

NAT'L INST. OF STAND. & TECH. R.I.C.



A11104 498556

NIST
PUBLICATIONS

NISTIR 5561

Asynchronous Transfer Mode Procurement and Usage Guide

Tim Boland

U.S. DEPARTMENT OF COMMERCE
Technology Administration
National Institute of Standards
and Technology
Gaithersburg, MD 20899

QC
100
.U56
NO.5561
1994

NIST

Asynchronous Transfer Mode Procurement and Usage Guide

Tim Boland

U.S. DEPARTMENT OF COMMERCE
Technology Administration
National Institute of Standards
and Technology
Gaithersburg, MD 20899

December 1994



U.S. DEPARTMENT OF COMMERCE
Ronald H. Brown, Secretary

TECHNOLOGY ADMINISTRATION
Mary L. Good, Under Secretary for Technology

NATIONAL INSTITUTE OF STANDARDS
AND TECHNOLOGY
Arati Prabhakar, Director

TABLE OF CONTENTS

1.0	Introduction to Paper.....	1
2.0	ATM.....	1
2.1	Introduction to ATM.....	1
2.2	ATM Switch.....	1
2.3	Physical Layer and ATM Layer.....	8
2.4	ATM Adaptation Layer.....	8
2.5	ATM Network.....	11
2.6	ATM Interfaces.....	11
2.7	Features of ATM.....	14
2.8	ATM Status Today.....	17
2.9	ATM Implementation.....	21
2.10	ATM Applicability.....	24
3.0	User Decisions Regarding ATM Products.....	24
3.1	ATM Access.....	24
3.2	ATM Switches.....	25
3.3	Security.....	26
3.4	Management.....	27
3.5	Additional Procurement Considerations.....	29
4.0	ATM Migration/Coexistence.....	30
5.0	Summary and Future.....	32
6.0	User Acquisition Considerations.....	35
	Appendix A Broadband ISDN Environment and ATM.....	37

ATM PROCUREMENT AND USAGE GUIDE

1.0 Introduction to Paper

ATM (Asynchronous Transfer Mode) is a revolutionary new data communications technology that involves a high-speed switching protocol to move cells (each containing 53 bytes) rapidly between two (or more) distinct points in a network. ATM uses cell switching for the high-speed transmission of data over a variety of subnetwork technologies. The media used and transmission rates are not a function of ATM. Use of the ATM technology should enable users to communicate more information in a fixed period of time. This paper will provide the necessary guidance to enable users to properly assimilate ATM products into the workplace. An important medium for ATM use is the Synchronous Optical Network Transmission (SONET) medium; SONET will be discussed in this paper.

Section 2 provides information on ATM. Section 3 explains the decisions that users have to make in procuring ATM equipment and the factors that should influence their decisions. Section 4 covers coexistence/transition issues involving ATM. Section 5 gives a summary and future of ATM. Section 6 gives a recapitulation of important user issues in acquiring ATM. Appendix A gives the relationship between ATM and the broadband ISDN environment, since broadband ISDN is a potentially widespread use for ATM.

2.0 ATM

2.1 Introduction to ATM

ATM works by performing two simple functions: cell-switching (in hardware), and path (next hop) determination (virtual connection management done mainly in software). ATM makes high speeds possible, and delay small, because it may be necessary to have only these two functions at a network node, without the overhead of additional functionality. A connection is created, called a virtual circuit (VC). A VC is similar in concept to an ordinary phone connection, which is in effect until connection termination. Each virtual circuit is characterized by a virtual circuit identifier (VCI), which distinguishes one virtual circuit from another. This is similar to distinguishing one phone conversation from another conversation on the same trunk line.

2.2 ATM Switch

ATM connections are most often implemented via a device called an ATM switch (shown in fig. 1a). An ATM switch can handle multiple bit streams. This is similar to a large phone switch handling multiple conversations from different points. A bit stream is assigned to a port (access point on the ATM switch), either

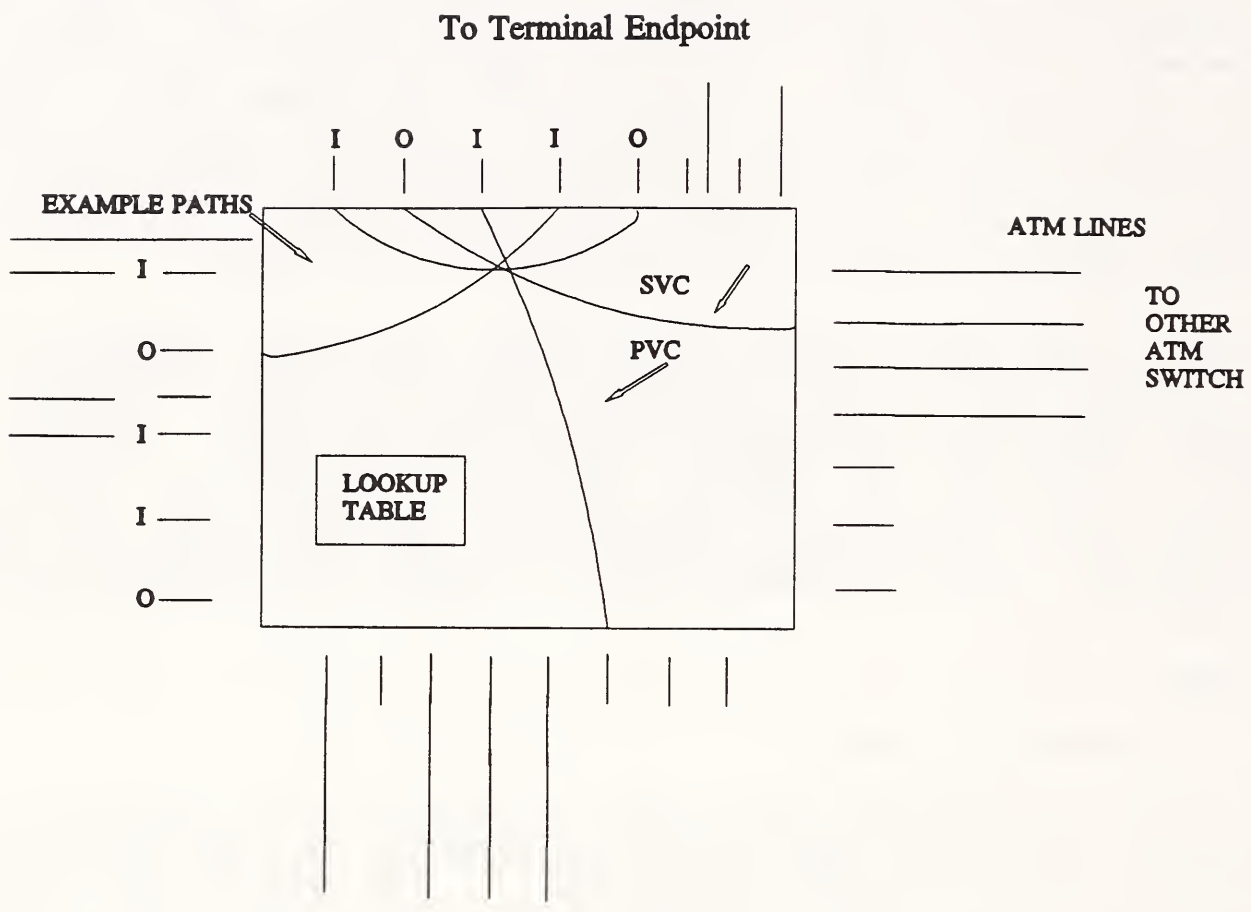


Figure 1a - ATM Switch

I = Input Component
 O = Output Component

statically or dynamically. Each port has a port identifier, which is known to the ATM switch. Different ports may support different data rates, depending on the physical medium to which the port is attached. There is always a pairing of input port and output port. Each port can handle multiple conversations (VCs); different conversations are distinguished by their VCI values. Semantically, a conversation may be a particular information type (e.g., voice or video), or have another meaning of interest to an application.

For cell switching, the bit streams coming into the ATM switch are examined, and cells are extracted as fast as they arrive. From the point of view of the switch, cells are just groups of bits, only distinguishable by their VCI values. Cells contain 5 bytes of header information (including a VCI for that cell), and 48 bytes of user information. The ATM switch performs no conversion on the user information; it just passes it on. The ATM switch fills the outgoing bit streams with newly-arrived cells to one or more output ports. Physically, the ATM switch takes input cells and redirects them to output ports using a "cross-connect" mechanism (see fig. 1a). The information in one cell may belong to a different conversation from the data in another cell on the same output bit stream, so the header of each output cell needs to contain an identifier for the conversation. The ATM switch may verify the correctness of this identifier, and may translate the identifier to make it presentable to the next system, because different systems may support different VCI syntaxes, and there is no specification supporting a common syntax for the VCI. If there is more information ready at an instant than cells available, the ATM switch may need to temporarily store the excess data in special memory areas called buffers before sending it out. There are different types of buffers; the number and size of buffers depends on the data rates supported on that port. In summary, the ATM switch performs statistical multiplexing, in which each cell is filled up with the first available information. This is opposed to traditional time division multiplexing, where time slots are allocated in a "round robin" fashion to alternate conversations, regardless of whether there is any information to send at that time.

An important function of an ATM switch is virtual connection management (virtual path or next hop determination). A virtual path is a collection of VCs possibly belonging to the same application. To convey the fact that different virtual circuits may belong to the same virtual path, a virtual path identifier (VPI) is also included in the cell header. The combination of VCI and VPI uniquely distinguish a conversation among all applications and conversations supported by a single ATM switch (which enforces this uniqueness by allocation of values); the VCI alone is not sufficient, because the same VCI value may be used incidentally for different virtual paths. For an illustration of this concept, see

figure 1b. Path (next hop) determination is accomplished by means of a lookup table in the ATM switch; entries in the table are consulted to determine where to send the cells next. These entries are accessed by a label, which is a concatenation of the VCI and VPI in the cell header. The entries may contain a new label (new VCI and/or VPI) which is inserted into the cell header leaving the switch. The output port identifier(s) is also determined from this VPI/VCI value. Some switches just use the VPI value for next hop determination (see fig. 1c), and ignore the VCI; thus all bit streams with the same VPI value are switched together. The cell header is recomputed before leaving the switch. The lookup tables must be set in advance (with space reserved); they may be pre-defined, dynamically assigned, or both.

A virtual circuit link is a means of unidirectional transport of ATM cells between a point where the VCI is assigned and the point where that value is translated or removed. A virtual circuit connection is a concatenation of virtual circuit links that extends between two terminal end points (see fig. 2). In figure 2, virtual circuits are called virtual channels, and ATM switches are called ATM nodes. Virtual circuit connections are established with reserved bandwidth before data transfer begins. The control traffic (signalling) for setting up and releasing virtual circuit connections flows on separate virtual circuit connections. This is in contrast with the in-band signalling mode of X.25, where control data is intermixed with user data on the same line. An ATM virtual circuit connection between two users is similar to a virtual wire between the two users. There are various ways of establishing a virtual circuit connection, including network-driven signalling, metasignalling (user-driven signalling), and provisioning (by subscription).

An ATM switch can handle two types of virtual circuits: a switched virtual circuit and a permanent virtual circuit. Control signals are just regular ATM cells with special (reserved) VCI/VPI values; they are used to set up and release virtual circuits. For a switched virtual circuit (SVC), when an ATM switch receives a connection request signal, it verifies that it has sufficient memory, port, and traffic capacity to handle the connection, examines the destination information in the request, builds it into a lookup table, and passes the signal on to the next system (which may or may not be a destination system). It is possible to pass cells on to multiple destinations (cell copying). Thereafter, all information that flows is identified and directed according to this information. When a signal (cell) is received that indicates a disconnection request, the ATM switch, upon receiving confirmation that disconnect will occur, erases the identification and path information for that connection. The above is analogous to phone call setup and teardown, and is suited to occasional bursts of information. For a permanent virtual circuit (PVC), there is always a dedicated path through the switch, that is reserved at all

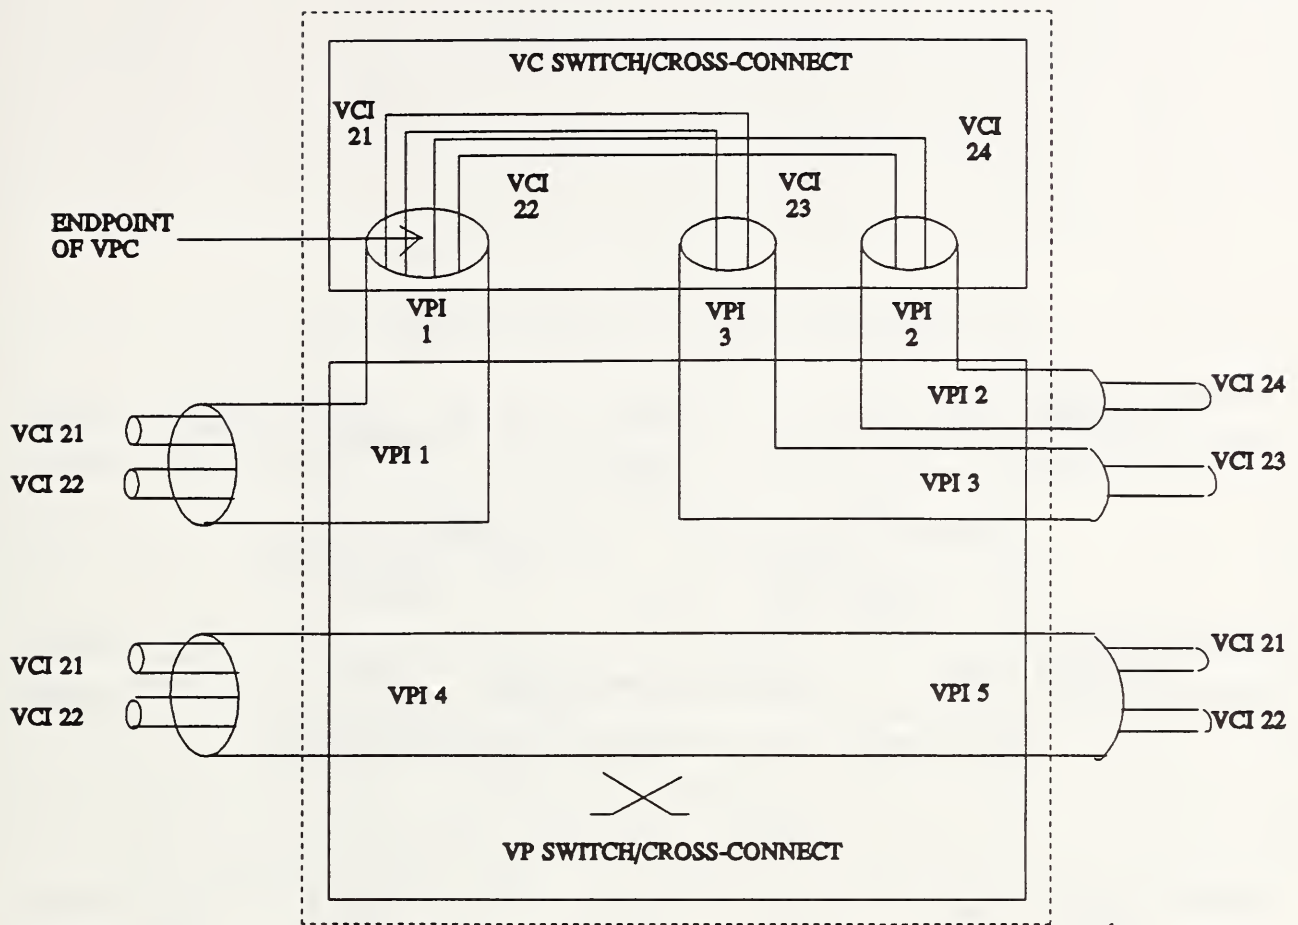


Figure 1b
Representation of VC and VP switching.

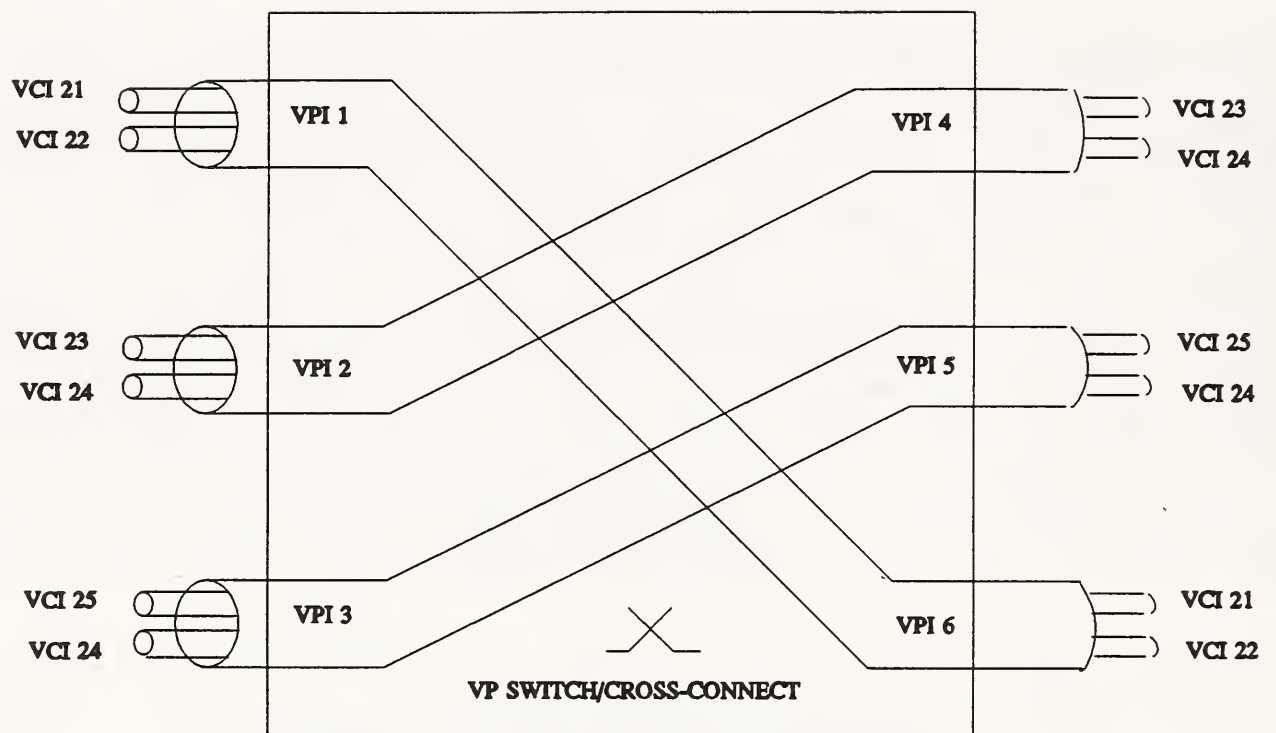


Figure 1c
Representation of VP Switching.

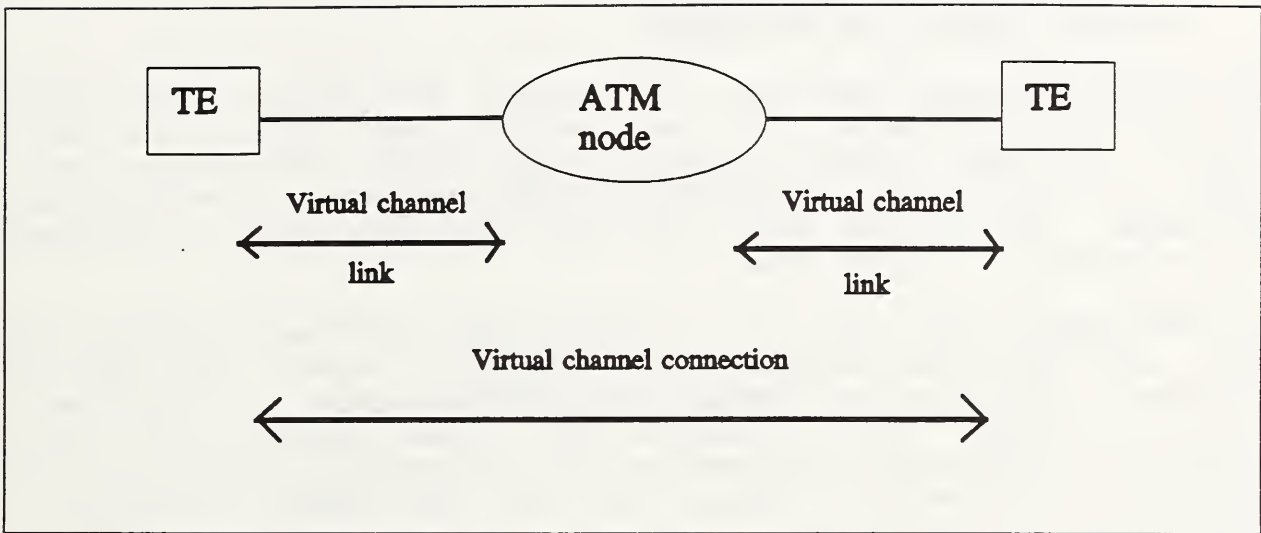


Figure 2
Virtual channel link and connection.

TE = Terminal Endpoint

times for information transfer; the necessary information is maintained in the lookup table. This means that input and output ports on the switch may be dedicated to a particular circuit, and are thus unavailable to other traffic. There is no connection setup as with SVC service. A PVC can be used for a steady high volume of data transfer (between fixed users and/or locations), and PVC service is faster than SVC service, but it is more expensive, and thus is used less often than SVC service. SVC service affords maximum flexibility in establishing dynamic connections. It is possible for a single ATM switch to have both SVC and PVC capability.

2.3 Physical Layer and ATM Layer

The physical layer (see fig. 3) supports the correct transmission of ATM cells on the bit streams leading into and out of the ATM switch. Various media (optical, electrical, satellite) may be used, but all of these must support the continuous transmission of ATM cells. The timing used, and thus the data rate, is dependent on the physical medium used.

The ATM layer (see fig. 3) is independent of the physical medium used, and thus independent of the physical layer. The ATM layer performs multiplexing of different circuits (identified by different VCIs) onto a single bit stream coming into the ATM switch, and performs corresponding demultiplexing of the bit stream leaving the ATM switch. Also, the ATM layer adds cell headers to cells passed from the ATM adaptation layer (see sec. 2.4). The translation of the VCI/VPI format to make them recognizable to a destination system by the ATM switch is a function of the ATM layer. The ATM layer performs cell relaying, delay handling (handling connections with different delay requirements), cell loss priority processing (determining which cells get dropped in case of overload), explicit forward congestion notification (determining the existence of congestion), and usage parameter control (traffic enforcement on the connection). Finally, the ATM layer may need to provide access flow control at the user network interface (UNI) (see sec. 2.6); the UNI connects a terminal endpoint to an ATM switch. The ATM layer offers the capability of realizing virtual wires (i.e., VC connections) in a dynamic way, which merely exist for the connection's duration.

2.4 ATM Adaptation Layer

Although nothing prevents a user from using the services of the ATM layer directly, they might not be directly usable for some applications. Thus, there can be a need to enhance the ATM service for distinct classes of applications; this is what the ATM Adaptation Layer (AAL) does (see fig. 3). AAL functionality may be




Figure 3
ATM Related Functional Layers

implemented at terminal endpoints (see fig. 2). The AAL enhances the service provided by the ATM layer according to the requirements of specific services. The AAL maps the user, control, or management protocol data units into the appropriate fields of the ATM cell and vice versa. For instance, in addition to segmenting and reassembling protocol data units (PDUs), the AAL can improve the quality of service provided by the ATM layer (e.g., by detecting lost cells). In sum, the AAL complements the ATM layer services.

To reflect the spectrum of services, four AAL types have been recommended by the International Telecommunications Union Telecommunications Standardization Sector (ITU-T, formerly known as CCITT): types 1, 2, 3/4, and 5. In addition, users can define their own AAL types to satisfy their own service requirements, if the ITU-T recommended types do not suffice. The service (mapping) requirements are categorized according to three parameters: the timing between source and destination, constant or variable bit rate, and the connection mode. Each AAL type performs enhancement according to a different set of characteristics, and thus requires different adaptation mechanisms to convert to ATM cells. The recommended types mentioned above may be used in other contexts than those mentioned below; uses stated below are recommended by ITU-T. Type 1 specifies a timing requirement, in that the terminal endpoints expect the recurrence of information from a particular user to be periodic. The bit rate is constant and the service is connection-oriented. A typical example is 64-Kbps voice, or circuit emulation. The adaptation must "remember" the recurrence relationships, and map the connection information to ATM signalling. Buffering and padding are used in this regard. Type 2 (not standardized yet) is (thus far) similar to type 1, but sources have a variable bit rate (such as variable audio). Thus the adaptation must accumulate a sufficient number of bits to be able to segment into ATM cells. In type 3 there is no timing requirement between source and destination, the bit rate is variable, and the service is connection-oriented (such as signaling information). Type 4 is identical to type 3, but supports connectionless capability (for example, LAN connectivity applications). Thus the adaptation must provide "datagrams" (self-contained units) on the ATM cells. Types 3 and 4 are so much alike that they were combined into type 3/4. In addition, there is a type 5, which is very similar to type 3/4 but may be more efficient (less overhead). It is possible to have an empty AAL for users who may find the ATM services sufficient for their requirements. The AAL for each type is divided into a SAR (segmentation and reassembly sublayer) and a convergence sublayer (CS). The functions of the SAR are the segmentation of a PDU into ATM cells and reassembling of cells into a PDU; these functions are common to all of the AAL types. The convergence sublayer is service dependent (that is, dependent upon the mapping requirements

dictated by each application).

For AAL type 3/4, two transport modes are possible: message mode service, where an AAL service data unit (equivalent to a PDU with header and trailer, see fig. 4) is passed across the AAL interface from the application in exactly one AAL interface data unit (PDU-like), and streaming mode service, which is identical to message mode except that multiple AAL interface data units are allowed per AAL service data unit. Segmentation and (later) reassembly then occurs on the AAL processed packet, as shown in figure 4. For both transport modes, the following two operations are possible: assured operation, and nonassured operation. In assured operation, all AAL service data units are correctly delivered by retransmission of erroneous or missing data, and flow control is mandatory. In nonassured operation, data can be lost or corrupted; there is no retransmission and flow control is an option. The nonassured operation is preferred, because simplicity is a fundamental condition to sustain high throughputs, and because, in most network configurations, a transport protocol with a complete recovery/flow control mechanism will be used.

2.5 ATM Network

There are two kinds of ATM devices: ATM switches, and ATM access devices. A network using ATM consists of one or more ATM switches, which connect different ATM access devices, where AAL functionality (and possibly "higher level" functionality) is implemented. ATM access devices send ATM cells that can be processed by the ATM switches. The term "ATM network" is used in this paper to denote a network using ATM as the underlying communications technology. Figure 5a indicates the position of an ATM network in an internetwork.

Since the ATM network is a subnetwork, the AAL functions are terminated at the ATM network boundaries, and so the AAL functions are not necessarily end-to-end. Thus a transport protocol is necessary to have full end-to-end interoperability. For example, a Transmission Control Protocol/Internet Protocol (TCP/IP) application may use an ATM network (including AAL services) to provide this end-to-end interoperability. Accordingly, the "higher layers" referred to in figure 3 denote various services or applications that may use the AAL. Examples of these are a connectionless service, a local area network (LAN) application, a broadband ISDN application (see Appendix A), or a user-defined or proprietary application.

2.6 ATM Interfaces

There are two levels of ATM interfaces: a network-to-network interface (NNI), and a user-to-network interface (UNI). AAL functionality, as well as "higher layer" functionality (e.g, routing, transport, and application functionality), may be




Figure 4
AAL

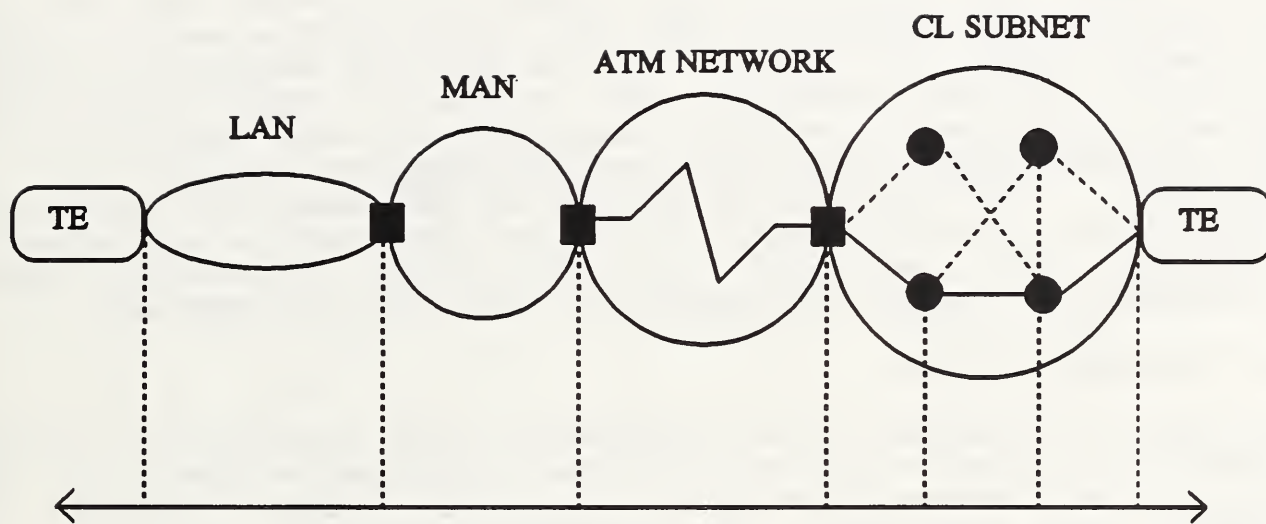


Figure 5a
An ATM Network in a combination of subnetworks.

MAN = Metropolitan Area Network
CL = Connectionless

implemented at the UNI. At the UNI, many terminals can share a given ATM access. Different services are integrated at the UNI, and each service may be associated with one dedicated terminal or a multimedia terminal. The interconnected terminals may form a customer premises network (CPN). When a shared medium is used for ATM access, there is a need to perform contention resolution and bandwidth sharing. There are many methods of implementing the UNI (see sec. 2.9); figures 5b and 5c illustrate three methods. The UNI is contained in ATM access devices. The NNI governs the interface between different ATM switches; it is more complex than the UNI and is yet to be found in many switches. Different switches may have differing requirements; the NNI must reconcile these differences.

Cell formats and signalling are different in the UNI and the NNI. The cell format leaving the UNI has a generic flow control (GFC) indicator field; the GFC field allows multiple terminals to share the same physical channel at the interface. This field is not present in the header passed between ATM switches. The ATM switch is responsible for recognizing and stripping this field, if necessary. The signalling at the UNI is concerned with connection control and administration; the signalling between switches is concerned with basic cell transfer. The ATM switch should recognize and change the signalling algorithm if necessary.

2.7 Features of ATM

ATM data is designed to be transported over high-quality digital lines with minimal cell errors; it is a "fast packet" service. ATM is an ideal transmission mode for bursty traffic prevalent in today's data networks. ATM provides bandwidth (capacity) on demand, and high bandwidth to every host. ATM reduces network response time, allows efficient network usage and provides for scaling to very large or very fast networks. ATM can accommodate both a wide range of transmission rates (each rate corresponding to a different information type, such as voice, video, etc.) and unpredictable traffic patterns. ATM can support narrowband as well as broadband uses. ATM connections can be point-to-point or point-to-multipoint. Small cell sizes minimize delay. Cell sequencing means that cells are delivered in order. ATM eliminates the need for separate voice, video or data networks. ATM provides circuit emulation (simulating a normal circuit), and allows flexible networks to be built using VP connections.

ATM can support a unified addressing capability among these systems. It can carry data, audio or full-motion video. ATM can be public or private, and can be included in wired or wireless systems. Total throughput is increased by adding more ATM switches and more links at higher rates. The simplicity of both the algorithm for forwarding ATM cells and the constant cell size leads to efficient hardware implementation. ATM is based on a

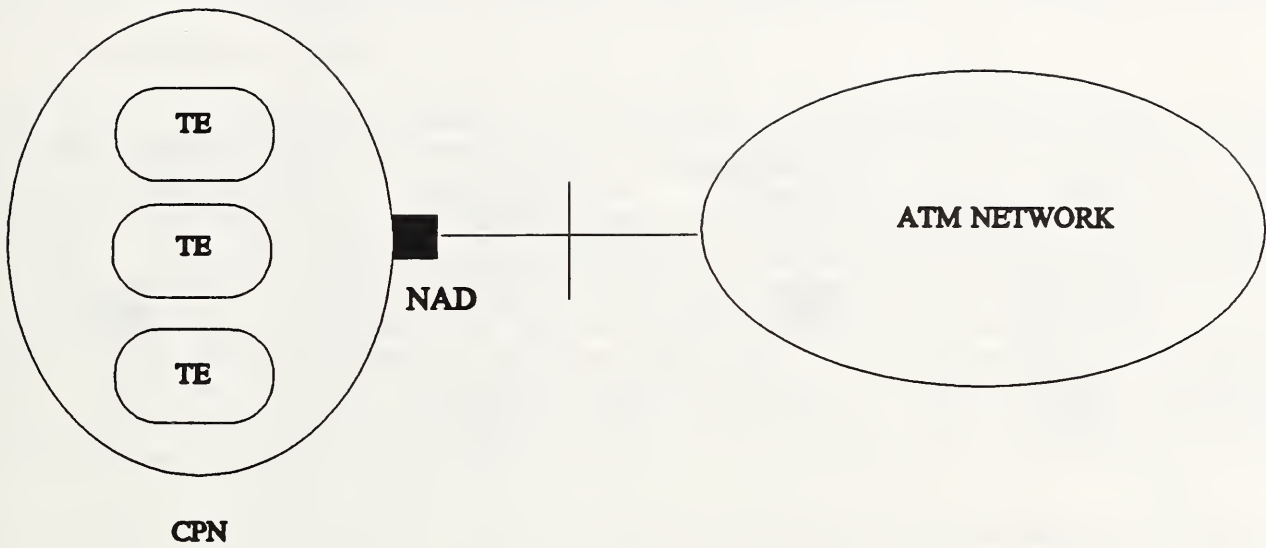


Figure 5b
A cluster of terminals sharing an ATM network access.

CPN = Customer Premises Network
NAD = Network Access Device

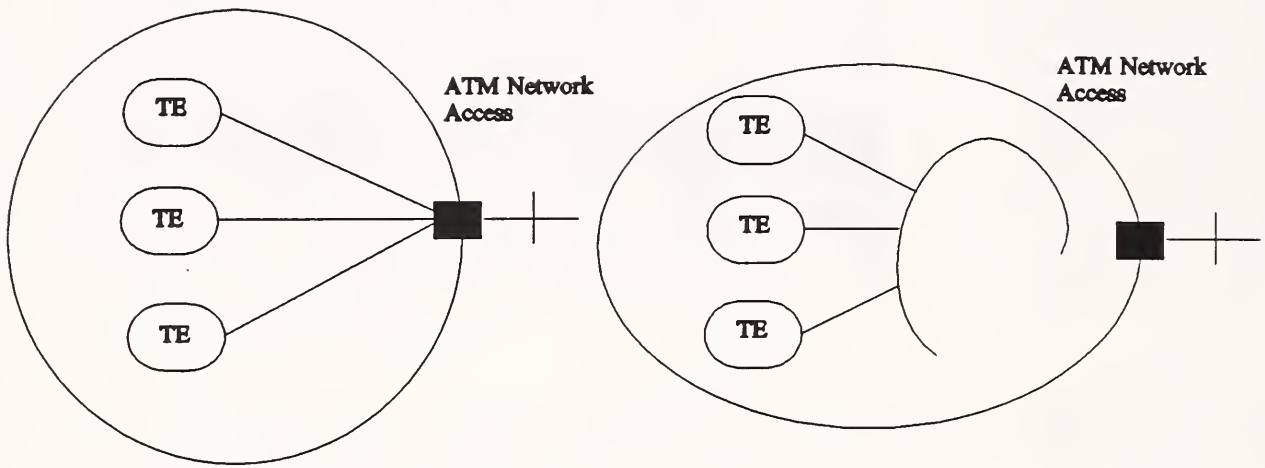


Figure 5c
CPNs with star and shared medium configurations.

CPN = Customer Premises Network

network of switches and dedicated links to each host, so that the aggregate bandwidth of an ATM network increases as hosts are added. Data transfer by a large number of hosts can occur in an ATM network. Since there is a large amount of bandwidth on an ATM network, a user should be able to send information on demand with minimal delay. A connectionless data service can be built on top of an ATM network; in this scenario, a predefined ATM connection is made at the UNI, and an "overlay" network of connectionless servers perform network layer functionality at terminal endpoints of the ATM network. ATM signalling at the UNI is not required, since that functionality is performed by the servers. Connection-oriented data services using ATM require ATM signalling and AAL protocols.

2.8 ATM Status Today

The ITU-T (CCITT) and American National Standards Institute (ANSI) Accredited Standards Committee T1 have completed standards for basic ATM-related capabilities. The ATM Forum and Internet Engineering Task Force (IETF) are tackling important ATM standards issues. The ATM Forum (sponsored by Interop) is a consortium which is promoting cooperation in the development of ATM implementation specifications, standards, testing, and marketing acceptance; it is made up of ATM equipment makers and many leading network vendors, and has over 300 member organizations. The IETF develops specifications such as the Internet Protocol (IP) to support the activities of the Internet.

In chronological order, ATM product offerings are likely to be (1) ATM local area networks (LANs), (2) campus ATM (connecting LANs), (3) ATM wide area networks (WANs), and (4) broadband ISDN and cell-relay services offered by subscription from public carriers. ATM LAN products exist now, but from only a few vendors. An ATM LAN comprises one or more ATM switches, dedicated high-speed circuits from each switch to networking elements, and the networking elements (such as workstations) themselves. All of the equipment is owned by a single organization, the distances involved are small, and the ATM switch may "broadcast" a cell to multiple destinations. An example of an ATM LAN is depicted in figure 6; ATM LANs differ from traditional LANs (e.g, Ethernet) in that they do not use shared media. In addition, ATM switches can be directly connected to LAN servers, LAN switches, or Ethernet hosts. ATM LANs can be used for client-server computing involving workstations and servers. Inhibitors to deployment are incompatibilities between ATM and NICs (network interface cards) for PCs and workstations, and the lack of development of a widespread user market for ATM LANs. These factors may set widespread deployment of ATM LAN systems back several years. Since campus ATM systems interconnect ATM LANs, hubs, bridges, and routers throughout a

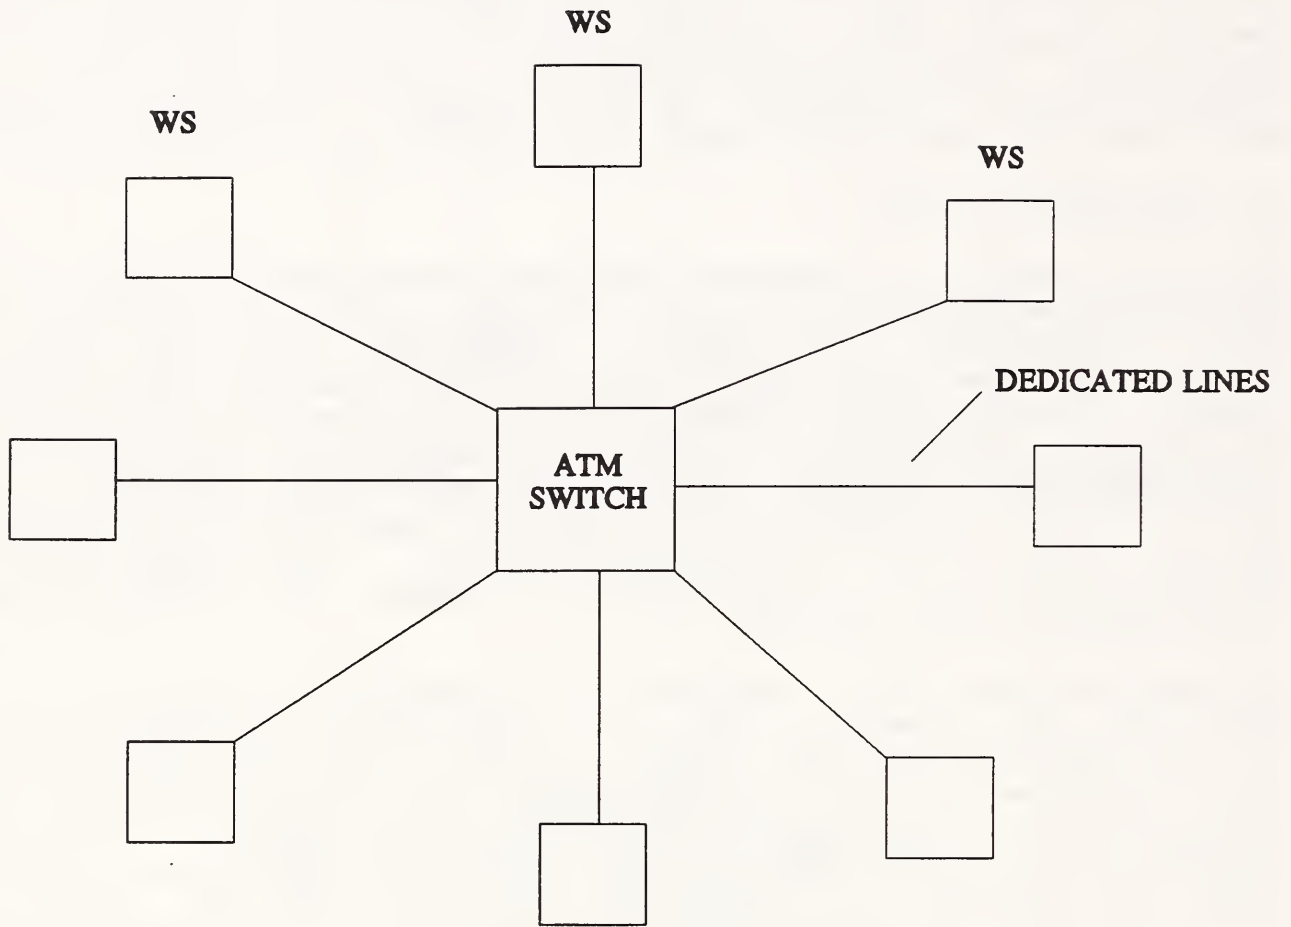


Figure 6
ATM LAN

WS - Workstation

campus, fewer campus ATM products exist than LAN ATM products. Widespread campus ATM products may also not be available for several years. A public ATM WAN is depicted in figure 7. ATM in the WAN refers to carrier service offerings based on switching ATM cells from a network-entry access location to a remote destination access location. The address information could be connection-oriented or connectionless. Widespread ATM WAN product offerings are a number of years away. The inhibitors to deployment include cost effectiveness concerns (economic feasibility), a possible shortage of applications which could benefit from ATM WAN technology, and the lack of availability of SONET (see 2.9). The introduction of widespread broadband ISDN and cell-relay services will not occur until near the end of the decade, and will await the evolution of broadband ISDN (see Appendix A) and the stabilization of access rates offered by all public carriers. In addition, signalling concerns and concerns over the variety of physical media may delay introduction of ATM on public carriers.

Currently, some users are buying ATM for supercomputers and high-end workstations (scientific uses). The Department of Defense (DOD) is involved in ATM procurement (e.g., the Army's private branch exchange (PBX) upgrade). The feasibility of ATM was demonstrated at Interop 94. Some users are developing pilot projects to test ATM interoperability. For example, the ATM Research and Industrial Enterprise Study (ARIES) project is a pilot project which will test the interoperability of ATM services and equipment from more than twelve vendors and carriers. The State of North Carolina has an ATM-based network (the North Carolina Information Highway), and the State of Connecticut has developed a plan for such a network. Nearly all switch, router, or hub vendors are in various stages of ATM product development. Vendors are also forming partnerships for ATM interoperability; some PVC interoperability testing has been done among ATM switch and ATM access vendors, as well as between different ATM switch vendors. There are pilots for LANs, campus networks, national networks, and prototype gigabit testbeds using ATM. The National Aeronautics and Space Administration (NASA) has a pilot ATM network using IP. The Department of Energy's ESNET is a national ATM network. Prototypes of ATM switching systems have been demonstrated to show technical viability. COMSAT has demonstrated the feasibility of ATM using satellite links. ATM networks are emerging at an opportune moment as high bandwidth applications and hosts are only now arriving on the scene.

Though promising, ATM has some obstacles. There are still some issues to decide regarding ATM signalling and ATM adaptation, and some additional AAL types are under development. Issues still to be decided include traffic control and connection usage enforcement. Adaptation introduces additional overhead and complexity. ATM may be costlier than other technologies.

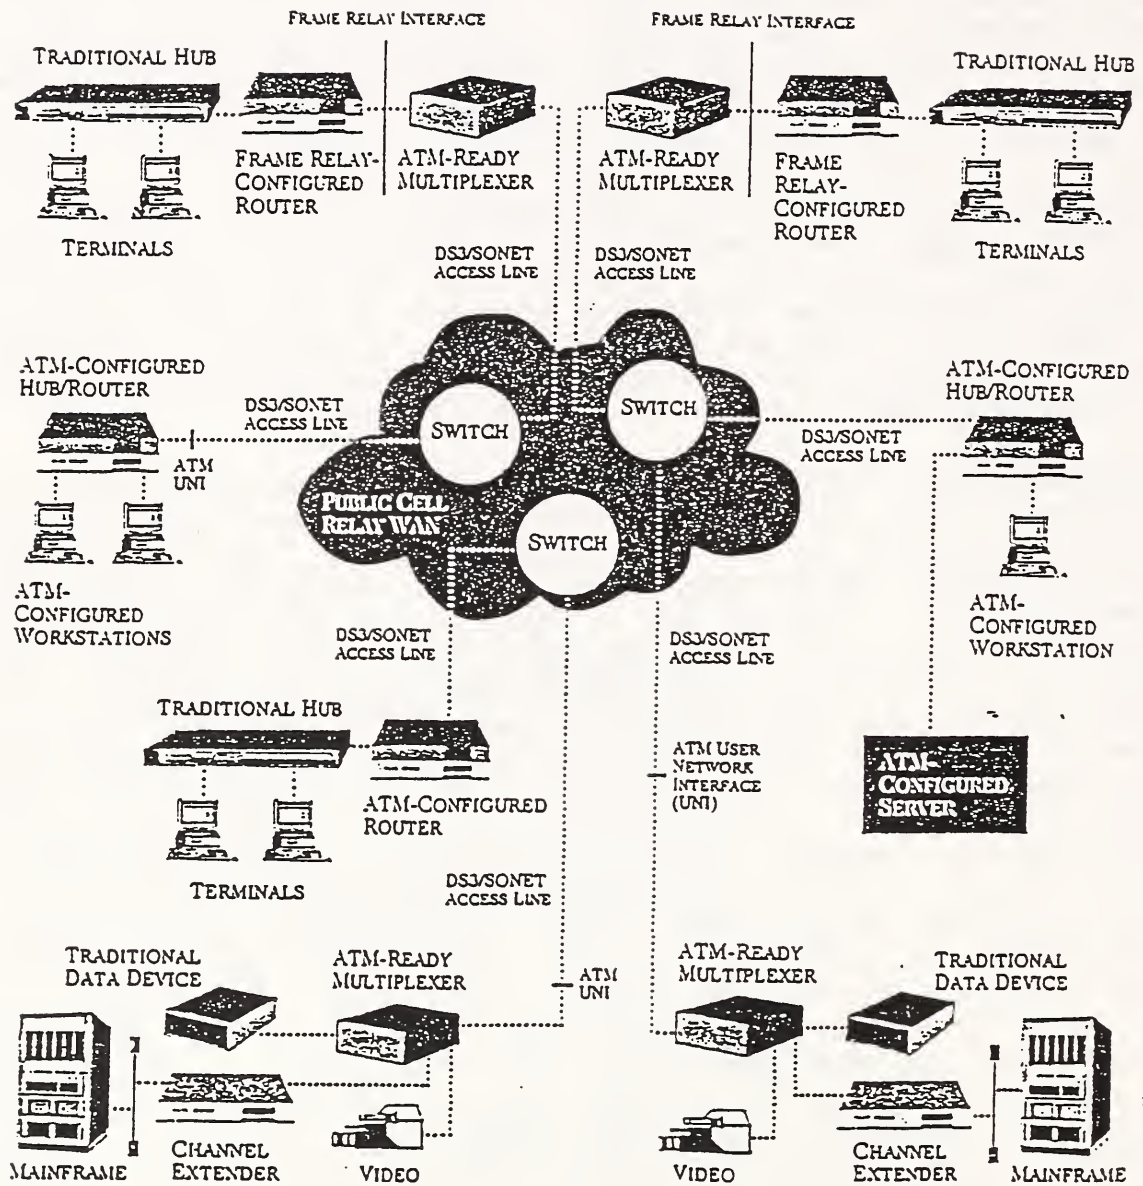


Figure 7
The Public ATM WAN: Task Links to Multiple Sites

Migrating to ATM is more complex than with products of other technologies. It is difficult to purchase ATM equipment from multiple vendors, hook it together and expect it to work. There are few standards for ATM switches. For interoperability, ATM access equipment needs to be from the same manufacturer as the ATM switch. It is prudent for large agencies to delay widespread installation of ATM until the above obstacles are overcome.

2.9 ATM Implementation

The ATM technology can be implemented in a number of ways; ATM supports many different physical layer media, including twisted pair, optical fiber, and satellite. The speed that is supported is dictated by the characteristics of the supporting physical medium; for example, certain lines can support speeds of 45 Mbps (T3), twisted pair can support speeds of up to 52 Mbps, and a fiber distributed data ("TAXI") interface can support speeds of 100 Mbps. ATM can use T1 lines (1.544 Mbps) also. Probably the best physical medium for ATM support is the optical medium of SONET. In fact, ATM was originally defined in the mid-1980's to run over SONET facilities.

SONET is a fiber optic transport system. The system combines a basic signal rate of 51.84 Mbps with a multiplexing scheme to create a hierarchy of higher-rate signals, with each higher-rate signal providing an integer multiple of the basic signal rate. SONET can transfer traffic from one network to another. SONET speeds are multiples of 51.84 Mbps; common multiples supporting ATM are 1, 3, 9, and 12 (particularly 3), but a multiple supporting ATM of 48 is possible in the future. The syntax used to denote the speed is commonly represented as OC-n, where n is the multiplier. SONET networks may be public or private. SONET capabilities are not yet widely available. There are different versions of SONET; one version that is popular is SDH (Synchronous Digital Hierarchy). ANSI has developed standards for SONET capability.

Other physical layer implementation issues are important in addition to the choice of physical layer media. For example, ATM can support inverse multiplexing of cells on a physical medium, use trunk lines (to connect PBXs), or use T1 lines to provide LAN connectivity.

Implementation decisions involving ATM include how many ATM switches to interconnect, and how to provide access to the ATM network through the UNI. For ATM access, one choice is to provide ATM access via an Ethernet switch, which can provide dedicated access to multiple Ethernets through one ATM UNI port. ATM access devices can also have LAN ports directly configured into them, for LAN-LAN connectivity. Transparent asynchronous transmitter/receive interfaces can be supported at the ATM UNI. ATM access can be

implemented in host interfaces, ATM multiplexers, ATM concentrators, ATM hubs, ATM servers, ATM adapters for workstations, DSUs connected to routers, standalone ATM routers, ATM drivers, and other ATM WAN access devices, all connected to an ATM switch via the UNI. These devices differ in the functionality they offer the end user, but all can provide access to an ATM switch; some of these devices may be able to place SVC calls or connect to LAN servers. ATM access devices may be standalone or module-based, and support various numbers of ATM interfaces at the UNI per module. In some case, these devices may be connected to a NT1 (network terminating device); in broadband ISDN terminology, these devices are part of the user's CPE (customer premises equipment). Prices vary, but the current price of an ATM adapter is approximately \$1000, the price of an ATM hub approximately \$4000, and the price of an ATM switch up to \$400,000 (\$40,000 per port, or depending on configuration). Other ATM access devices can cost up to \$25,000. Per-port prices for different ATM interfaces offered by a single switch provider do not vary significantly, but there are wide disparities when considering offerings of different vendors; some ATM carriers may have higher prices than other types of ATM vendors. These prices are subject to change.

ATM switch implementation can be implemented in a small workgroup switch, which could be supported by users who want a small ATM, or in a PBX-like switch. Both of these switch types have fewer numbers of ports and less complexity. ATM does not have to be centralized in a huge switch. The size of the ATM "backplane" is proportional to the number of interconnected ATM switches. ATM switches are multislots, using a "rack-mount" or stand-alone chassis. All ATM switches support PVC capability, and many are adding SVC capability in addition. For reliability reasons, ATM switches may offer a redundancy in components. In practice, ATM switches may have more capacity than users can tap, because the number of ports that can concurrently access the internal switching matrices is limited.

Traffic characteristics are important in determining how ATM is implemented. For example, steadier traffic ATM could be carried over dedicated T3 lines, whereas, for bursty traffic, ATM devices could support a connectionless capability (e.g., support of switched multimegabit data service (SMDS)).

Entry-level ATM can run at speeds of 25 Mbps in a client-server environment. Desktop ATM can run at 155 Mbps. ATM using metallic interfaces can run at 1.544 Mbps (called "T1" or "DS-1") or 45 Mbps (called "T3" or "DS-3"). The speed is determined by the physical medium used.

An example of ATM implementation in an enterprise network is depicted in figure 8.

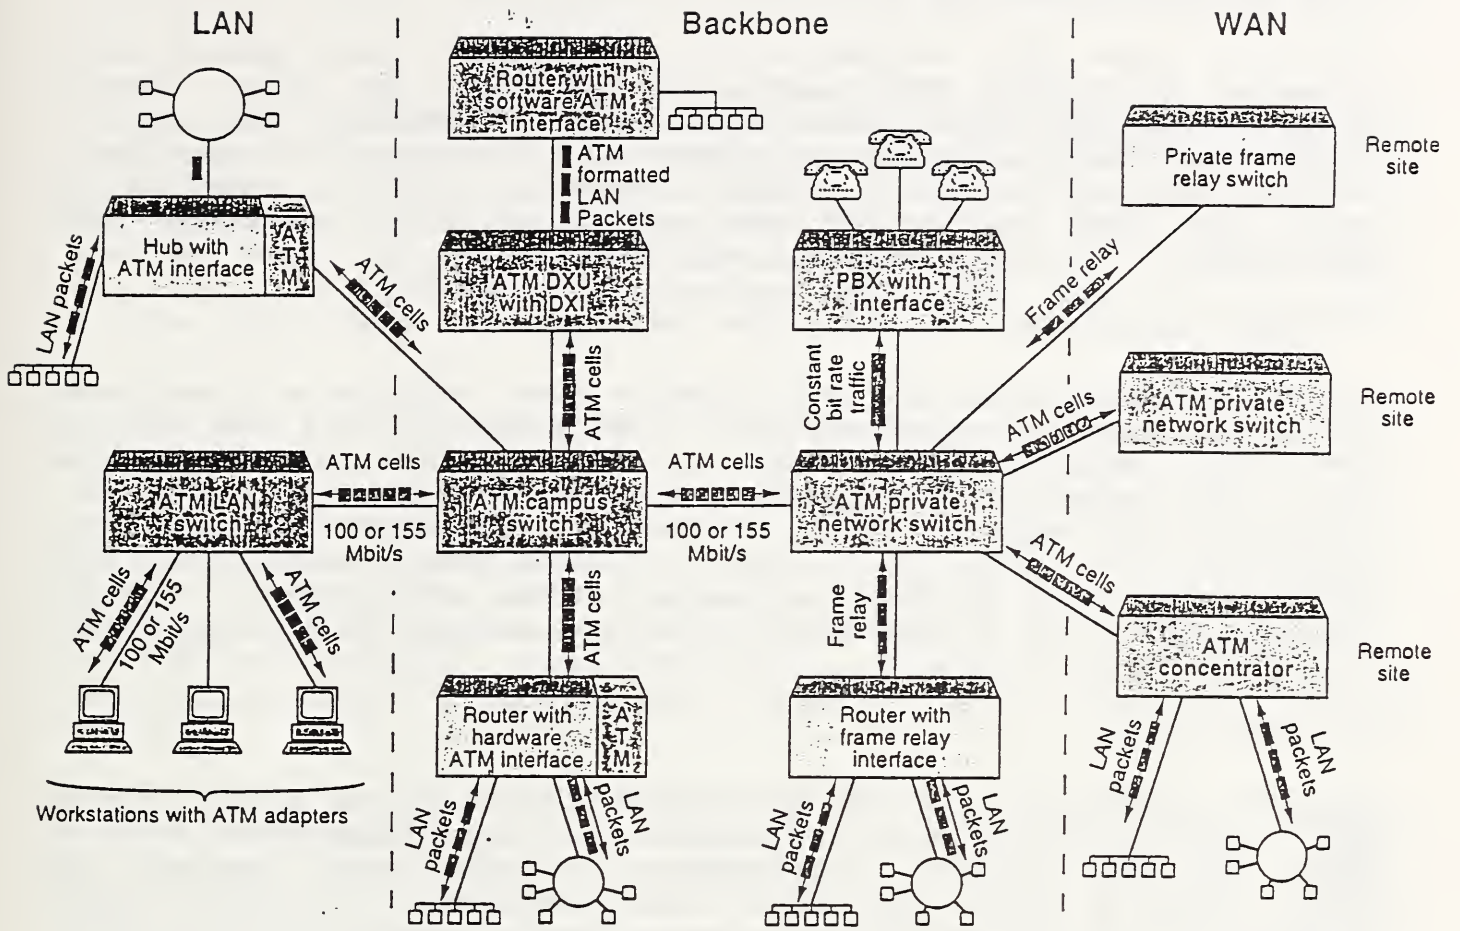


Figure 8
ATM Enterprise Networking

2.10 ATM Applicability

Assuming users have access to an ATM switch, when should users consider procuring ATM access equipment? If users have a requirement for high-speed bandwidth on demand, or have a requirement for multimedia, or want to link a number of LANs, WANs, metropolitan area networks (MANs), and other kinds of networks and units, or want to realize the advantages of a broadband ISDN environment, then use of ATM should be considered.

3.0 User Decisions Regarding ATM Products

There are a number of important issues about which users must make decisions when procuring and installing ATM products. The choices that users make will affect the operation of the ATM product for them. Following is a list of concerns and the factors to be considered for each concern. Users should understand the factors affecting each issue, decide which factors are important to them, and execute their procurement/installation accordingly. Certain factors which may be important to some users may not be important to other users.

3.1 ATM Access

The characteristics of ATM access are important to users because user applications have specific transfer requirements and need to know what capabilities to request of an ATM network through the UNI, including type and number of connections. In most cases, users will not procure ATM switches directly, but will have access to ATM switch service provided by others. Users need to determine the types of physical media in their ATM environments, because this affects cost and performance. For example, ATM can be carried over one of the SONET versions, single-or-multi-mode fiber, or a metallic interface.

Another user concern is the kind of ATM access equipment being purchased. Users should buy ATM access equipment appropriate to their configuration, application and functional requirements, and appropriate to interfaces with other technologies coexistent with ATM. Users need to look at the state of ATM product offerings to meet the above requirements. For example, users may want to replace private network connections with access to public ATM services; users should ensure that those services are available before considering such a replacement.

Users should be concerned about the number and type of ATM access ports (bit streams) supported by the ATM access device. ATM access ports can be directly connected to multimedia workstations and central servers, or can be connected indirectly to hosts through shared-access LANs such as Ethernet. ATM access devices can have

other ports (connected to diverse devices) in addition to ATM access ports. ATM access devices could be connected to gateways, routers, or hubs. Synchronous interfaces can be built into an ATM access device.

Users need to determine the possible destinations of ATM connections in their ATM access devices. For example, ATM access devices can link to local telephone companies and cable system providers. In an enterprise network there are a number of different devices connected via an ATM network to an ATM access device, and the ATM access device must recognize these other devices. The more types of devices recognized, the more flexible is the ATM access device.

Another user concern is the number of ATM access ports to assign to a specific ATM access device. The more ports on the device, the more users and traffic can be supported, but the more complex the device. If more ports are requested than are actually used, then resources will be underutilized. However, it might be desirable to have the extra ports built in order to dynamically handle varying traffic loads. The number of users per port is also an issue. Increasing the number of users per port will increase the contention for ports in the access device. The speeds supported by the ATM switch will affect the number of ports needed on an ATM switch.

Users should ensure that the ATM access devices are able to handle a wide range of data rates with a high quality of service. There may be more than one way of providing ATM access to an ATM network for a particular application. Cost may also be important in selecting ATM access capabilities. Users should specify speed and cost requirements for ATM access devices as appropriate. Users may also need to specify call control (ATM signalling) capabilities supported by ATM access devices.

3.2 ATM Switches

If a user needs to procure an ATM switch, users should specify the number of input and output ports supported by the switch, as well as the speeds supported for each port. There are several ATM switch types, and not all of them interwork. For signalling purposes, an ATM switch needs to know to which kind of ATM switch it is connected. Users should determine to which switch types their ATM equipment will be connected, and ensure that the switches can effectively interwork. A specification on inter-carrier interoperability nearing completion by the ATM Forum will likely promote ATM switch interoperability. All ATM switches support PVC and UNI capability, but more and more switches are supporting SVC and NNI capability; users should specify if these additional capabilities are needed.

ATM switches must (if the application demands) be able to handle

SVC calls, and must be able to handle PVC calls if necessary. Applications where SVC capability is important include videoconferencing and transmitting trading information, as well as transmitting 3-D images over multiple routes to multiple destinations.

If ATM users are losing data because the switch software cannot keep up with huge data volumes, then the problem is probably a lack of coordinated buffering and congestion control schemes among the different vendors' equipment. Users should then ensure that their switch types minimize this possibility by working with vendors to develop an approach to solving this problem. The ATM Forum is developing specifications for flow control and buffer management to alleviate congestion; this functionality is being implemented in some switches.

Ports on ATM switches should be flexibly assigned, and must keep traffic counts on different service classes so that usage-based billing can be supported. ATM switches can have non-ATM interfaces (e.g., Ethernet) directly connected to them. ATM switches must also segregate traffic with differing delay-tolerance characteristics (for example, file transfer vs. interactive video). The time delays supported by ATM switches vary widely from nanoseconds (jitter) to microseconds, and could pose problems for continuous bit-rate traffic, such as voice and videoconferencing. Users need to specify acceptable delay-tolerance characteristics for ATM switches.

Users should decide (based upon management, security or other requirements) whether to request full-bandwidth broadcast or multicast capability in their ATM switch products. Broadcasting involves sending a single transmission from a source to everyone on a network, and multicasting involves sending a single transmission from a source to a selected group of recipients (not necessary everyone). Broadcasting and multicasting allow simultaneous transmission to multiple hosts. Fully-loaded switch ports can participate in multicast calls. Cell loss is rare, and cell delay minimal.

3.3 Security

When procuring and installing ATM equipment, security is important because users want to ensure that data that passes through an ATM switch has its integrity preserved, is subject to access control restrictions, and is subject to authentication and identification restrictions. The scope of security for ATM includes aspects of network security, as well as administrative and operational security. ATM products are still in their infancy, so a number of features have yet to be defined.

SONET services may offer users service protection for high-capacity circuits in emergencies. SONET can be implemented with a double

ring backup capability, with two duplicate connections rather than a single connection; if one completes acceptably, the other is discarded. For security, there is also a SONET Ring Service, which lets users design a special fiber optic ring for their special switched access services.

High-speed ATM traffic poses difficulties with network security. Current methods of data encryption are unable to keep up, because in real-time processing, encryption times are too slow for the data rates. With ATM, (possibly replicated) cells can be sent on different paths to a destination, thus "scattering" the data; this is a possible network security mechanism. Users should ensure that their network (and other) security concerns are satisfied in ATM products.

There is work underway at the National Institute of Standards and Technology (NIST) to further define security issues for ATM and other emerging telecommunications technologies. Users should reference this work as appropriate.

3.4 Management

Users want to access resource management, performance, accounting, and configuration information for ATM networks. They will want to know, for a call placed through an ATM access device, did the call complete? The payload ID field of an ATM cell header distinguishes management information from user and other signalling information in an ATM cell. Users should ensure that management concerns (such as resource management) are satisfied by placement of appropriate information in this field. Resource management involves the proper allocation of memory, processors, and devices to support the requirements of an ATM network.

Performance is important in ATM networks. ATM networks must be able to ensure consistent performance in the presence of widely varying traffic characteristics, because many high-speed applications (video) require guaranteed throughput.

ATM access devices are responsible for guaranteeing good performance for a mix of traffic types in an ATM network. ATM access devices can perform traffic management. Traffic prioritization and load balancing are examples of performance metrics for ATM access devices. Other performance metrics include device data rate (throughput) and full (peak) ATM cell rate, both in number of cells/second, number of PVC and SVC connections supported per module, and the time it takes to set up an ATM connection. ATM access devices can also offer different connection qualities of service. For example, the connections used for signalling can have a preferential treatment in the case of overload. Users should specify a desired quality of service. ATM access devices should be efficient, in that all existing paths should be completely filled with cells before new connection

requests occur. ATM connections are refused if insufficient bandwidth is available at connection request time, so users should specify the bandwidth reservation capability.

ATM switches may be implemented with inadequate congestion management capabilities. The buffers may be too small and they may discard data when the network becomes too congested. There is also concern about traffic delays. An emerging specification from the ATM Forum describes traffic management and the use of congestion control indicators; this should alleviate some concerns. An alternative to a congestion management algorithm is over-provisioning of existing lines. Cell delay variance, mean cell transfer delay, cell loss ratio and basic error rate are important switch performance metrics. Cell loss may be due to buffer overflows (see section 3.2), physical medium problems, or errors in the cell routing. ATM switches usually discard any cells in excess of the negotiated capacity. A cell loss priority field in the cell header controls which cells are lost. Users should ensure that these rates are kept low (for example, around 10^{-8}) by appropriate detection and correction measures. It is possible to recover from cell loss. Other ATM switch performance metrics are the number and type of input ports supported, the number and type of output ports supported, the traffic load (for example, number of simultaneous connections) supported, and the cell switching time. In terms of performance, an ATM switch's output buffer capacity may vary from 240 to 4000 cells. Currently, wide-area ATM switches have buffers with capacities of only 30 to 60 cells. The ATM switch also needs to enforce traffic "regulations" by means of traffic usage parameters tied to VCI/VPI values. Implications of the above for users are that users should ensure that these concerns are adequately addressed, either in AAL functionality or in a transport service.

Configuration may involve setting up buffer management parameters. Buffer management is important, as mentioned above, particularly in a wide area network environment. Buffers are not as important in ATM LAN environments, because LAN traffic is fairly consistent. Specific buffer modules correspond to specific physical media attached to the ATM switch. Different types of traffic require different size buffers and different algorithms for managing each class of service. Based on user-defined criteria, the ATM switch may choose traffic to be discarded. Buffer management ensures that different types of traffic will be handled according to the quality of service each requires. Users should specify appropriate buffer management capabilities.

Users may want user-configurable calling capabilities, as well as automatic call control for signalling purposes. ATM switches should be automatically configurable for error recovery. For example, if there is a "break" in the network (e.g., a destination switch fails) an alternate destination can be specified in the lookup table. ATM access devices should be self-initializing,

programmable, and critical software should be stored in non-volatile memory. ATM access devices can be configured for automatic routing. Users should ensure that configuration requirements are met in ATM products.

Accounting information should allow the user to monitor the status of ATM activity, such as the current traffic status, device utilization, or existence of network failures. Since ATM is still an emerging technology, the exact nature of accounting services is still unclear. A useful service would be customized billing. ATM has experienced troubleshooting difficulties, and appropriate diagnostic tests are not yet in place, partially because the parameters of behavior on an ATM network are not yet completely determined. Users should specify appropriate accounting procedures for their ATM access devices.

ATM switches and ATM access devices may be managed via SNMP agents. ATM products can access network management capabilities via SNMP over ATM cells, by a local Ethernet interface, or by a local serial port.

3.5 Additional Procurement Considerations

Users should require the ATM vendors to conduct all appropriate interoperability tests. Users should conduct basic interoperability tests as part of acceptance. The P-NNI (Private Network-to-Network Interface) is being defined by the ATM Forum for interoperability between multivendor equipment using SVCs. Users should reference this specification if appropriate. In addition, the SONET Interoperability Forum (SIF) has been created to demonstrate multi-vendor SONET interoperability. All interoperability testing adds to the cost of procurement, but may be money well spent for the success of the procurement.

Users need to decide whether to add ATM access to existing systems, or to procure entirely new ATM access devices. Users should ensure that ATM access devices support application addressing requirements. Users also need to specify which AAL types (for example, AAL type 3/4) if any to employ. Appropriate application programming interfaces (APIs) should be defined for ATM networks. A carrier-based ATM network must be finely tuned to each user's needs. It is unclear what supplementary services will be added to ATM products; users could drive the creation of such services. IP interfaces could be configured into ATM devices to permit access to TCP/IP applications. Forward compatibility should be maintained with future ATM-based systems, and backward compatibility should be maintained with existing configurations.

ATM access should be made transparent to (hidden from) the application (in other words, just another circuit connection). The application is insulated by access devices which offer the flexibility of supporting a wide range of WAN services by adding or

chaining feature modules. Users don't have to change applications when the underlying networking infrastructure changes (for example, from ordinary circuit-switching to use of ATM cell switching).

Cost is an important consideration in ATM. ATM access devices should monitor the status of connections to keep the total costs down. It is unclear how the business market or tariffs for ATM access will emerge. For example, it is unclear how aggressively public ATM carriers will price their services, or whether higher-priced cables will be replaced with consolidated ATM access at a reduced price as an inducement. Interexchange carriers may cut prices competitively to gain business. There is a big difference in the per-port cost of LAN and WAN switches to ATM switches because of traffic management and buffer management differences between the two. Consolidation of services and equipment should save ATM users money. Users can reduce access costs by aggregating local and long distance traffic on a single service. ATM charges may be a flat monthly fee, be usage-based, or be speed-based. Currently ATM access is provided by carriers on a special service agreement basis. Users want simple ATM pricing, which is distance-insensitive. Installation fees for ATM networks are approximately \$5000; prices are subject to change. Generally, the higher the speed of the physical medium, the higher the monthly charge. Competition is driving down the per-port price of ATM switches. ATM users may need to negotiate with telephone companies, ATM equipment makers, and even application software developers to get the best prices and most functionality in their procurements, and users should monitor ATM regulation and pricing issues as they are resolved. It is desirable for long distance carriers, local phone companies, and competitive access carriers to come to agreement regarding ATM pricing. There should be a single point of contact for circuit provisioning, network management, and billing. Users can make choices in ATM network configuration and access which can save them money; for example, ATM access provided by a local carrier may be less expensive than ATM access on a dedicated line.

4.0 ATM Migration/Coexistence

The coexistence of other technologies with ATM is important. This is because the availability of other technologies along with ATM allows users to make their networking decisions based on technical factors, rather than worrying about whether the sites will be able to communicate with each other. For example, ATM switches support non-ATM interfaces (allowing both constant and variable bit rate data to be communicated), and any "higher level" functionality protocol (or a protocol which can be encapsulated in such a "higher level" protocol) can operate using an ATM network; examples are Ethernet, FDDI, and X.25 (as a data link) and IPX, DECNET, and Appletalk (at a higher level). The IETF has issued a Request for Comments that defines IP over ATM; the IETF is taking a leadership role in mapping its protocols to ATM. It is also possible to run the Connectionless Network Protocol (CLNP) over ATM. In order for

a protocol to use an ATM network, the AAL needs to segment, reassemble, and "frame" packets; addresses need to be converted to ATM layer addresses, and the "connection" features of the protocol need to be mapped to ATM connections. The ATM switch may need to "reconstitute" control information (such as congestion indicators) for the protocol using the ATM network.

Interest groups for other technologies (such as frame relay) are developing specifications for promoting interoperability of their technologies with ATM. These specifications are being evaluated by the ATM Forum, and include provisions to translate cells/packets of other technologies to ATM cells. Connectionless SMDS may be combined with the bandwidth guarantee of ATM's virtual circuit technologies, and frame relay can pass data between routers and ATM switches. There are additional methods of linking other technologies with ATM. Legacy protocols (e.g., Systems Network Architecture protocols) may be carried over ATM circuit-emulation services (using ATM access devices) which mimic existing support of these protocols. One way of doing this is protocol encapsulation, but this may be complicated.

Users can choose a number of strategies for migrating to ATM networks. For example, entry-level client or server ATM adapters can be used for smaller networks or systems to introduce ATM access gradually. It may be necessary to transition from the statistical multiplexing of ATM switches (see sec. 2.2) to traditional time division multiplexing (to support certain LAN plus videoconferencing applications). The conversion can be in an ATM access device (preferable), or the carrier can do it. From existing T1 lines, migration is possible to wide area ATM with T1 or T3 emulation. This would enable users to connect existing applications transparently. Carriers and equipment makers could migrate their systems to support ATM networks in different ways, depending on the degree to which their networks differ from ATM networks. For wide area networking, carriers are looking for major cost reductions as they combine various data services over a single ATM backplane (see sec. 2.9).

Migration to ATM networks depends upon whether an existing system is LAN or WAN. For WANs, migration is fairly straightforward. Standalone WANs will be replaced with ATM access devices (multiplexers) to public or private ATM modules supporting an ATM backplane. Existing technologies (T-1, frame relay) would be feeder technologies instead of primary technologies. IP routers can be replaced with ATM hubs. Internet access can be supported by ATM access devices.

For LANs, migration to ATM networks is more complex, depending on base technology. Each base technology offers its own evolution path. Ethernet switched hubs support ATM access, at a fixed price per port. A user may need to choose between routers and LAN switches for providing LAN connectivity to an ATM access device.

ATM LAN migration strategies include use of ATM interface modules and switched Ethernet. Users should require vendors to propose a migration strategy to ATM LANs if appropriate. ATM Forum specifications address "LAN emulation" over ATM; LAN emulation lets existing LAN types be connected transparently, usually in a bridging mode.

Some users are reluctant to rely on a single-vendor solution, especially before standards for advanced features (such as virtual LANs) are complete. Users over the next two or three years are preparing themselves for ATM with infrastructure improvements (e.g., installing fiber) and strategic plans based on parallel (ATM pilot) networks that handle production traffic as the technology and the vendors prove themselves. Some users cannot afford to invest in pilot ATM networks, and are waiting for other users to lead the way. Other users are employing incremental strategies by employing MANs (e.g., SMDS) or narrowband ISDN as a stepping stone to a full broadband ISDN environment. ATM access may be implemented gradually, for example, starting with entry-level 25 Mbps ATM adapters, then moving to desktop 155 Mbps ATM drivers, ATM LANs, campus ATM, and then WAN ATM. Users need to develop a long-term plan to implement the ATM technology, to take advantage of ATM benefits.

5.0 Summary and Future

The ATM technology is a unifying communications paradigm; it insulates protocols from underlying changes in the physical medium. Use of the ATM technology promises:

- cost savings;
- speed improvements;
- flexibility gains in the design and operation of networks;
- consolidation of voice, data, video and other information types.

Whether the user's need is for wide area connectivity or LAN links, ATM promises to bring the swift and sure transport of information. Any application or service can be supported by cells on an ATM network. ATM makes the interface between LANs, WANs, and MANs transparent to users. ATM provides better application integration and speeds up application development.

By far the most important advantage that ATM offers over competing technologies is its open-ended growth path. As a LAN architecture, ATM is not locked into a single physical medium or speed. Instead, it can utilize any combination of speeds, from 25 Mbps to SONET OC-48 (2.5 Gbps), and a variety of media, including metallic and optical (SONET). There are many opportunities for cost reduction in designing ATM networks. ATM can be implemented without affecting the local user environment. The ATM technology

separates the basic transfer mechanism (that is, cell relaying) from policies governing routing, quality of service, and security. This separation facilitates the efficient interconnection of administrative domains. The broadband ISDN environment (see Appendix A) allows the user to request services transparently of the network, and to be assured that these services will be provided.

ATM using wireless capability and satellite links is under very serious consideration. In these scenarios ATM could be the "glue" for interconnecting networks to build the National Information Infrastructure (NII).

ATM networks provide access to centrally stored video information through the network, two-way or multipoint videoconferencing, and remote classroom instruction using video. For distributed computing, ATM is being used as the interconnection network for a distributed memory multiprocessor. ATM networks can also be used in:

- LAN interconnectivity;
- publishing (newspapers, journals);
- healthcare (such as telemedicine);
- multimedia libraries;
- financial applications;
- government;
- entertainment.

Cable companies are experimenting with ATM for delivering real-time interactive TV and video on demand. Some ATM LAN applications are:

- transport of real-time video;
- development of new transport protocols;
- parallel computation on ATM-connected workstation clusters;
- development of distributed multimedia applications.

For multimedia, the real-time traffic can be sent directly over an ATM connection, with or without AAL processing, as appropriate. Corporate applications using the ATM technology include:

- desk-to-desk videoconferencing;
- multimedia conferencing;
- multimedia messaging;
- distance learning;
- imaging tasks (including CAD/CAM);
- animation;
- cooperative work (for example, joint document editing);
- desktop applications for enterprise networks.

Alternatives to ATM do exist (including narrowband ISDN, cell-based

MANs, such as SMDS, frame relay, and LAN switching hubs). ATM is most likely to succeed if (1) there is a LAN-based ATM presence in the workplace and an established need for very high bandwidth over the wide area, (2) the wide area cost is reasonable, (3) wide area ATM service is a ubiquitous and uniform offering (a lesson learned from narrowband ISDN), and (4) there is a graceful migration from other technologies and services to ATM. Early users of ATM will be those for whom new enterprise-wide applications are possible using ATM that are currently impractical with other technologies. The ATM LAN and campus ATM market will be much more active in the short term than the wide area ATM market, which is expected to develop slowly, although due to corporate downsizing and the success of frame relay, users would like to replace their private networks with access to high-speed public WAN services. Prices for ATM service will likely decrease in the future, particularly for local ATM access. ATM is not yet available from all major local and long distance carriers. Thus far user demand for ATM LAN and campus ATM has yet to reach projected levels, due to (among other reasons) incomplete specifications, lack of interoperability, and cost.

Broadband ISDN (see Appendix A) will continue to evolve and change. SONET may continue to serve as the physical layer in the broadband ISDN environment, but the large number of choices possible may delay progression of the broadband ISDN technology in the wide area networking market. In the future, the broadband ISDN architecture will offer public and private network operators an affordable alternative to bridge the gap between LAN-based ATM products and large ATM central office switches. Deployment will be permitted as close as possible to points of greatest customer demand and will allow easy expansion as usage patterns change and grow. As a result, ATM service providers will not be required to make a high initial investment that is commonly required when establishing ATM service offerings.

In terms of future technological advances, by placing call processing in software rather than in hardware, substantial flexibility and generality will be achieved without sacrificing throughput. Positioning will occur for future growth, and bandwidth management will be addressed more easily in the future. ATM research items for the future include:

- signal interfaces;
- cell-pacing circuitry;
- cell-splitting;
- multiport ATM routers;
- call control signalling for broadband ISDN;
- generic flow control;
- operations and maintenance;
- resource management (queueing and congestion control);
- high-speed switching;
- metasignalling.

As computers become faster and ATM interfaces move "closer" to the CPUs, applications will be able to access the network at the standard SONET speeds of 622 Mbps, 2.5 Gbps, and above. SONET will have an economic impact, in that the existence of a standard for optical transmission will eventually make optical-electrical equipment a commodity, and commodity pricing will speed deployment of fiber-based facilities by public telecommunications providers. The SONET physical layer will likely play a role in the gigabit networks of the mid to late 1990s.

It may be some time before users make a complete migration to ATM (or achieve a complete coexistence with ATM). There is resistance from the operational side, and standards and management issues must be resolved. Users are reluctant to place mission-critical applications over ATM networks. ATM technology is not yet completely tested (especially for SVC capability) and has yet to prove its economic and technical viability. In addition, ATM products are not yet widely available. The extent of ATM product deployment in the future is a question. A full spectrum of ATM products will not likely be available for several years. There is the possibility of divergence between the computer and telecommunications industries in ATM development, divergence in different regions of the world in ATM implementation, and divergence of relevant ATM standards (such as those produced by ITU-T (CCITT) and those produced by ANSI). As with narrowband ISDN, it may be expensive to provide ATM service locally to all customer premise sites. However, competition between access carriers and Regional Bell Operating Companies (RBOCs) may help reduce local ATM charges. ATM may be key in future systems, such as the National Information Infrastructure (NII) and FTS-2000. Users should position themselves for the transition to more advanced networking technologies when they become available, but should delay implementing ATM on a widespread basis until issues described above involving ATM are resolved.

6.0 User Acquisition Considerations

To recapitulate, there are several major determinations that users must make when acquiring ATM capability. The user should determine the desired quality of service of the networking environment in which ATM is used. The higher the level of service needed, the higher the cost is likely to be, but in general the more capabilities will be provided. For example, the AAL could provide protection against lost cells, or provide an end-to-end transport capability if needed. Users should specify AAL functionality in accordance with their networking quality of service requirements. Users need to determine when to include ATM in their networking plans, since there is a cost in acquiring ATM capability. For example, users need to be able to use productively the high bandwidth that ATM provides, to ensure proper resource utilization. Users need to be assured that performance requirements, "economy of scale" requirements, applications requirements, compatibility (with

existing system) requirements, and communications (interoperability) requirements justify the use of ATM. The simplicity, ease of installation, and availability of ATM products, as well as the coexistence of ATM with legacy systems (for example, legacy LANs), are all important issues. Narrowband ISDN, fast Ethernet, switched Ethernet, SMDS, or frame relay may be useful in migration to ATM use. Users should be aware of ATM switch interoperability issues. For ATM to succeed, users need to be convinced that use of ATM provides for transfer of data with differing characteristics more efficiently and more economically than current technologies.

Appendix A Broadband ISDN Environment and ATM

Integrated Services Digital Network (ISDN) transmits voice, video, fax, and data in an all-digital environment over a single line. All kinds of equipment can be connected to ISDN, eliminating the need for duplicate lines, to provide an enhanced level of service to the end user. There are several methods by which ISDN technology can be employed, including those of narrowband ISDN and broadband ISDN. Many concepts are similar for both, but the interfaces and reference models used to describe behavior are different; these differences arise mostly from the characteristics of the transmission media used.

Narrowband ISDN uses copper cabling connected to digital voice networks. It takes advantage of existing technology widely available today, but is limiting in terms of bandwidth. The interface consists of a number of 64 Kbps data channels controlled by a separate signalling channel.

As technology has improved, the quality of transmission and switching systems has increased. Within the network, this has had the effect of reducing errors, increasing bit rates and increasing distances supported. One ramification is that it is not necessary to repeat functions throughout the network. If network services can be guaranteed, then the functions are only implemented once at the network boundary. Thus, the internals of the network become much simpler, because the high quality of the physical media make many common network functions unnecessary at each node.

The term "broadband" refers to data rates of greater than 1.544 Mbps ("T1"), and can apply to technologies other than ISDN. In broadband ISDN, error recovery functions may be shifted to the network end points, and fixed-length cells can be used, in keeping with the simplicity of the internals of the network. Like narrowband ISDN, out-of-band signalling is used to control the operation of the channels, but instead of controlling multiple 64-kbps channels, a single physical channel (bit stream) is filled with cells. Narrowband ISDN uses time division multiplexing on the 64 kbps channels; broadband ISDN fills the physical channel with bits, and the timing is dependent on the physical medium used. Otherwise, narrowband ISDN concepts are relevant, in that all different types of devices may be connected to a broadband ISDN, but they all need to "speak the same language" when connected to the broadband ISDN. The interface to the broadband ISDN is standardized in the same fashion as that for the narrowband ISDN. Cells are created, cell formats are defined, and access control is enforced at the broadband ISDN interface. Out-of-band signalling occurs for call control at the interface (just as in the narrowband ISDN case). Broadband ISDN is designed to give users access to multiple networks in an integrated environment. Users are able to request a specified level of service transparently from the network

(or networks) and be assured that the requested level will be satisfied.

Broadband ISDN connects with existing technologies at an interface as defined by ITU-T (CCITT); that is, the formats and conventions of other technologies must be converted to broadband ISDN formats and conventions at that interface. Large information entities must be segmented into fixed-length cells; protocol address, and control information may need to be converted (to ATM signals, for example). The data rate requirements of the broadband ISDN interface must also be supported. Multiple configurations must support the single physical channel of the interface. The amount of control and address information that needs to be converted is dependent upon the level of service requested by the user.

The ATM technology is one of a number of technologies which can support the requirements of broadband ISDN. A telecommunications service using the ATM technology standardized by ITU-T is included in the broadband ISDN Reference Model. The layers referenced in figure 3 fit conveniently into this reference model. The ATM technology could also be used in an environment other than that for broadband ISDN. ATM has been standardized for broadband ISDN use; thus far, broadband ISDN is a popular use for ATM. ATM is often mentioned as using SONET technology, because SONET physical media capabilities are advocated for the lowest functional layer of the broadband ISDN reference model; however, ATM may use other physical media capabilities. The broadband ISDN approach is still evolving; as other technologies or services are defined or implemented in the future to support broadband ISDN, appropriate broadband ISDN reference models for these may be developed as well.

