, does not have the on-time delivery capabilities that video requires.

If WilTel's initial offerings seem unlikely to improve companies' ability to send integrated voice, video and data quickly and cheaply, it's not yet clear that any other carriers' early offerings will either. Yet cost reduction with ATM is far from certain. All ATM WAN services are sold on individually negotiated and confidential contracts, and only general pricing policies have been released. AT&T, MFS Datanet and Sprint have confirmed published reports that the price of Class C service for an average T-3 speed, three-node, trans-U.S. ATM network runs between $105,000 and $125,000 a month, not counting the cost of local access circuits. That's as much as 51% less than the cost of a circuit-switched T-3 network. Even discounting for overhead, there's still savings of as much as 35%. But that example doesn't explain how a hybrid of a Class C variable rate service and Class A's dedicated bandwidth would be priced. Furthermore, at distances of less than 600 miles, ATM is likely to be more expensive than dedicated circuits, estimates Richard J. Malone, a principal at the Vertical Systems Group in Dedham, Mass.

ATM's constant-rate services will have an advantage over Frame Relay because they can carry videoconferences and multimedia. But the earliest that users will be able to capitalize on that advantage is in about a year. That's when carriers are expected to implement the newly approved switched virtual circuit standard, which emerged from the ATM Forum in July. Then T-1 access circuits will be able to handle both Frame Relay and ATM and also be able to switch ATM video connections to the correct channel in the ATM backbone.

Despite this coming ability to share access circuits, it is misleading to say that ATM can "aggregate" data with video traffic on backbones. Both video and data pass through the ATM service, to be sure, but most experts say video must travel in Class A. While data also can travel in Class A, that's generally uneconomical, and so it must go Class C.

If Class A and Class C services could be set up and torn down instantly, on command of a user's application, that would be bandwidth on demand and integration of voice, video and data. But currently that's not possible. The upshot still is that, to use ATM for video or voice, a dedicated capability must be established. And when video or voice is not being sent on this facility, users pay for empty cells to be transmitted.

It's theoretically possible to be less wasteful by moving normally Class C variable bit rate data onto the empty Class A circuit when it's idle. But that would involve a kind of statistical multiplexing of cells by the service or by user equipment, and neither is yet available.

The current pricing and structuring problems of WAN services are likely to fade in a few years as carriers aggressively implement ATM. They're much more eager to leap into ATM than they were to move to the Integrated Services Digital Network (ISDN). That eagerness is augmented by the recent statements from several of the large cable TV firms that they are investigating ATM as the technology for distributing programs over telephone circuits.

A side from the most zealous partisans of the Asynchronous Transfer Mode, many potential users are hanging back from wide area ATM service, concerned about its expected pricing. That includes active members of the ATM interest group known as the Enterprise Network Roundtable.

Even representatives of Hughes Aircraft Co. and Motorola Inc., among ATM's strongest backers, indicate dissatisfaction with WAN ATM as it currently is offered, on a contract basis. Although most carriers insist that only constant bit-rate services, whose prices approach those of leased circuits, are suitable for voice or video services, Motorola sees no advantage in doing this and wants to try out variable bit-rate services to see if they can do the same job more cost effectively.

6.1 TCP/IP Not Ready For ATM-Like Speeds

First the good news: Asynchronous Transfer Mode has been proven to handle throughputs into the hundreds of megabits per second. Now the bad: for wide area high-bandwidth applications, that fact is meaningless because workstations with ATM connections generally use the popular Transmission Control Protocol/Internet Protocol. A recent test uncovered some serious problems with current TCP/IP products when they're used for high-speed data transfers over long distances.

Researchers at Sandia National Laboratories in Albuquerque, N.M., running a wide area version of some tests that had been completed successfully on a local area net, found that they were able to achieve at best an 8.5 megabit-per-second throughput in a file transmission over 1,700 miles using the full TCP/IP stack. This low throughput--comparable to a clear Ethernet link --occurred despite the fact that the data were transmitted across a 155Mbps ATM connection (which has a throughput of 135.6Mbps because of packet overhead).

The fault lies not with ATM. It worked well for voice and video in the test between Sandia and the Supercomputer '93 show in Portland, Ore. But for data, the problem lies with TCP/IP and its "window"--the amount of data that can be sent before receiving an acknowledgment. Currently that's 64 kilobits, maximum.

In the test, a 51 KB window was used. Putting that much data on the net took only 3 milliseconds. But transmitting it 1,700 miles and getting an acknowledgment back took 34msec because fiber-optic circuits have a delay of about 1 msec per hundred miles. At that rate, the circuit was being used less than 9% of the time.

The obvious way to solve the window problem is simply to use bigger windows, and work is underway on extensions to the protocol to do just that. However, new TCP/IP products able to use the proposed extensions are about a year away.

Larger windows mean larger buffers on workstations--as large as 3 megabytes for transcontinental use--and larger buffers may not be supported by applications. Many installed UNIX products default to 16 and 8KB. Microsoft Corp.'s new Windows NT defaults to an 8KB.

Ultimately, a whole series of changes will have to ripple through networks and applications before they'll be capable of exploiting the promise of long-distance broadband networking.

6.2 ATM WANs: What's Happening?

Real-world applications of asynchronous transfer mode will go on-line this year, following blessings from long-distance telephone providers Sprint Communications Corp., U.S. West Communications, and Bellcore.

The carriers late last year outlined plans to create a broadband infrastructure that can transmit data, along with audio/video information, at speeds ranging from 45M bps to 2.488G bps over non-dedicated lines that can be used to interconnect LANs. Previously, such high-speed bandwidth was possible only through dedicated leased lines.

Sprint was the first to announce availability of 45M-bps ATM connections at more than 300 network locations across the country. Customers who want this increased bandwidth must upgrade their routers' network interfaces, purchase an ATM data service unit, and order the 45M-bps access, Sprint officials said.

Although prices for the service range according to distance and usage, Sprint representatives in Washington estimated costs for a hookup from Seattle to Miami at between $51,000 and $62,000 per month. That doesn't include an installation charge of approximately $1,500 per location or the average $8,000 monthly local access fee per location.

Bellcore (Bell Communications Research Inc.), of Livingston, N.J., announced in October its specifications for RBHCs (Regional Bell Holding Companies) to provide ATM services to their customers. The specifications are intended to ensure consistent service nationally. Bellcore outlined how LANs and WANs should connect to ATM services, and how T-1 lines can be used to access these services.

Pacific Bell in San Francisco is one of the first RBHCs to offer ATM services according to Bellcore's specifications. The cost of its service in the San Francisco area includes a $5,000 installation fee and a $4,850 monthly fee for unlimited use of 45M-bps services.

7. ATM as a LAN

7.1 The LAN Environment

Over in the LAN environment, analysts say that new "virtual LAN" features of switched LANs are similar to ATM in the way that they deal with transforming switched connections into a LAN-like environment. Indeed, virtual LANS are "pre-ATM" both in being similar to ATM and in providing what seems likely to be a reasonably priced migration strategy for obtaining higher bandwidth networks before embracing ATM. If virtual LANS are something of a precursor to ATM, they will also be a necessity with ATM, because ATM's switched, connection-oriented protocol lacks the broadcast and segmenting capabilities of Ethernets or Token Rings. To emulate LANs, ATM must use multicasts and software-defined segments. And so, the ATM vendors' standards body, the ATM Forum, has a "LAN emulation" project under way, a project proposed by IBM to the group in February.

Much the same concept is embodied in one of the first ATM switches--the ATMX, from Network Equipment Technologies Inc. of Redwood City, Calif. It features a proprietary way of forming the equivalent of broadcast packets, multicasts and software-defined segments, which NET has called "a virtual LAN." Other ATM switch vendors have delivered some of the same capabilities under different names.

No matter what it's called, however, ATM needs a virtual LAN capability. Paradoxically, and usefully for a migration to ATM, some real LANS share the same requirement. That's because high-band-width switched LANs, which provide, for example, the equivalent of 10 megabits per second to each user (in contrast with shared Ethernet's mere 10Mbps for all users) cannot run in large configurations without software to recreate the segmenting capabilities of shared-bandwidth LANS.

Currently, the most conspicuous player in the virtual LAN cloud is Ungermann-Bass because of its many papers and free conferences on the subject. But UB is by no means the only hub vendor that sells virtual LANS. It wasn't even first, says Fred McLimans, program director for local area communications at the Stanford, Conn., research house Gartner Group Inc. The first hub with something like a virtual LAN, analysts say, was the PowerHub from San Jose-based Alantec Inc., which had a virtual LAN function a year before UB. Alantec, however, called the function "port subnet mapping," which is a little like calling sushi "cold, raw fish."

McLimans and other analysts say that the virtual LAN concept is considerably more. It can be a tool for quickly forming task-oriented work teams with shared resources and security barriers around their network. It can unite such teams and other workgroups regardless of their location within a building or a campus network (see diagram, "A Vision Of Virtual LANS"). When ATM comes along, it will be able to connect them across WANs.

Unlike Alantec, UB immediately and forcefully began building a public awareness of the virtual LAN concept when it unveiled what it calls its Virtual Network Architecture (VNA) in a

February announcement in which it also launched its latest hub product, the DragonSwitch line of switched Ethernet modules. An eight-port version began shipping last month, and a 16-port model is due by the fourth quarter. These convert UB's Access/One line of shared Ethernet hubs into switched Ethernets. UB also promises a switched Token Ring.

"The virtual workgroup concept is really important. UB has implemented it in this pre-ATM product, and it's coming to all the ATM products," says Nicholas J. Lippis IH, president of Strategic Networks Consulting Inc. in Rockland, Mass., and publisher of the "Internetwork Advisor" newsletter. "A tremendous innovation," comments James Herman, a principal at Northeast Consulting Resources Inc. in Boston.

The analysts give us special praise for recognizing that virtual LANS can play a role across wide area networks, as well as for setting up a joint development project with router maker Well-fleet Communications Inc. of Billerica, Mass. "We think this area has considerable potential," says Well fleet president Paul Severino.

Virtual LANS work on the basis of port-mapping, which means setting up tables of communications entitlements for each port. In contrast to port-switching, which uses software to physically place ports upon different internal buses within a hub (and requires bridges or routers to unite those buses), virtual LANs define segments purely with tables of permitted and not-permitted connections.

Despite the flexibility of virtual LANs, it's difficult to see users of shared LANS switching over to switched LANS simply for flexibility's sake. Gaining experience with a product similar to ATM is another possible motive. But since switched LANs still cost $800 to $1,000 per port more than shared LANs, it isn't much of a reason for swapping over many nodes.

However, that situation will probably change rapidly, analysts say, because switched LANs are likely to be the new competitive area for hub vendors, and there's apt to be a sharp decline in the cost in the next two years--down to around $300 a port, in McLimans' view. That would make switched Ethernets competitive with the new 100Mbps Ethernets, whose interface cards are expected to be similarly priced. It would also make switched Ethernets considerably cheaper than FDDI or ATM, which cost between $2,000 and $8,000 a node, depending on the type of workstation.

For these reasons, UB and other hub vendors are betting that the switched LAN market will soon break out of its present niche of merely providing bandwidth to power users and enter a broader plain, furnishing the high bandwidth and low latency needed by many users for video and client/server applications.

UB is currently lagging behind switched LAN vendors such as Santa Clara-based Kalpana Inc. and Artel Communications Corp. of Hudson, Mass. But UB is trying to parlay its virtual LAN technology and the functionality and ease of management that it brings into strong presence in the high end of the switched LAN market. UB has ATM product, but hasn't said when it would deliver it.

While UB has been a leader in pioneering virtual LANs, its ability to remain a stable force delivering this technology over the next few years is open to some question. Analysts say increasing competition in the hub market is hitting hard at UB's sales. A wholly owned subsidiary of Tandem, UB claims to have shipped the first intelligent Ethernet hub ever manufactured. But now, including the products made by UB'S subsidiary--the 51%-owned low-end hub vendor Net-Worth Inc.--UB accounts for only about 5% of the hub market, down from 6% in 1991, according to figures by the San Jose-based market research firm Dataquest Inc.

There have been reports that Tandem is considering selling UB, which Benoit strongly denies. "They've been very supportive of our restructuring," he says, adding that Tandem is "engaged in a vigorous search to find a replacement for Ralph [Ungermann]." Although UB did increase sales of hub equipment and its overall revenues last year, adds Stephen M. Diamond, UB director of marketing, the DragonSwitch and Virtual Network Architecture "continue to be critical" to the company.

But, other vendors are crowding into the virtual LAN market. This month Lannet Data Communications Inc., the Huntington Beach, Calif., subsidiary of Lannet Data Communications Ltd. of Tel Aviv, is scheduled to ship virtual LAN software for its 10Base-TV line of hubs. This, claims Lannet president Avi Fogel, will be the first shipped product that provides virtual workgroup segmentation across multiple hubs.

And, DATAMATION has learned, hub market leaders SynOptics Communications Inc. of Santa Clara, which will incorporate the Kalpana switching technology in its hubs, and Cabletron Systems Inc. of Rochester, N.H., which in late June announced a deal to use Artel's Ethernet switching technology, are working on virtual LAN products. ADC Fibermux Corp. of Chatsworth, Calif., expects to ship a virtual LAN box by year's end.

8. Sample ATM Product Plans

The following section will summarize the plans and products of a single representative ATM vendor. (The appendix contains very detailed plans and product descriptions of a second vendor, for those wishing to examine these issues in detail.) As with all companies, these plans, products and prices will change with time. Use the data presented here as representative only. It is likely that this data will quickly go out-of-date. Always contact vendor sources for the most current products and prices.

DEC's ATM Strategy:

On October 26, 1993, Digital Equipment Corporation announced the availability of three asynchronous transfer mode (ATM) products for enterprise-wide networks, as well as the details of a two-year product plan. These products build on Digital's ATM strategy and will be delivered over the next two years for use in both local- and wide-area networks (WANs).

The three ATM products consisted of a premises ATM switch to control the flow of data; a turbochannel adapter that allows devices--such as Alpha AXP workstations--to communicate in an ATM environment; and a GIGAswitch module that allows Digital's recently announced GIGAswitch network switching platform to convert FDDI traffic to ATM for transmission in the LAN or wide-area network.

The Potential Of ATM--Digital's ATM Network Strategy:

Digital views ATM as a critical component of the next generation of networks. Digital will provide a line of standards-based products for the future--products that will support existing multiple protocols (i.e., FDDI, token ring, and Ethernet), as well as the evolving technologies and services that combine information from voice, video, data, and imaging applications.

ATM is a network technology that permits an unlimited number of users to have dedicated high-speed connections with each other, and with network resources. On conventional shared-media networks--such as token ring, FDDI, and Ethernet--information is transmitted in variable-sized information packets over a fixed bandwidth. ATM converts information packets to fixed-size ATM cells and establishes a dynamic virtual pathway between the source and destination before transmission. When all of the cells are successfully received at the destination, the information packet is reassembled, and the individual links of the pathway become available for reconfiguration and reuse.

The use of fixed cells--which permit hardware switching among heterogeneous (WAN, LAN, remote, and mobile) networking nodes--is the key to ATM's advantages. Virtual workgroups can be set-up from a network-management console without altering the physical network. Even better, the ATM network can dynamically allocate bandwidth on demand for bandwidth-intensive applications, such as image-based communications, desktop-video-conferencing, interactive imaging applications, interactive CAD/CAM, remote- training, research collaboration, and financial applications.

Unlike other emerging high-speed technologies, ATM can accommodate mixed data types, because the fixed-sized ATM cells from various kinds of applications can be easily intermixed. Digital expects global corporations to adopt the use of ATM to support larger, more mobile user populations, multiple technologies, distributed computing, and multimedia information services.

Digital's ATM Product Strategy:

Actively involved in research on ATM-related products at its Systems Research Center in Palo Alto, CA--and also in the evolution of ATM standards--Digital has been developing ATM technology and products that will support the spectrum of applications for ATM. These products will be used by workgroup and departmental servers, the desktop (workstations and PCs), and enterprise-wide servers.

Digital's most recent announcements initiate a three-stage roll-out of ATM products, which will take place over two years:

- In Stage 1, Digital will offer products to make ATM LAN technology available to workgroups using high-performance computing applications
- In Stage 2, products will be offered to enable connection of legacy Ethernet and PCI-bus servers into an ATM LAN backbone
- In Stage 3, products will enable desktop systems to integrate with the ATM LAN

Digital claims to have developed solutions to many system-level problems not yet resolved by the ATM Forum or other ATM vendors. These features will be included and enhanced in future Digital ATM products. Digital is also pioneering a new technology that will enable information to smoothly cross the link between a private LAN and a public WAN, regardless of the differences in Quality of Service (QOS), flow control, or security between the two networks.

Digital'S ATM products fall into three categories:

- ATM switches--These switches will connect links from individual nodes--such as workstations and servers--to form an ATM LAN; the Premises ATM Switch provides both high-speed connections between ATM nodes and interoperability with standards-oriented vendors
- PC and Workstation ATM Adapter Cards--Adapter cards connect data- terminating equipment--such as workstations and servers--to the ATM LAN; Digital's TURBOchannel Adapter will enable DEC 3000AXP workstations running either OSF/1 AXP or OpenVMS AXP to connect directly to ATM networks; Digital's PCI adapter cards will be able to connect future RISC workstations from Digital, Hewlett-Packard, and Apple Computer, Inc., as well as Intel's Pentium-based PCs, to the ATM network
- ATM connections to existing platforms--These products bridge or route traffic from conventional LANs to ATM-based networks; Digital has three products already shipping in this category:
  - The DEChub 900 MultiSwitch system is designed to provide the bandwidth needed for the higher-speed emerging technologies, such as ATM

- The GIGAswitch two-port ATM line-card connects local FDDI networks to an ATM LAN or across the WAN internetwork
- The DECNIS Multiprotocol Router high-speed ATM interface will allow connections from legacy multi-protocol LANs and to emerging ATM services

Industry Activities:

Digital is an active member of the ATM Forum, an international consortium of network users, equipment vendors, and service providers. Founded in 1991, the ATM Forum is chartered to accelerate the use of ATM products and services through a set of interoperability specifications. The ATM Forum promotes industry cooperation by focusing on ATM interface issues. Digital is also working with the International Organization for Standardization (ISO) and the American National Standards Institute (ANSI) to speed ratification of ATM standards.

Earlier this year, Digital announced that it will be participating in a joint commercial and military testing program at the University of Kansas in Lawrence. Conducted under the auspices of the Multidimensional Applications and Gigabit Internetwork Consortium (MAGIC) of vendors and users, the three- year effort is testing ATM-switching equipment, advanced LANs, and image servers on a WAN network that includes equipment from Digital, Sprint, Northern Telecom, and other suppliers. MAGIC is funded in part by the Defense Advanced Research Projects Agency (DARPA).

Two of the high-speed network technologies being tested in this program are ATM and Synchronous Optical Network (Sonet), a physical interface already defined by the ITU-TSS (International Telecommunication Union-Telecommunications Standard Sector--formerly the CCITT). Sprint is providing the SONET WAN backbone that connects Minnesota Supercomputer's Minneapolis facility with the Earth Resources Observation Systems (EROS) Data Center in Sioux Falls, South Dakota, the University of Kansas, and the U.S. Army Future Battle Laboratory in Leavenworth, Kansas.

Digital, which has more than six years of experience with high- speed switching networks that use both fixed-length and variable- length information packets, is providing part of the ATM LAN for the MAGIC program. The predecessor for that ATM LAN was constructed from "off-the-shelf" components and has been in operation for four years at Digital's Systems Research Center in Palo Alto, CA.

Technology Leadership:

The valuable experience gained from operating a 30-switch internal network and managing the production and application of the GIGAswitch, a high-performance network packet switch, has helped Digital solve technical challenges presented by ATM networking.

With ATM, the successful transmission requires the successful delivery of all the ATM cells formed from the original information packet. If any one cell is lost, the entire packet must be transmitted again. When the ATM station repeatedly tries to resend the original information packet, the network becomes increasingly unstable as more network congestion occurs, and more cells are lost. As a result, ATM networks can only be stable when there is no cell loss. Digital has developed new technologies to both maximize bandwidth utilization and ensure that all cells transmitted into the ATM network are delivered to their destination.

FLOWmaster flow control is Digital's credit-based flow-control system that allows the network links to operate near capacity without cell loss and instability. SWITCHmaster queue control is said to improve the efficiency of a switch by maintaining a high rate of utilization. SWITCHmaster maximizes the volume of information transmitted through the switch fabric [of the network,] thereby significantly decreasing the cost of transmission.

9. Summary

ATM is a technology that is in the process of birthing. While ATM will eventually emerge from the development labs into the standardized, commodity marketplace, it may be several years until it attains this status. Until then, ATM will remain a specialty technology, useful in niche markets that can profit from its unique capabilities, while tolerating the costs and problems attendent with any emerging technology.

Even when ATM does reach commodity status, it must not be viewed as a panacea. ATM's small packet size will always place it at a disadvantage relative to protocols specifically enginered for the transmission of non-interactive data. The bottom line: *track the development of ATM, and purchase components employing existing technologies that will allow migration to ATM with a minimum of breakage*.

Appendix A: Fore Systems Plans

In this section, significant detail is provided regarding the internals of the Fore Systems product line. This detailed material is provided for those wishing a deeper insight into the lower level details of ATM. The reading and understanding of this material is not necessary to the understanding of the applicability and future of ATM.

Fore Systems' local-area network products are based on the international Asynchronous Transfer Mode (ATM) standard. The ATM standard was developed by the CCITT (Consultative Committee for International Telephone and Telegraph), with contributions from vendors and carriers worldwide, as the key component for the Broadband Integrated Services Digital Network (BISDN).

With its ForeRunner series of ATM LAN switches and computer adapters, Fore Systems provides users with ATM LAN workgroup, backbone, and wide-area networking products. The ForeRunner family of ATM computer adapters are available for Sun, DEC, HP, IBM, NeXT, Silicon Graphics, EISAbus, and VMEbus platforms.

The ForeRunner ASX-100 switch is the first in a line of ATM switches from Fore Systems and delivers desktop ATM connectivity to up to 16 computers. Larger LANs can be built using the ForeRunner ASX-100 SwitchCluster, which supports up to 64 ATM attached devices, or by interconnecting multiple ForeRunner ATM switches. Alternatively, the ForeRunner line of ATM LAN switches can be used for building backbone networks interconnecting routers, hubs, and bridges. Furthermore, its ATM wide-area network (WAN) interfaces--the ASX-100 family--supports direct connection from the ATM LAN workgroup, or campus backbone, to ATM wide-area DS-3, E3, and SONET/SDH networks. The 200-series adapters include the SBA-200 SBus adapter for Sun SPARCstations, the HPA-200 EISA adapter for HP workstations and high-end PCs, the VMA-200 VMEbus adapter for SGI workstations and other VMEbus computers, and the MCA-200 MicroChannel adapter for IBM RS 6000 workstations and high-end PCs.

200-Series Hardware Architecture

The 200-Series hardware architecture is divided into three distinct sections:

- The Network Interface, which includes a physical-medium- dependent section, FIFO buffering to isolate the network interface from the remainder of the adapter, CRC hardware for AAL processing, and control and status for the network interface

- The Intel i960 Control Processor, which implements the segmentation and reassembly functions and manages the transfer of data between the adapter and the host computer; the operation of the i960 software is described in the software architecture section below

- The Bus Interface, which provides a high-speed interface between the 200-series adapter and the I/O system of the host computer; the bus interface is re-implemented for each different I/O system but provides FIFO buffering to isolate the bus interface from the remainder of the adapter--it includes special- purpose DMA control hardware

The building blocks of the 200-series architecture are as follow:

Processor

An Intel 80960CA ("i960") processor running at 25 MHz is used to support segmentation and reassembly, packet buffer management, status monitoring, and error recovery. The i960 contains an instruction cache and data RAM on-chip.

To further enhance performance, the 200-series architecture also supports "fly-by" burst transfers: this enables the i960 to issue a burst read inquest that causes a burst of data to be transferred directly between the bus interface and the network interface. This fly-by ability improves the performance of the 200-series architecture, since it alleviates the i960 from having to read and write every word of data that is transferred between the network interface and the bus interface.

Memory

The 200-series architecture includes 256 kilobytes (KBytes) of static RAM memory. This memory is not used for buffering packets during the segmentation and reassembly process; instead it is used by the i960 for program memory, and for data structures that it needs for control of segmentation and reassembly.

In addition, this memory is accessible from the host computer and is used to communicate control information (such as lists of packets to be transmitted or received) between the host computer and the i960.

Network Interface

The network interface is a modular interface that is decoupled from the control processor and bus interface through the use of a pair of FIFO buffers. There is a 16-KByte buffer for incoming cells and a 1-KByte buffer for outgoing cells.

The decoupling between the network interface and the control processor allows the 200-series architecture to support multiple PMD (Physical-Medium-Dependent) interfaces, including:

- 100-Mbps Multimode Fiber Interface (as per the ATM Forum UNI Specification)--this interface is a full-duplex 100-Mbps ATM interface (a 140-Mbps version of this interface is also available); it is implemented with the AMD TAXI chipset and uses 4B/5B encoding to send data over multimode fiber; the line rate is therefore 125 MBaud; note that the multimode fiber specified for this interface is the same as specified for FDDI
- SONET OC-3c Interface (as per the ATM Forum UNI Specification)-- this interface is a full-duplex, 155.52-Mbps ATM interface; it is implemented using the PMC-Sierra SUNI chipset

Additional PMD interfaces planned by Fore Systems include a UTP (Unshielded-Twisted-Pair) Category 5 interface at 155 Mbps, and a UTP Category 3 interface at 25-50 Mbps; both of these copper interfaces will conform to ATM Forum specifications as they are completed.

CRC Support

The network interface also includes hardware support for computing and checking header and payload CRCs (Cyclic Redundancy Checksums). Support for computing AAL 3/4 and AAL 5 in parallel is provided. For computing AAL 5 payload CRCs--which must be computed over the complete PDU and not over individual cells as in AAL 3/4-- the CRC block maintains a partial CRC sum that is used as the basis for computing the CRC of the next cell. The i960 is responsible for ensuring that the correct partial CRC is in the CRC block for incoming packets.

Bus Interface

The bus interface of the 200-series provides both a master and a slave interface. The master interface allows the 200-series to do high-performance DMA (Direct Memory Access) transfers between host computer memory and the 200-series adapter. The slave interface allows the host to read and write control and status information on the 200-series adapter.

The bus interface is implemented using a pair of FIFOs--the IN FIFO and the OUT FIFO--in addition to the master and slave control logic. The control logic is specific to a particular host computer I/O bus and thus must be re-implemented for each bus to which the 200-series is ported. However, the IN and OUT FIFOs provide a bus-independent interface to the i960.

In addition to containing the actual data that is transferred over the host computer I/O bus, the IN and OUT FIFOs are also used by the i960 to control the DMA transfers, as follows:

- For a block of data to be written from the 200-series adapter to host memory, the i960 writes the destination address for the block in the OUT FIFO, followed by the block itself (note that the i960 can transfer the block to the OUT FIFO using the fly-by DMA feature discussed above); when the i960 has written the last word of the block, the DMA control logic executes the DMA transfer
- For a block of data to be read from host memory, the i960 writes the address of the block in the OUT FIFO; immediately thereafter, the DMA control logic executes the DMA transfer; the i960 can then start reading the words of the block from the IN FIFO

DMA Performance Tuning

The bus master logic for each implementation is tuned to the particular I/O architecture of the target host. Thus, on the Sun SBus, the bus master logic uses the burst transfers that are supported by the SBus (any of 8-, 16-, 32-, and 64-byte burst transfers, depending on the particular Sun workstation). On the VMEbus, the bus master logic is tuned to use block-mode transfers to transfer complete ATM cell payloads (48 bytes) over the bus.

Interrupts

In addition to its master and slave capabilities, the bus interface also has support for the i960 and the host computer to interrupt each other. This allows the i960 to inform the host computer that a completely reassembled packet is available.

Software Functions

As shown above, the 200-series Software Architecture splits the ATM protocol processing between the host computer and the i960 control processor on the 200-series adapter.

The software that runs on the i960 is responsible for ATM cell processing and ATM Adaptation Layer processing (specifically AAL 3/4 and AAL 5). The i960 uses the bus-master DMA and AAL- processing hardware of the 200-series adapter, as described above.

The host computer software--or device driver--is responsible for ATM sign-offing and addressing (including Fore Systems' SPANS protocol), and for providing the interface to the upper-level protocol modules in the operating system. These upper-level modules include existing network protocols--such as IP/TCP--as well as Fore Systems' ATM API (Application Programming Interface). In addition, the device driver includes an SNMP agent that keeps track of ATM statistics and reports them to network- management platforms via SNMP.

In the UNIX operating system environment, the device driver supports the TCP/IP protocol and performs the following functions:

- Address resolution of IP addresses using the ARP (Address Resolution Protocol)
- Establishment of virtual connections for IP destinations, and for applications using the ATM API
- Encapsulation and multiplexing/demultiplexing of PDUs (Protocol Data Units, i.e., packets) over virtual connections for PDUs at either the network level (e.g., IP packets) or the datalink level (e.g., 802.2 packets); for IP packets, encapsulation and use of virtual connections is as specified by the IETF IP-over-ATM working group

Packet-Level Interface

The interface between the device driver and the i960 is a PDU interface: in other words, the driver handles only complete upper- level PDUs and does not need to be concerned with ATM cells--the i960 handles all of the cell-level processing. In addition, the i960 maintains full statistics about cells received and transmitted, as well as the number of cells with incorrect CRCs (header or payload), or with connection identifiers for connections that have not been established.

To transmit a PDU, the device driver supplies the i960 with a connection identifier (VPI/VCI) and a set of buffer descriptors. Each descriptor consists of an address in host memory and a length indication. The i960 uses the 200-series DMA to transfer PDU data from host memory. As data arrives on the adapter, the i960 creates the ATM cells by prepending the cell header. It then transmits the cells over the network interface by writing them to the transmit queue. In normal operation, it can use the fly-by DMA features of the 200-series adapter to copy the data from the bus interface to the network interface.

The interface, and the i960, operate as follows for incoming cells:

- To receive cells on a connection, the device driver first specifies to the i960 that a particular connection (VPI/VCI) has been opened; in addition, it must provide the i960 with a set of buffer descriptors that form a free pool of buffers for use by the i960; the device driver is responsible for ensuring that the i960 always has a pool of free buffers

- For each open incoming virtual connection, the i960 maintains a reassembly context which contains information about the state of the connection, as well as a list of buffers (in host memory) that are being used to reassemble an incoming PDU; when the i960 receives a cell for an open connection, it checks to see whether there is a buffer allocated to the reassembly context--if there is not, it allocates a new buffer for the context and then checks to see whether there is sufficient free space in the buffer; if there is not, it allocates a new buffer for the context and then uses the DMA hardware on the adapter to offer the cell's payload to the free buffer space in host memory; and, finally, it updates the state of the reassembly context to reflect the arrival of the new cell
- When a PDU has been completely reassembled and all of its data has been copied (using DMA) to host memory, the i960 interrupts the host to inform it of the arrival of the packet--note that the i960 transfers cell payloads to host memory as cells arrive, instead of waiting until a PDU has been completely received

200-Series Performance

According to Fore Systems, initial performance results for the 200-series architecture are as follows: SBA-200 SBus ATM adapter (100-Mbps PMD)

TCP/IP performance measured using the TCP test program:

TCP window size/SPARCstation 2/SPARCstation 10

16 KByte/24.8 Mbps/29.8 Mbps

32 KByte/35.6 Mbps/42.3 Mbps

51 KBytes/38.6 Mbps/56.6 Mbps

UDP/IP performance measured using the TCP test program:

SPARCstation 2/SPARCstation 10

63.2 Mbps/77.5 Mbps

HPA-200 HP ATM Adapter (100-Mbps PMD)

TCP/IP performance measured on HP Model 735 workstations using the TCP test program: 55 Mbps

ForeThought Management Software For ForeRunner ATM Switches

The ForeThought management software--the brains of the ForeRunner ASX-100 switch--is comprised of two main functions: the ForeThought Integral Connection Management Software and the ForeThought Integral Switch Management Software.

All critical functions for ATM LAN operations are performed by these software elements. The ForeThought software is resident on an integral RISC-based Control Processor engine within each ASX-100 switch on the network.

Essential network-wide operations--such as topology discovery and routing--are performed in a distributed, peer-to-peer fashion among all intelligent ASX-100 switches.

ForeThought Integral Connection Management Software

The ForeThought Integral Connection-Management Software performs all connection-management functions, including topology discovery, connection establishment, VCI and VPI routing, bandwidth allocation, multicast connectivity, connection teardown, and rerouting of virtual channels and virtual paths. Connection management is performed within the ForeRunner ATM network using the SPANS UNI and NNI protocols developed by Fore Systems.

The ForeThought feature set includes Automatic Network Configuration, Optimized Routing, and QuickConnect connection establishment and rerouting, as called-out below.

Automatic Network Configuration is designed to eliminate operator intervention for network installation, network moves, adds and changes (M.A.C.), and maintenance of routing tables. As new switches, ports, or trunks are added to the network, the network automatically "learns" of their presence and adds them to the topology map for future connection routing decisions.

ForeRunner switches automatically detect failures (trunk failures, switch hardware failures, workstation or switch reboots, and power cycling), dynamically calculate alternative paths, and reroute connections.

Optimized Routing means that every routing decision is made optimally, taking into consideration existing network conditions that include interswitch trunk utilization, port-bandwidth utilization, priority of connection, and number of network hops. Links are automatically load-balanced to reduce network congestion and allow for burst-mode communications.

ForeThought software supports multiple classes of ATM services. A dual-priority scheme allows delay-sensitive applications--such as multi-media or video--to pass through the network with a predetermined quality of service (QOS). Connections can be established with negotiated peak bandwidth, average bandwidth, and burst-length parameters.

QuickConnect connection establishment and rerouting utilizes distributed intelligence to ensure that each switch within the network is continually updated so that routing decisions can be quickly made by each local switch.

SPANS Signaling Protocol

At the core of the Integral Connection Management Software is Fore Systems' SPANS protocol (Simple Protocol for ATM Network Signaling). All switch-to-end station and switch-to-switch communications (ATM UNI and NNI signaling, respectively) are provided via the SPANS suite of signaling commands. SPANS is architected for simplicity; each operation is performed in a single message.

SPANS is architected so that on-demand connections--or Switched Virtual Circuits (SVCs)--can be created between any pair of end stations (computers, routers, bridges, or hubs) that need to communicate across the ATM network. The use of SVCs in ATM LAN networks is essential;

using Permanent Virtual Circuits (PVCs) in a LAN environment would require continual configuration and manual routing of connections. A PVC-only ATM LAN would be analogous to antiquated voice systems controlled by switchboard operators.

Furthermore, SPANS supports key ATM networking features, including connection-oriented transport with guaranteed bandwidth and quality of service (QOS). In addition, SPANS offers the same capabilities as legacy LANs, such as multicast, broadcast, and connectionless service. Multicast functions are achieved without using the congestion-prone cell-copying techniques deployed by other ATM network architectures.

The ATM standards in the signaling arena are still incomplete, so Fore Systems developed the SPANS protocol to meet requirements for ATM local-area networking today. SPANS provides similar functions to the draft ATM Forum signaling specifications. Once standard signaling protocols are established by the ATM Forum, the ForeRunner switches will be updated to support these protocols.

Congestion Management

The ForeRunner ATM Switch provides congestion management in three main ways:

- Congestion avoidance
- Congestion reporting
- Flow control

The ASX-100 implements a number of features for congestion avoidance:

- Connection requests are only accepted up to the maximum allowable port speed
- Should a given connection exceed its negotiated rate, any traffic beyond that rate is transferred through the network at a reduced priority
- Interswitch trunks are load-leveled, reducing the probability of trunk congestion
- Buffering at each output port (approximately 32 KBytes) is provided in the event of temporary congestion
- Elimination of bandwidth-producing, cell-copying techniques for multicast operations

Congestion reporting--in the form of bandwidth utilization--from the ASX-100 via Switch Management Software is provided on a port, VCI, and VPI basis for policing by the network-operations manager.

Flow control is currently left to the end stations, typically running TCP/IP protocols which will throttle the applications based upon network-congestion conditions. In addition, the ForeRunner ASX-100 provides output port buffering.

ForeThought Integral Switch Management Software

The ForeThought Switch Management Software contains management functions--explained in detail in the following subsections--that include:

- Configuration management

- Network-performance monitoring and reporting
- Fault management
- Security management
- Inventory management

All management information collected by the Switch Management Software is stored within the ASX-100's internal 120-MByte hard drive and can be subsequently accessed by any password-protected management station on the network via inband ATM ports or out-of- band Ethernet, FDDI, or serial ports.

The ATM Management Software is architected using the SNMP open- management protocol, which has been specified by the ATM Forum for management of ATM networks. Each ForeRunner switch contains an integral SNMP agent. Currently deployed SNMP network-management systems--such as HP OpenView or SunNet Manager--can access and display ATM network information.

Configuration Management:

In most cases, configuration of the ForeRunner ATM network is automatic. No maintenance of routing tables is required. Operator configuration is required only for initial naming of network resources, permanent virtual circuit establishment, and Ethernet and FDDI interface configuration.

All software for the ForeRunner switch can be downloaded via the ATM, Ethernet, or FDDI interfaces. Furthermore, this software can be downloaded and updated while the switch is in operation for "on- the-fly" updates. Furthermore, all network interfaces maintain their software images during a power failure.

Network-Performance Monitoring:

A set of performance parameters are continually collected and stored for each element of the ATM network. Utilization, status and error statistics for each virtual channel, virtual path, switch port, and interswitch link are continually compiled and available for performance reporting from desktop workstations or SNMP management stations.

Fault Management:

In the event of network failures--including interswitch link outage, switch or workstation power cycling or rebooting, or switch hardware component failure--the SCS automatically recognizes the event and reconnects applications over available network paths. Event history is logged on the internal hard drive and reported to the network-management systems.

Security Management:

End-stations control over incoming connections requests is provided by the SPANS protocol. Since the caller address is provided during connection establishment, the ForeRunner network is in a position to accept or deny connection requests, or even to create virtual subnets within the single ATM network fabric. Furthermore, the ForeRunner network can dedicate bandwidth for critical applications.

Inventory Management:

The Switch Management Software makes available information that assists in remotely maintaining inventory of the ATM network. Parameters available include name, serial number, module type, address, and hardware and software version.

Specifications:

Network Management SNMP MIB II supports the Interim Local Management Interface Specification (ILMI). SNMP MIB information includes:

- Inventory management--switch hardware serial number, switch hardware version, software version, board and network module count, number of ports on switch, max VPI/VCIs, output buffer number, type, status, logical size; adapter serial number, HW version, HW speed, firmware version, software version, buffer size, queue length, operational status, address, carrier detect
- Configuration management--network switch-topology map, VPs originating on switch, VCs passing through switch, switch address, uptime counter; port number, port status, ATM address, IP address of port-connected device, max incoming/outgoing VPs and VCs; for virtual paths--number VCs currently allocated, max VC allocated
- Bandwidth utilization--ATM cells transmitted/received per end station, current cell count, allocated bandwidth, max bandwidth (cells/sec), current port bandwidth utilization, VP allocated bandwidth, VP current bandwidth, VP max bandwidth
- Errors--errored cell physical-layer framing, bad header CRCs, VCIs out-of-range, inactive VCIs, AAL checksum errored cells, AAL protocol errored cells, AAL discarded cells, VPI/VCI look-up errors
- Ordering Information--SCS: Switch Control Software for ASX-100; SCS/SC: Switch Control Software for ASX-100 SwitchCluster; SCS/SRC: Switch Control Software source code

E3 Support For The ForeRunner ATM Switch Family

Fore Systems plans to enhance its ForeRunner ASX-100 ATM (Asynchronous Transfer Mode) Switch offering by providing support for a standards-based E3 (34-Mbps) switch interface for use in European and Asian networks who have deployed the CCITT Plesiochronous Digital Hierarchy (PDH) standards.

The E3 Network Module is being developed for connecting the ForeRunner ASX-100 and ForeRunner ASX-100 SwitchCluster ATM switches to equipment and services supporting ATM at the E3 data rate. ForeRunner ATM switches will be able to connect to other ForeRunner switches, ATM multiplexers, ATM DSUs, private PDH/E3 links, or E3 ATM services.

Need for E3 Best Described in the Following Quote from the CCITT Recommendations:

> "Existing transmission networks are based upon the Plesiochronous Digital Hierarchy. ... The Synchronous Digital Hierarchy (SDH) will form the basis of transport of the ATM cells. During the transition period, there needs to be transport of ATM cells using existing PDH transmission networks."

The standards-based E3 interface will be engineered using an E3 chip developed in a partnership with PMC-Sierra of Burnaby, BC, Canada. Fore Systems and PMC-Sierra will be jointly funding this effort. This development will produce the world's first CCITT- standard E3 ATM chip and will be available for use by all ATM vendors wishing to add E3 support to their product lines. Dubbed the SUNI-PDH chip, the feature set will provide a superset of the currently available PLPP chip--also from PMC-Sierra--supporting DS-3, T-1, and E1 ATM speeds.

The E3 chip will be compliant with ATM Forum User Network Interface (UNI) E3 specifications, which identify CCITT Recommendation G.705 and Recommendation G.804 specifications for mapping of ATM cells into PDH.

The E3 interface will complement the family of ForeRunner wide-area network (WAN) interfaces which include the DS-3, SONET OC-3c, and SDH (Synchronous Digital Hierarchy) STM-1 Network Modules.

E3 Network Module features will include:

- E3, 34.368-Mbps data rate
- 2-port or 4-port interface module
- Coax connector
- Field-installable
- Compatible with ATM Forum UNI specification
- Compatible will CCITT G.705 and G.804 specifications
- SNMP network management

Unshielded-Twisted-Pair and SONET/SDH Support For ATM Adapters

Fore Systems plans to enhance its line ATM Adapters to include support for unshielded-twisted-pair and SONET/SDH physical layer interfaces. Currently, the line of ForeRunner ATM adapters (for Sun, DEC, Silicon Graphics, Hewlett Packard, IBM, EISAbus, and VMEbus computers) supports the 100-Mbps (4B/5B) ATM Forum User Network Interface (UNI) specification and the 140-Mbps speed using fiber-optic interfaces.

Twisted-Pair

Recent submissions by industry vendors and work within the ATM Forum have resulted in twisted-pair recommendations (UTP and STP) to allow customers to use their installed copper wiring plant. Fore plans to support both the lower-speed voice-grade UTP-3 (unshielded-twisted-pair) specification, and the high-speed, 155- Mbps, data-grade UTP-5/STP specification, once they are ratified by the ATM Forum

SONET/SDH

The SONET/SDH LNI specification--which calls for a 155-Mbps data rate using OC-3c/STM-1 framing--is currently part of the ATM UNI. Recently available chips will allow Fore Systems to incorporate this SONET/SDH functionality into its existing ATM adapter product line. Fore has already shipped SONET/SDH functionality with its family of ForeRunner ATM switches and will use similar technology to provide the adapter features. Multimode fiber will provide the physical media interfaces for the adapters.

ForeRunner ATM Developers Kit

The ForeRunner ATM Adapter Kit--designed to connect two workstations--includes two 100-Mbps ATM adapters and one 50-foot multimode fiber-optic cable. The ForeRunner ATM LAN Kit-- designed to connect four workstations in a LAN workgroup-- includes an ASX-100 2.5-Gbps ATM switch, one ASX-100 Network Module with four ATM Forum UNI-compliant, 100-Mbps ports, four 100-Mbps ATM adapters, and four 50-foot multimode fiber-optic cables. Adapters are available for Sun, Silicon Graphics, DEC, Hewlett Packard, NeXT, and VMEBus workstations.

ForeRunner workstation adapters support TCP/IP applications and come equipped with the ForeRunner API (Applications Programming Interface) library. This offers applications access to the features of ATM, such as guaranteed bandwidth reservation, per- connection selection of AAL 5 or 3/4, and multicasting with dynamic addition and deletion of recipients.

Both kits include example application source code, a software developer's documentation package, and 12 months of ForeMan technical support. Vendors will also receive a free listing in the ForeRunner ATM Applications Directory.

ForeRunner ATM Developers Kits are available now to qualified applications vendors and are limited to one kit per customer site. The price of the ATM Adapter Kit is $5,995; the price of the ATM LAN Kit is $39,995.

Fore, Sprint Team On Complete ATM Service Offering

In an effort to provide a complete end-to-end ATM service offering, Fore and Sprint have agreed to co-market the Fore Systems' line of ForeRunner ATM switches.

Sprint has rolled-out an ATM service, available in 300 Points of Presence (POPs) throughout the U.S. (Details are available from Sprint.)

The ForeRunner ATM switches will be available as Customer Premises Equipment (CPE) to Sprint's customers, including local- area connections at 45, 100, 140, and 155 Mbps.

Through the use of SVC (Switched Virtual Circuit) tunneling, customers will be able to use Fore's ForeThought connection- management software in both the local-area campus environment and the wide-area network.

3Com And Fore Partner

As also mentioned earlier, 3Com and Fore Systems have established a strategic-marketing partnership. In addition, the two companies will investigate opportunities for potential areas of technology cooperation.

In the initial phase of the multi-year agreement, the two companies will co-market Fore's ForeRunner family of ATM LAN workgroup and backbone products. 3Com and Fore also will conduct interoperability testing to ensure interoperable products based upon both companies' products.

The companies have completed initial compatibility tests running NETBuilder II router traffic through the ForeRunner ASX-100 ATM switch via an ATM DSU. Plans are also underway for formal LAN interoperability testing of both companies' full product lines.

Fore Systems Partners With Cabletron

Fore Systems, Inc., has also announced a long-term relationship with Cabletron Systems to deliver ATM (Asynchronous Transfer Mode) products for LAN backbone and LAN workgroup applications. The arrangement between the two companies includes a technology partnership and an OEM agreement.

The technology partnership includes two major components:

- Using the ForeRunner ATM adapter card products, Cabletron will offer an ATM backbone interface from Cabletron intelligent hubs to Fore Systems' ATM switches
- Fore Systems will join the SPECTRUM Partners' Program to allow the SPECTRUM network-management platform, via SNMP, to manage the ForeRunner ASX-100

The non-exclusive OEM agreement with Cabletron entitles the company to become a worldwide reseller of the ForeRunner ATM product family.

The currently shipping ForeRunner ASX-100 will be used as a high- speed backbone to connect users residing on MMAC intelligent hubs. In addition, users requiring dedicated ATM bandwidth to the desktop will connect directly to the ASX-100 switch using the ForeRunner ATM Computer Interfaces.

EISA And Microchannel Adapters For HP, Silicon Graphics, And IBM Computers

The ESA-200 EISAbus ATM adapter is available immediately, with initial software drivers for Hewlett-Packard 700 series and Silicon Graphics Indigo-2 series workstations. The ESA-200 adapter hardware has also been tested in EISA-based PCs and will be targeted for use in high-end PCs and servers.

The MCA-200 ATM adapter provides support for IBM Microchannel computers. Initially, AIX drivers will be available for IBM RS6000 computer platforms, with additional Microchannel platforms being supported in the future.

Both the ESA-200 and MCA-200 products utilize the Advanced Cell Processing Architecture developed by Fore Systems and used in the existing SBA-200 ATM adapter for Sun

SPARCstations. This architecture is based on an embedded Intel i960 RISC processor with special-purpose AAL cell-processing and DMA hardware. This architecture provides the highest available throughput, as well as the flexibility to support a wide range of computer platforms.

As with all ForeRunner ATM adapters, the ESA-200 and MCA-200 include software that supports a range of networking applications, including standard protocol suites that support existing software applications and applications that use the ForeRunner ATM API (Applications Programming Interface). Furthermore, the adapters support PVCs (Permanent Virtual Circuits) and the Fore Systems' SPANS SVC (Switched Virtual Circuit) signalling protocol, which offers customers on-demand call set-up, guaranteed bandwidth reservation, per-connection selection of AAL (types 3/4 or 5), and multicasting.

ESA-200 pricing for HP 700 series and Silicon Graphics' Indigo-2 series workstations starts at $2,495 per adapter in quantities of six. Pricing for PC and server platforms using the ESA-200 hardware will be announced at a later date. MCA-200 pricing, including AIX drivers for the IBM RS6000 product line, starts at $2,495 per adapter in quantities of six.

Adapters are currently available with ATM Forum 100-Mbps and 140- Mbps TAXI UNI fiber-optic physical interfaces. Unshielded- twisted-pair and SONET physical-layer interfaces will also be provided.

SONET/SDH ATM Interface

Fore Systems has announced the general availability of its OC-3c Network Module--SONET/SDH ATM interface--for use with its family of ForeRunner ATM Switches.

Both LAN and WAN connections are made possible with the SONET/SDH module. Users can use the modules to connect to other ForeRunner switches, ATM-capable SONET/SDH desktop adapters, ATM-ready hubs and routers, campus backbone muxes, private SONET/SDH links, or SONET/SDH services.

The OC-3c Network Module operates at the 155-Mbps speed and conforms to the ATM Forum UNI specifications and Bellcore/CCITT standards. Framing is software-selectable and compatible with both the SONET OC-3c and SDH STM-1 standards.

Up to four interfaces are available on a single network module. A variety of physical media interfaces are available:

- Multimode fiber for desktop use and campus backbones
- Short-reach singlemode fiber
- Long-reach singlemode fiber for wide-area connectivity

The long-reach singlemode fiber module is also available in a two- port module. Pricing for the OC-3c SONET/SDH module starts at $6,995.

The OC-3c Network Module complements the DS-3, 100-Mbps UNI, and 140-Mbps TAXI interfaces, which are currently available for the ForeRunner ASX-100 and ASX-100 SwitchCluster ATM switch products. This range of interfaces gives Fore Systems the most complete array of ATM interface types on the market.

SBA-100 Sun Adapters Pricing

In March of 1993, the SBA-100 unit price was cut from $3,995 to $1,995 in quantities of six. This latest price-reduction offers the adapters for $1,295 per adapter in quantities of twelve.

Price reductions represent continuing cost reductions experienced by Fore as a result of large volume shipments, as well as engineering cost reductions. To date, Fore has sold in excess of 900 ATM adapters to its customer base.

DS-3 Network Module--ForeRunner ATM Switch Interface

The DS-3 Network Module provides ATM connectivity via the ForeRunner ASX-100 ATM LAN switch. Operating at the standard 45- Mbps rate and conforming to the ATM Forum UNI (User-Network Interface) specifications, the DS-3 Network Module can connect directly to a DS-3 (T3) service, a DS-3 Data Service Unit (DSU), a DS-3 multiplexer, or Digital Crossconnect System (DCS).

In conjunction with the ForeRunner ASX-100 ATM switch, the DS-3 Network Module provides interswitch connections for DS-3 campus backbone, metropolitan, and wide-area networking. In addition, the DS-3 Network Module enables existing router-based LAN networks to connect to the ASX-100 via an ATM-capable DSU.

Up to four DS-3 interfaces are available on a single, compact Network Module which can be plugged into any of the network module slots on the ASX-100 switch. Up to 16 DS-3 interfaces can be configured on a single ASX-100 switch, or up to 64 on the ASX-100 SwitchCluster. The DS-3 physical layer carries standard ATM 53- byte cells, thereby providing true ATM connectivity. The DS-3 interface conforms to ATM Forum UNI and ANSI/Bellcore specifications.

ForeRunner SBA-100--ATM SBus Adapter For Sun SPARCstations

The SBA-100 ATM SBus adapter provides Sun SPARCstations with dedicated fiber-optic connections to ATM switches that include the ForeRunner line of ATM switches. Operating at data rates of up to 140 Mbps, the SBA-100 is suitable for distributed applications requiring high-bandwidth and low-latency networking.

The SBA-100 is a single-slot SBus card supporting standard ATM cell-processing, including segmentation and reassembly. Included with the adapter is a SunOS device driver supporting all TCP/IP and OSI protocols, as well as ATM Adaptation Layers (AAL) 3/4 and 5. With the SBA-100, SPARCstation users can add ATM networking capabilities to their applications, leaving the low-level ATM cell processing, segmentation and reassembly, and signaling to the SBA-100 hardware and device driver.

The SBA-100 uses the SPANS SVC protocol to give applications end- to-end ATM connectivity on-demand. SPANS supports workstation- initiated unicast, multicast, and broadcast. Standard network- management functions are available through an SNMP ATM MIB.

Applications developers are provided with a number of tools designed to simplify the migration to ATM networking. Documentation of the hardware interface is included with the SBA- 100. Source code for the SBA-100 driver is also available.

ForeRunner GIA-100--ATM GIO Bus Adapter For Silicon Graphics Workstations

The GIA-100 is a single-slot GIO Bus card supporting standard ATM cell-processing, including segmentation and reassembly. Included with the adapter is an IRIX device driver supporting all TCP/IP and OSI protocols, as well as ATM Adaptation Layers (AAL) 3/4 and 5. With the GIA-100, Indigo users can add ATM networking capabilities to their applications, leaving the low-level ATM cell- processing, segmentation and reassembly, and signaling to the GIA-100 hardware and device driver.

The GIA-100 uses the SPANS SVC protocol to give applications end- to-end ATM connectivity on-demand. SPANS supports workstation- initiated unicast, multicast, and broadcast. Standard network- management functions are available through an SNMP ATM MIB.

Applications developers are provided with a number of tools designed to simplify the migration to high-performance ATM networking. Documentation of the hardware interface is included with the GIA-100. Source code for the GIA-100 driver is also available.

ForeRunner TCA-100--ATM TURBOchannel Adapter For DEC Workstations

The TCA-100 is a single-slot TURBOchannel card supporting standard ATM cell-processing, including segmentation and reassembly. Included with the adapter is an ULTRIX device driver supporting all TCP/IP and OSI protocols, as well as ATM Adaptation Layers (AAL) 3/4 and 5. With the TCA-100, DECstation users can add ATM networking capabilities to their applications, leaving the low-level ATM cell processing, segmentation and reassembly, and signaling to the TCA-100 hardware and device driver.

The TCA-100 uses the SPANS SVC protocol to give applications end- to-end ATM connectivity on-demand. SPANS supports workstation- initiated unicast, multicast, and broadcast. Standard network- management functions are available through an SNMP ATM MIB.

Applications developers are provided with a number of tools designed to simplify the migration to high-performance ATM networking. Documentation of the hardware interface is included with the TCA-100.

Appendix B: References

Direct / edited references used in this report:

Journal: UNIX Review Oct 1992 v10 n10 p28(7)
Title: Speeding to the ATM: the next-generation LAN architecture
Asynchronous Transfer Mode, uses cell-relay technology to move data at high-speeds across networks.
Author: Lamb, Chris
Journal: LAN Magazine August 1994 v9 n8 p101(5)
Title: Open Issues
(ATM standards) (includes related article on speed and cabling standards)
Author: Feltman, Charles
Journal: Datamation Jan 21, 1994 v40 n2 p20(5)
Title: Rocket science or lost in space?
(Asynchronous Transfer Mode in wide area networks)(includes related articles on user complaints that ATM is too costly, problems with using current TCP/IP under ATM)
Author: Strauss, Paul
Journal: PC Week Jan 31, 1994 v11 n4 p21(2)
Title: Industry, carriers dialing for ATM: phone companies' adoption of high-speed standards spurs pilot application development.
(Sprint Communications Corp, U.S. West Communications; Asynchronous Transfer Mode) (PC WEEK Special Report: ATM/ISDN)
Author: Crowley, Aileen
Journal: Digital News & Review Dec 20, 1993 v10 n24 p9(1)
Title: ATM Forum considers congestion management. (Asynchronous Transfer Mode)
Author: Lawton, Stephen
Journal: The LocalNetter Dec 1993 v13 n12 p55 (1)
Title: Special report: Digital Equipment Corporation's ATM strategy and products. (DEC's asynchronous transfer mode products for enterprise-wide networks)
Journal: The LocalNetter Dec 1993 v13 n12 p56(1)
Title: Special report: Fore Systems, Inc.'s ATM technology and products.
Journal: Datamation August 15, 1993 v39 n16 p20(1)
Title: Virtual LANs pave the way to ATM.