ATM Technology

Geoffrey David Rothman

Independent Study

February 27, 1998

Introduction

Ever since my Data Communications class, I have wondered what ATM entailed. It had to be something other than just a bank teller machine! I didn't realize the importance of this technology until a professor at the University of Kentucky, Brent Seales, asked me if I had done any research in this area. He recommended this topic to me so this is what prompted my independent study. I have found most of my information from books at the Ball State library, on Internet Relay Chat, and from an email I posted to the comp.dcom.cell-relay Usenet newsgroup.

Abstract

The ATM (Asynchronous Transfer Mode) protocol is a networking technology for communications which uses a 53-byte fixed length cell to carry information such as voice, video, data all on the same network. This protocol can be used with a variety of transmission mediums such as copper wire or fiber optic wire, and allows any speed and even variable rate connections. To achieve its high transfer speed of 100 Mbps to gigabit speeds, an ATM network must have special hardware and software. "First, an ATM network consists of one or more high-speed switches that each connect to host computers and to other ATM switches. Second, ATM uses

optical fibers for connections, including connections from a host computer to an ATM switch" (Comer 36). ATM is designed with scalability in mind; it will be able to handle new services that arise.

ATM networks use ATM switches instead of routers. The high-speed switch does nothing but send the cell on its correct path. The speed of ATM networks is also helped by the fact that it uses fixed length cells. Most of the time-consuming error correcting is done at the end nodes, hence eliminating excess time spent going from switch to switch. All ATM switching uses 53-byte cells that contain fields for the VCI (virtual channel identifier), VPI (virtual path identifier), GFC (generic flow control), PT (payload type), CLP (cell loss priority), and HEC (header error check), and then the 48-byte information field (the payload). These cells pass through three ATM standard layers: the physical, ATM, and ATM Adaptation layer. The physical layer provides the physical means for cells to travel across the network. The ATM layer switches the different cells around the network based on the information in their header, and the ATM Adaptation layer assembles and disassembles the services into a stream of cells.

ATM was developed because of emerging trends in the field of networking. The industry needed a protocol that would adapt to diverse services with sometimes unknown requirements. Examples of emerging services that ATM can support are: High Definition Television, video conferencing, high speed data transfer, videophony, video library, home education, and video on demand. There needs to be support for not only high speed services, but also low speed services. ATM can provide this. This research paper gives you a general below-the-surface look at this new technology. I intend to focus on: what ATM is, the concept of ATM, general uses of ATM, and then a few short case studies of ATM implementation.

What is ATM?

Asynchronous Transfer Mode, better known as ATM, was introduced in 1983 by CNET and AT&T Bell Labs in one of their research publications. In 1984, Alcatel Bell of Antwerp, Belgium, began developing this concept (Kumar 141).

To define ATM, I will break its name into two parts: asynchronous and transfer mode. A transfer mode is a telecommunications technique where information can be transported or switched from one network to another. To best define asynchronous, I'll contrast it with the concept of synchronous transfer. In a synchronous transfer, cells (information) are sent across the network constantly whether there is information contained in the cells or not. In an asynchronous transfer mode, the cells are only sent when there is information besides the header contained inside. Cells and header information will be discussed later. ATM technology is based upon high speed (greater than 100 Megabits per second) connection-oriented networking, and is the most recent of telecommunications switching techniques (Clark 1). It is a single communication technology that operates in all network environments and over both short distances (Local Area Networks) and long distances (Metropolitan Area Networks & Wide Area Networks) (Kumar 142).

ATM is the means for carrying information not limited only to data, but also for telephone, video, and other services yet unknown. It has the potential for superseding the many conflicting electronic communication technologies into an integrated whole (Martin 3). One of the strengths of ATM is its scalability in the sense that developers will be able to use ATM technology to transport any new service they create in the future because ATM is "not associated with a single media or transmission speed" (Davidson 87). It guarantees successful transport of any services regardless of the differences in bit rate, quality of service, or the bursty nature of its traffic (Kumar 142). It can support constant bit rate (CBR) services such as voice an video as well as variable bit rate (VBR) services such as the bursty data coming from Local Area Networks (Davidson 87).

Since ATM is so flexible, it will be the basis of Broadband Integrated Services Digital Network (B-ISDN), a telecommunications network that offers simultaneous switching of different information services for varied multimedia applications. Many people speak of ATM and B-ISDN interchangeably because ATM is a subset of B-ISDN. ATM is the protocol or switching technique used by B-ISDN, and B-ISDN is the "complete network and management control architecture" (Clark 2).

The Concept of ATM

In order to connect two hosts or the CEQ (customer equipment), the hosts must have some sort of connection between them. ATM provides for two forms of connection: a SVC (switched virtual circuit) or a PVC (permanent virtual circuit). A good example of a SVC is the abstraction of a telephone call. First, the host "calls" the local ATM switch and gives it the address of the remote computer and the QoS (quality of service) that it needs. The ATM switch communicates by sending signals to other ATM switches across the network. Eventually, the remote computer is reached and must accept the call in order to establish a virtual circuit. More specifically, the first ATM switch looks at the address and quality of service for the request, and if it agrees to forward the data to another ATM switch, the switch then records information about the circuit and then forwards the request to the next switch. This happens until the remote computer is reached, and then the local ATM switch of the host tells both ends that the virtual circuit is complete (Comer 306). With a SVC, once a call is setup, you pay for the entire duration of the call until you disconnect or quit. If there is no sound or nothing transmitted, then you have wasted this resource and money as well. There is another alternative that lets you pay only for what you use. This is called a permanent virtual circuit.

A PVC is similar to a private line that has the ability to stay connected indefinitely. For the PVC, the following procedures are performed. The host calls the service provider with a request for a permanent connection, and then the user gives the destination address, QoS (average bandwidth requirements or bit rate), and the duration of the PVC. The service provider simultaneously enters the information into their system in order to setup the circuit path immediately. The host (user) pays a monthly fee for the circuit's use. You can think of this as a monthly telephone bill. Money and resources are not wasted because you only pay for what you use of the circuit (Kumar 150). Many carriers such as AT&T, MCI, Sprint, Wiltel, and MFS now offer ATM PVC services (Kumar 157).

The means by which information is transferred through these circuits is through 53-byte fixed length cells. The header contains five bytes and the information field (payload) contains 48 bytes. A good analogy for the reasoning behind fixed length cells is given by Balaji Kumar. He gives the example of railway trains. All coaches are the same size, no matter if it carries passengers or cargo. One would think that the passenger coach would be a different length than the cargo coach. It's easier for the train company to build a fixed-length coach so it can add or take away coaches in order to put on different trains when they arrive at the different junctions. This is the same situation as with ATM cell sizes. The routing, adding, discarding, and multiplexing (mixing on one path) of cells can be done faster because you don't have to worry about what is carried inside the cells (142).

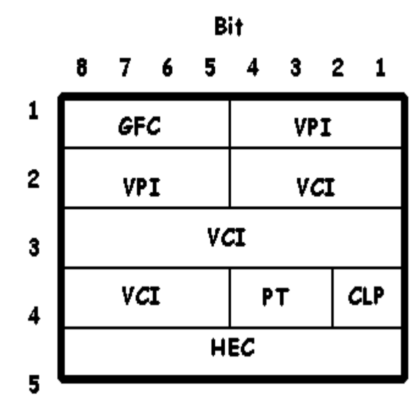

Fig. C

As you can see from Fig. C, there are six different fields in the header. The GFC (generic flow control) field is a four-bit field that is used for flow control from one host to the other. The VPI (virtual path identifier) is an eight bit field that contains all the VCI's (virtual channel identifiers). The VCI field is 16 bits and gives the logical addresses for the connection between two ATM switches or nodes. The PT (payload type) is a three bit field that tells if the information field contains user or network management information. The one-bit CLP (cell loss priority) field defines whether a cell can be eliminated if network congestion is heavy. The last field in the header is the HEC (header error control) field. It's an "eight-bit error code to correct single-bit errors in the header, and to detect double-bit errors" (Davidson 89).

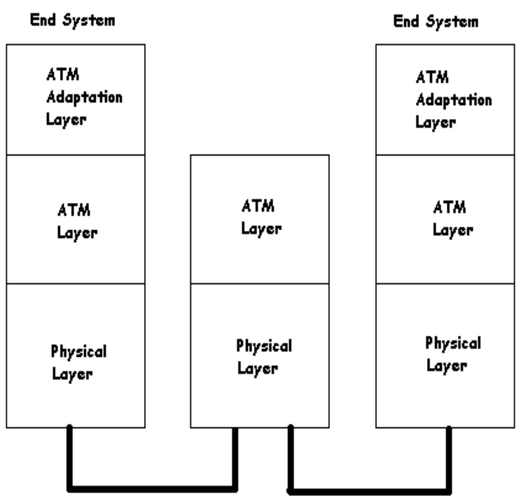

Fig. D

In the ATM protocol, there are three layers by which cells are transported from end user to end user. These layers are the Physical Layer, ATM Layer, and the ATM Adaptation Layer (AAL). Fig. D shows you how the cell is transported from one end system, through its AAL, to its ATM Layer, to the Physical Layer of that system, and then to the Physical Layer of the

ATM switch. Then the cell goes to the ATM Layer of the switch back to its Physical Layer, and then to the Physical Layer of the remote end system. It travels up through the ATM Layer and then the AAL of the remote system until it reaches its destination.

The Physical Layer consists of the PM (physical medium) and the TC (transmission convergence) sublayers. The PM sublayer is responsible for accurate transmission of bits upon the appropriate physical medium (copper, fiber, etc). This sublayer also has to "guarantee a proper bit timing reconstruction at the receiver" (Rad). The transmitting peer must insert the required bit timing information and line coding for everything to run smoothly. In the TC sublayer, the bits are already recognized. The TC sublayer adapts to the transmission system used, generates the HEC in each cell at each transmitter and verifies this field at the receiver. It also performs cell delineation, which is based on the HEC. If many consecutive cells' HEC is correct, then it's assumed that the correct cell boundary is found. To avoid accidental deleting of user information (payload) in the cell, the information field is scrambled at the sending end and descrambled at the receiving end. This ensures that the TC will not find a HEC in the payload field. After the cell delineation has been found, there is a mechanism that use the HEC field to detect errors in the header. If there is only a single bit error, it is corrected (Rad 8).

The ATM layer is where all the work is done with moving the cells through the network. The ATM layer primarily uses endpoint devices and ATM switches. A UNI (user-network interface) connects the CEQ (customer equipment) or endpoint to the ATM switch. The NNI (network-network interface) is the connection between two ATM switches. An example of this setup is seen in Fig. E. Cells are sent by ATM switches along different VP (virtual paths) along a physical medium. A cell may go through many virtual paths until it reaches the other endpoint device. There are many VC (virtual channels) in a virtual path. The reason for having different channels is a strength for ATM. One can send multimedia that uses voice, video, and data along the same virtual path. Each service would have its own virtual channel. Appendix A shows an example of a physical link that contains virtual paths and virtual channels. When an end user wants to send video across the network, it must send a QoS (quality of

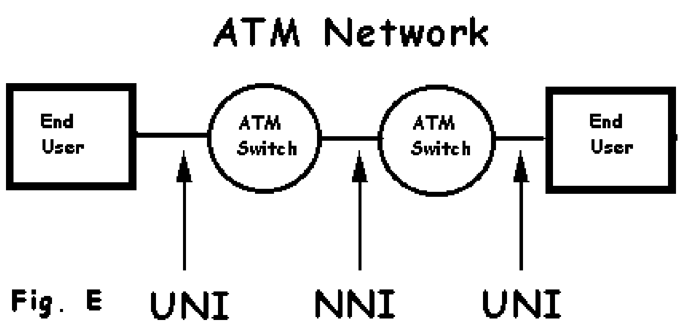

Fig. E

service) information to the receiver. This will determine if the network has the necessary bandwidth to make this type of connection for the service. ATM also uses a feature called statistical multiplexing along the transmission line. It can take information (cells) from different sources and put them all on the same line. The cells are then sorted into the proper virtual paths and channels when they hit an ATM switch. An example of virtual path and virtual channel use is shown in Appendix A and Appendix B.

The ATM Adaptation Layer (AAL) is the highest layer in the ATM architecture and is an interface between user network software and the ATM network. There are four different classes of service for the AAL and five different AAL Layers as shown in Appendix C. The adaptation layers have to distinguish between data, voice, video, and their different transmission requirements. The AALs have two logical sublayers called the CS (convergence sublayer) and the SAR (segmentation and reassembly) sublayer. "The CS receives data from the ATM user interface and gives the units back to the user after receiving them from the SAR. It makes sure the different types of services receive the right level of service at the UNI and NNI. The SAR divides each data unit into cells on the segmentation side and reconstructs incoming cells into data units on the reassembly side" (Davidson 91).

General Uses of ATM

ATM's high speed networks have opened up the doors for many corporations and enterprises to become more efficient. In most cases, ATM was used to solve existing problems and needs while enabling the company for future expansion and to later enable multimedia (Gadecki 100). In the next few paragraphs, I will look at various situations in which ATM can be used to solve problems.

ATM can be used to "uncork" bottlenecks in the network because of its higher speeds. According to Cathy Gadecki and Christine Heckart, the most common implementation of ATM is to relieve the bandwidth constraints in a campus LAN (local area network) backbone. An ATM backbone provides high speed connectivity for file servers that may get used frequently; users can run programs or access files without any delay. They state that the second most common implementation is using an ATM WAN to support the existing router infrastructure. Having a DS1 (T1) connection alone is not enough to support the traffic during peak demand. Increasing the bandwidth from T1 to T3, regardless of the transmission technology is very costly. This extra expense can be offset by the ability to eliminate backbone routers or to consolidate separate networks such as voice and data. The overall cost of your network will decrease in the long run. As a bonus, you will already have the capability to expand effortlessly when needed. This extra bandwidth and speed will make the difference in distance between a LAN and a WAN transparent to applications (Martin 18). One will think that the programs are run locally because ATM connections are 34 Mbps, 45 Mbps & 155 Mbps and are faster than typical 10 Mbps Ethernet and 16 Mbps token ring networks.

Another general use of ATM is for the consolidation of networks. ATM is the only technology capable of this. One economic factor that will help push corporations to ATM adoption is the fact that they can upgrade from a T1 to a T3 and use ATM technology. In doing this, they can get rid of their existing voice services and use their new ATM service for voice as well. Corporations are currently limited to calls within their own ATM LAN until public carriers start providing connectivity between the public switched telephone and ATM network infrastructure (Gadecki 101).

Another sector where ATM technology can be used is with Internet Service Providers. Fig. A is the old network structure of an ISP and Fig. B is a diagram of the new ISP network. In the old ISP diagram, a packet may have to go through three or more routers which turns into more delays and disappointed end-users. "A network that only has routers becomes very difficult to manage and optimize as it grows in size" (Gadecki 103). With the new ISP diagram in Fig. B, the virtual connections interconnect each router with all the routers on one physical path. This new setup: uses bandwidth more efficiently, eliminates many backbone routers, lessens delay because packets go through a maximum of two routers (switches), and handles physical link failures within seconds by having the ATM switch move the virtual circuit to an alternate physical path. "Internet providers will find ATM backbones for their router networks to be the highest performing, most cost-effective architecture" (Gadecki 104). It is expected that the

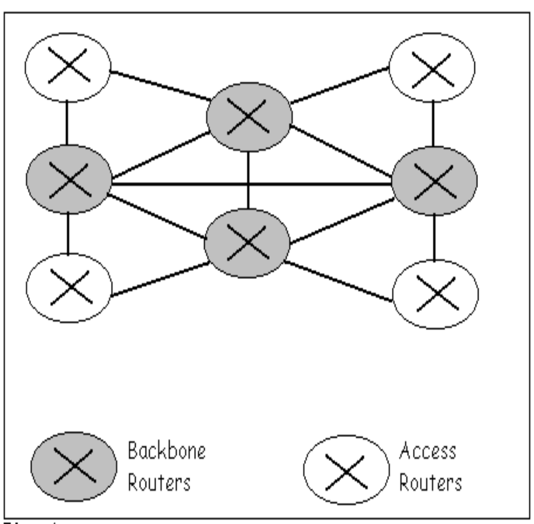

Fig. A

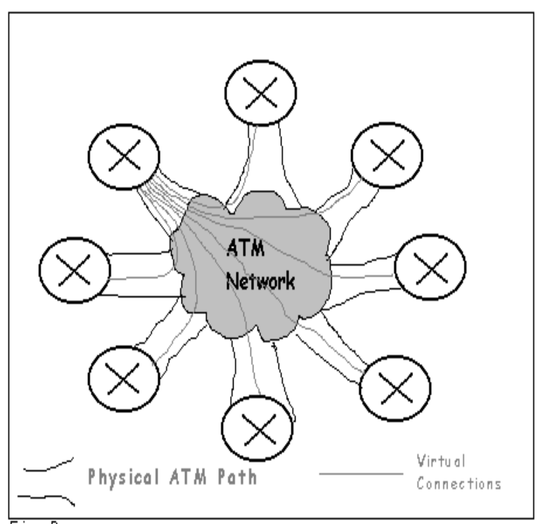

Fig. B

earliest apps coded strictly for ATM technology will be used for workgroups.

It will enable users in remote geographical areas to work together on time-critical projects. The initial focus may be on videoconferencing but could later include scientific experiments by remote staff, or even remote surgery. Entertainment companies could use this to edit movies or film clips remotely. Stock brokerages could use this technology to have updates to the stock market down to the very second. The possibilities are unlimited.

Case Studies in ATM

The first case study was that of William Flusek of McLeodUSA Network Services in Cedar Rapids, IA. His previous network architecture before switching over to ATM was a combination of Ethernet and token ring LANs. His current hardware includes Cisco Lightstream 1010s and Catalyst 5000s. His intranet currently has 5 LS1010s and 27 Catalyst 5000s, and his Internet network has 2 LS1010s and 2 Catalyst 5000s. His chief reason for switching to ATM protocol was the fact that he wanted a way to have multiple networks extended over a single connection, and the option to do voice and video services in the future. He achieved this via ELANS running LANE (LAN emulation) software. They had experimented with Frame Relay and some SONET Ethernet extenders, but ATM proved their best option. They also supply some networking to customers via the same ATM switching cloud at this time. They currently use their ATM architecture for data, but have been looking at a couple of different ways of adding video to the network. On the Internet side, they specifically went with ATM as a way of getting features in future access servers. They are within reason achieving the results they expected. They had a number of challenges with implementing the CISCO LANE implementation and have learned a great deal about the right ways to deploy LANE. William definitely sees a future for ATM technology. "It currently is the only way to get a real QoS (quality of service) and I have doubts about the ability of frame based services to provide this near term, or as effectively as ATM. Certainly, ATM is not as efficient for data as frame based services, but it provides an efficient way to do multiple service types.

The second and more brief case study was with Adam Griffin, the Network Manager for MicroComputer Support Services at Anderson University in Anderson, IN. Anderson University has one OC3 circuit running at 310 Mbps in full duplex between the basement of Decker and Dunn Hall. That is the only ATM link on campus and the only link at all between sites on the north side of third street and the main part of

Anderson's campus. They chose ATM technology because it was "cheap and simple". It is a NNI link on both sides of the connection.

The third case study was with Dr. John Hearns of Guys and St. Thomas' Hospitals in London, England. His previous network architecture was a shared Ethernet coax cable (10base2) using bridges and routers. For his current ATM setup, he uses five FORE systems ATM switches, all with 155 Mbps ports. It has both multimode fiber and UTP (untwisted pair) plus seven Ethernet switches each with an ATM uplink at 155 Mbps and inter-site ATM link at 155 Mbps. The chief reason for the switch to ATM was because it was the "fastest data transfer speed available at the time and also scalable". The ATM technology is mostly used for data, but some departments broadcast quality video conferencing to operating rooms. They are achieving very good results. It is a highly stable and reliable network. However, the raw network bandwidth wasn't what was expected, as they are using LANE. Dr. Hearns definitely sees a future for ATM technology. "ATM to the desktop will be reserved for high end/power applications which need video streams. Most desktops will have Ethernet, through switches, then on to ATM backbones".

Works Cited

Clark, Martin P. *ATM Networks Principles and Use*. New York: Wiley-Teubner, 1996.

Comer, Douglas E. *Internetworking with TCP/IP*. New Jersey: Prentice Hall, 1995.

Davidson, Robert P. *Broadband Networking ABCs for Managers: ATM, B-ISDN, Cell/Frame Relay to SONET*. New York: John Wiley & Sons, Inc., 1994.

Flusek, William. Email interview, 1998.

Gadecki, Cathy, and Christine Heckart. *ATM for Dummies*. Chicago: IDG Books Worldwide, 1997.

Griffin, Adam. Email Interview, 1997.

Hearns, John. Email interview, 1998.

Jain, Raj. An ATM Overview. 1997. http://www.cis.ohio-state.edu/~jain.

Kumar, Balaji. *Broadband Communications: A Professional's Guide to ATM, Frame Relay, SMDS, Sonet, and B-ISDN*. New York: McGraw-Hill, Inc., 1995.

Martin, James. *Asynchronous Transfer Mode: ATM Architecture and Implementation*. New Jersey: Prentice Hall, 1997.

RAD Data Communications. "A Little Bit of History .. or Why Do We Need ATM?."1994. http://www.rad.com/networks/1994/atm/tutorial.htm (1994).