ATM vs Gigabit Ethernet

(Research Essay #2)

Seok Bum Ko

kosby@ele.uri.edu

ELE548 Computer Architecture

5 March, 1999

1. Introduction

In the past several years, Local Area Networks have evolved from simple file and print sharing to applications that include large files, multimedia, and Internet access. Furthermore, the numbers of users on backbone LANs are increasing due to the lowering cost of the computer. As the volume of LAN traffic increases, typical 10 Mbps shared Ethernet backbones are becoming insufficient to handle the traffic. This has led to the deployment of faster technologies such as Fast Ethernet and FDDI (Fiber Distributed Data Interface) in the backbone. This allows for 10 Mbps desktop connections and 100 Mbps in the backbone. A combination of lowering costs of Fast Ethernet adapters, and faster desktop machines has led to the deployment of Fast Ethernet adapters in desktop machines. The two dominant technologies competing for dominance in the network backbone are ATM (Asynchronous Transfer Mode) and Gigabit Ethernet. This essay will explain these two technologies as well as give a technical and economic comparison of the two technologies.

2. Overview of ATM

2.1 ATM Defined
ATM (Asynchronous Transmission Mode) is a connection oriented cell-based switching technology using 53-byte cells to transport information. ATM is a relatively new approach to handling information transfer within a single network or between networks that may span the globe. The ATM Forum (the ATM standards body) completed implementation agreements in mid-1993 on the cell format for transmitting user data. The ATM Forum is continuing to develop and approve standards that will allow ATM to be the bridge between different types of legacy networks, while still maintaining the ability to offer a guaranteed quality of service. ATM products have been shipping for the past three to four years.
ATM does not actually transmit cells asynchronously as the name may suggest. Fixed-length ATM cells are transmitted continuously and synchronously, with no break between cells. When no user information is being transmitted each ATM cell is filled with a specific bit pattern indicating that it is an empty or idle cell. The asynchronous nature of ATM comes from the indeterminate time when the next information unit of a logical connection may start. Time not used by one logical connection may be given to other connections or filled with idle cells. This means that the cells for any given connection arrive asynchronously. Each cell is routed to its proper destination by addressing within that cell.

2.2 ATM Protocol Stack
The ATM protocol stack can be broken down into three layers: the physical layer, the ATM layer, and the ATM adaptation layer. The bottom layer is the physical layer. This layer consists of the physical transport used to transfer the ATM cells from one node to another. This includes converting the signals into the appropriate electrical or optical format and loading and unloading cells into the appropriate transmission frames. ATM has been adapted to work with the following physical transports: 155 Mbps SONET/STS-3c over multimode fiber and Category 5 UTP, 622 Mbps SONET/STS-12c over multimode fiber and single mode fiber, and 25 Mbps over CAT 3,4, or 5 UTP. ATM will be extended to 2.5 Gbps (OC-48).
The next layer is the ATM layer. This layer provides the switching and routing of the ATM packets according to their VCI (Virtual Channel Identifiers) and VPI (Virtual Path Identifiers) labels. The ATM layer is also responsible for generating the headers for the ATM cells and extracting this header from incoming cells. The top layer is the ATM adaptation layer. This layer maps various types of traffic into and out of the ATM cells. There are different types of adaptation layers for different traffic types. This is necessary due to the different transmission characteristics of the specific traffic. Currently there are five types of adaptation layers: AAL1 through AAL5. AAL1 is used to support real time, constant bit rate traffic like voice and video

traffic. AAL2 is used to support real time, variable bit rate traffic like MPEG video traffic. AAL3/4 is used to provide support for non-real time data, and was originally intended to carry LAN traffic. It has since been replaced by AAL5, since AAL5 has lower overhead per cell. By dropping some of the features of AAL3/4, AAL5 provides 48 payload bytes per cell.

2. 3 ATM Capabilities

ATM's main selling point has been the fact that ATM, unlike most other protocols, can carry voice, video, data, imaging, and graphics either separately or simultaneously on the same link. This is due to ATM's small, fixed-length cell as well as its QoS (Quality of Service) parameters.

2. 4 The ATM Cell - Key to Operation

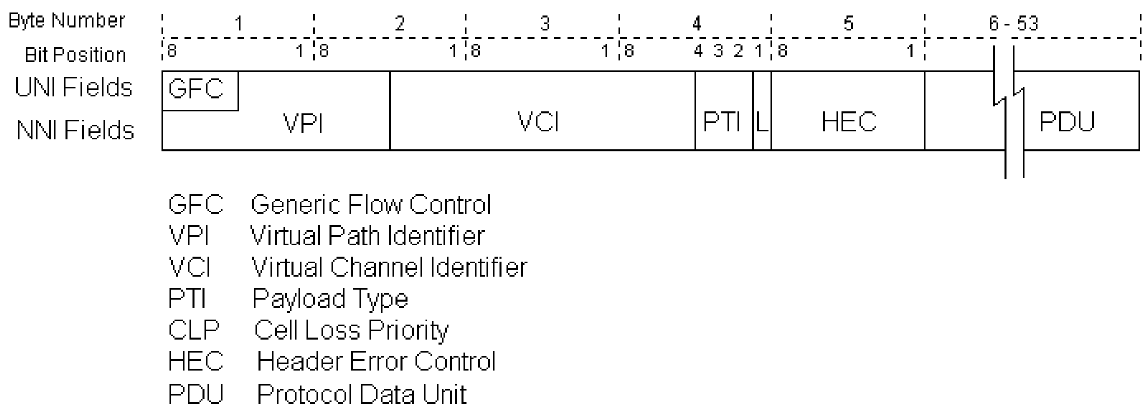

Figure 1. The ATM Cell

The first 5 bytes are the header, containing the addressing and control information, while the last 48 bytes, called the payload, carry the revenue-bearing information. The header was made as small as possible for efficiency. Similar thinking would dictate a long payload. This was not to be the case, though, for ATM. Designers had to be concerned with the impact the payload length would have on all types of carried information. The design team was concerned not only with efficiency, but also with packetization delay. Packetization delay is a result of the time it takes to fill a cell with 64 Kbps digitized voice samples. The 53-byte cell size was a compromise between efficiency and delay.

2.4.1 The ATM Cell Header

The UNI (User to Network Interface) cell header provides addressing information, limited flow control, and header error checking. The first four bits are called the GFC (Generic Flow Control). These bits would be used for local flow control for multiple users on the customer side of a switch sharing an access line. This function is generally not used at this time and these four bits are set to all zeros. For an NNI (Network to Network Interface) these bits are overwritten with addressing information. The UNI header has 8 bits for the VPI (Virtual Path Identifier), while the NNI header has 12 bits. This gives 255 possible paths for the UNI and 4095 possible paths for the NNI on the same port. The VCI (Virtual Channel Identifier), further extends the addressing with 16 more bits or 65536 possible connections within each path address. Some addresses are reserved for specific functions such as signaling, only slightly reducing the available addresses. The next three bits are the PTI (Payload Type Identifier). The PTI is used to distinguish user data cells from OA&M (Operations, Administration, and Maintenance) commands and statistics. If network congestion is experienced, the PTI is modified as it passes from switch to switch. The network can then relieve congestion by discarding any cells that are in excess of the guaranteed connection rate.

The CLP (Cell Loss Priority) bit is a two-state priority indicator telling the network which cells to discard first in the event network congestion. All CLP bits are initially set to zero. The CLP bit may be set to one by the first node in the backbone if usage exceeds that set upon initial connection. Any cell with a CLP bit set to one will be discarded within the network before any cell with a CLP bit set to zero.
The 8-bit HEC (Header Error Control) has the capability to correct single-bit header errors to ensure proper addressing. Headers that contain multiple-bit errors will simply be discarded. It is important to note that the HEC does not check for errors in the information payload. This function would be the responsibility of a higher level protocol, typically in the transport layer.

3. Overview of Gigabit Ethernet

3.1 Gigabit Ethernet Defined
Gigabit Ethernet is an extension to the 10 and 100 Mbps IEEE 802.3 standards. Like its predecessor, Gigabit Ethernet is a connectionless protocol. This means that each Ethernet frame contains a destination address, and is routed through the network. This differs from ATM, where a connection is established prior to data transfer. Gigabit Ethernet will provide a raw-data bandwidth of 1000 Mbps, and maintain full compatibility with the previous Ethernet standards. It will provide support for both half and full-duplex operation. In the case of half-duplex operation, carrier-sense with collision-detection (CSMA/CD) will be used, as in the previous Ethernet versions. It will operate over Category 5 UTP, MultiMode Fiber (MMF), and Single Mode Fiber (SMF).

3.2 Gigabit Ethernet Protocol Stack
The Gigabit Ethernet protocol stack can be broken down into two layers: the physical layer and the data link layer. The bottom layer is the physical layer. The physical layer of Gigabit Ethernet is built upon the Fiber Channel standard. Fiber Channel is a standard defined by the American National Standards Institute (ANSI) in the ANSI X3T11 specification. The Fiber Channel standard allows for a wide range of speeds, currently from 12.5 to 4000 megabytes per second over a variety of media including twisted-pair and fiber optic cable. Gigabit Ethernet uses the lower two layers (FC-0 and FC-1) of the Fiber Channel protocol stack. FC-0 defines the connector, fiber, and electrical parameters needed to support gigabit speeds. FC-1 defines the 8B/10B encoding/decoding scheme. The two additional bytes of overhead are used to integrate the data with the clock information required by serial transmission techniques. Since, Fiber Channel uses 10 bits to represent each 8 bits of data, it must operate at a speed sufficient to accommodate this 25 percent overhead. This means that Gigabit Ethernet will have to run at wire speeds of 1.25 Gbps to attain a true 1 Gbps. The fact that these physical layer standards are predefined will allow Gigabit Ethernet to be standardized and be deployed quickly.
The next layer is the data link layer. The data link layer consists of the Media Access Control (MAC) layer and the Logical Link Control (LLC) layer. The LLC sits on top of the MAC layer, and it serves as a common interface to the network layer. The MAC layer defines the Ethernet frame format, and the mechanism to gain access to the physical link. Gigabit Ethernet is designed to operate in full-duplex and half-duplex modes. In full-duplex mode Gigabit Ethernet uses IEEE 802.3x, which defines the MAC layer specifications for full-duplex operation. In half-duplex mode, Gigabit Ethernet uses the standard IEEE 802.3 CSMA/CD (carrier sense multiple access with collision detection) specification.
In full-duplex mode (Switched Ethernet), packets travel in both directions over the same path for a total bandwidth of 2 Gbps. In full-duplex mode, all connections are point-to-point, and as result multiple sessions can exist on the LAN without collisions. Full-duplex connections can exist between ports on two switches, a switch and a workstation, or between two workstations.

In the case of shared-port connections, Gigabit Ethernet operates in half-duplex mode. This situation may arise when several workstations are connected to a single gigabit port through some sort of shared repeater device. Since half-duplex operation uses CSMA / CD to operate, collisions can occur on the backbone. This results in performance degradation. Gigabit Ethernet is most effective when running in full-duplex mode.

3.3 Distance Limitations
A wire speed of 1.25 Gbps brings up the issue of the distance limitations of Gigabit Ethernet. Gigabit Ethernet will be designed to operate over Category 5 cable, multimode fiber, and single mode fiber. Using low cost short-wavelength (780 nm) transceivers limits Gigabit Ethernet to 300 meters over multimode fiber. If higher costs, long wavelength (1300nm) transceivers are used, the distance limitation can be extended to 550 meters. Furthermore, distances up to 2 kilometers can be attained over single mode fiber. These distances should be sufficient for backbone connections of Gigabit Ethernet. A problem arises, however, when connecting servers and workstations to the LAN. These connections have primarily been category five cable in the past. Current specifications limit Gigabit Ethernet to 25 meters over category five cable. This may be a problem for the current infrastructure wiring in some areas. There are current efforts in place to extend this maximum distance to 100 meters, to make it more feasible to run gigabit Ethernet to the desktop.
In half-duplex operation Gigabit Ethernet will be designed to have network diameters of up to 200 meters over twisted pair, and a maximum of one repeater. The minimum Ethernet frame size is 512 bits. At 10 Mbps this corresponds to a slot time of 51.2us. The slot time is the time it takes for the station to transmit a 512 bit frame. It is important that the slot time be large enough to allow for a station at the farthest end of the network to detect collisions. Ethernet LANs can have a maximum network diameter of 2500 meters. At 100 Mbps (Fast Ethernet) the slot time drops to 5.12us. This lowers the network diameter to 250 meters for Fast Ethernet LANs. Therefore, for Gigabit Ethernet the slot time drops to .512us, and the network diameter drops to 25 meters. In order to support the projected 200 meter diameter, the slot time has been extended. When sending small frames, the Gigabit Ethernet sending station adds extra time between frame transmissions to ensure collision detection on the LAN. The frame size itself is not changed, only the time between frame transmissions is changed. This lowers the performance of Gigabit Ethernet, but is required for CSMA / CD to operate at 1Gbps.

4. Technical Comparisons

4.1 Overhead
In order to carry traffic from higher level protocols, both ATM and Gigabit Ethernet must encapsulate that traffic. When carrying typical LAN traffic, Gigabit Ethernet has lower overhead per higher level packet than ATM. For instance, in the case of a 1500-byte IP datagram, Gigabit Ethernet would add 26 bytes of overhead, resulting in 1526 bytes to transmit the IP datagram. ATM AAL5 adds an 8-byte trailer and a variable pad size to ensure that the AAL5 protocol data unit (PDU) is a multiple of 48 bytes. For a 1500-byte IP datagram, this results in an AAL5 PDU equal to 1536 bytes. The AAL5 then segments the AAL5 PDU into 48 byte segments to be carried in 53 byte ATM cells (each ATM cell has a five-byte header). Therefore, 32 ATM cells (1696 bytes) are required to transmit the 1500 byte IP datagram. The corresponding efficiencies when compared to the IP network layer are about 98% for Gigabit Ethernet and 88% for ATM.

4.2 Support for Existing Network Protocols
Since Gigabit uses the same IEEE 802.2 LLC layer as standard Ethernet, existing network protocols like IP and IPX can operate without modification over Gigabit Ethernet. In order for existing protocols to work over ATM, they must be adapted operate directly over the ATM

adaptation layer. ATM offers to solutions to this problem: Classical IP over ATM and LAN Emulation.

A current standard known as Classical IP over ATM maps IP network layer addresses to ATM addresses, and enables ATM-attached devices to send IP packets over an ATM network. While this solution has the obvious advantage that it allows IP traffic to be routed on an ATM LAN, it has several disadvantages. First, Classical IP only supports the IP network protocol. Secondly, unlike IP over Gigabit Ethernet, Classical IP over ATM currently has no support for multicast traffic. This is because unlike Ethernet, ATM is a connection-oriented technology. Finally, and most importantly, Classical IP over ATM does not solve delay or congestion problems, because Classical IP over ATM cannot take advantage of ATM's Quality of Service (QoS).

LAN emulation takes a different approach than Classical IP. Rather than adapting the network level protocol to ATM, LAN Emulation emulates the MAC layer above ATM AAL5. In the case of Ethernet LAN emulation, this LAN emulation interface looks like a standard Ethernet interface. The advantage of this method is that it allows all existing network layer protocols that supported Ethernet to operate over ATM. Furthermore, unlike Classical IP, LAN emulation supports multicast traffic. This is implemented by the LAN Emulation Broadcast / Unknown Server (BUS) which accepts broadcast / multicast traffic and transmits to all the ATM stations through a point-to-multipoint connection. The main disadvantage of LAN emulation is that since it is transparent to higher level protocols, it cannot take advantage of ATM's QoS.

When upgrading existing Ethernet networks, Gigabit Ethernet offers a simpler solution. Gigabit Ethernet will operate with existing network protocols without the setup times associated with Classical IP over ATM and LAN Emulation.

4.3 Maturity

One disadvantage of Gigabit Ethernet is that the standards are not complete. In order to implement Gigabit Ethernet today, pre-standard products must be used. This may cause a problem if products from different manufactures are used on the same LAN. ATM standards for data traffic are complete. This means that ATM products from different manufactures that adhere to the standards will operate with each other.

4.4 Quality of Service and Scalability

When using native ATM applications, ATM can offer a QoS guarantee. This is something that is currently not offered in Gigabit Ethernet. It is clear that ATM offers some real advantages over gigabit Ethernet in functionality, but functionality is not the only consideration a user must examine when trying to determine which technology is best suited to their needs. Cost and ease of migration are also important considerations.

Table1. Gigabit Ethernet and ATM

Feature	Gigabit Ethernet	ATM
Overhead	Lower	Higher
Connection to existing LAN	No conversion needed, simple to integrate	Need frame-to-cell conversion, LAN emulation; complex to integrate
Standards-based position	Pre-standard products available	Standard-based products already available
Quality of Service	Non-connection oriented	Connection-oriented
Typical topology	LAN, workgroup, backbone	WAN, backbone, campus
Cost	<$2000 (US)/port	>$2000(US)/port
Typical Speed	1 Gigabit (1000Mbps)	155Mbps, 622Mbps
Technology	Frame-based	Packet switching
Advantages	Less expensive, easier to use	Reliability and flexibility
Disadvantages	Less scalable and reliable	Requires retraining staff, relatively expensive

5. Conclusions

In reality, ATM and Gigabit Ethernet are not equal substitutes for each other and should not be considered as such. Each technology is appropriate for specific applications. Simply stated, Gigabit Ethernet will be deployed in areas where fast Ethernet, and other technologies like FDDI, are no longer able to provide the bandwidth needed for pure data traffic. In other words, Gigabit Ethernet will be used in areas where high data throughput is required, but quality of service is not a main concern. ATM will be used in environments where video, voice, and other delay sensitive traffic exist. Some applications for ATM may be video conferencing and video on demand. In addition to these LAN applications, ATM is also starting to be implemented in the WAN. This could lead to networks that use ATM in the LAN and WAN, or possibly networks that use Gigabit Ethernet to carry LAN traffic and ATM to carry WAN traffic.

References

1. “Are ATM, Gigabit ethernet ready to for prime time?”, IEEE computer, May, 1998
2. "Why Gigabit Ethernet?" Austin, Texas: Jato Technologies, January 1997.
3. Flanagan, William A. ATM User's Guide (1st ed). New York: Flatiron Publishing, 1994.
4. Freeman, Roger L. Telecommunication System Engineering (3rd ed.). New York, New York: John Wiley & Sons, Inc, 1996.
5. Axner, David. "Gigabit Ethernet: A Technical Assessment". Telecommunications March 1997: pp. 31-34.
6. "Gigabit Ethernet White Paper." Cupertino, California: Gigabit Ethernet Alliance, August 1996.
7. http://www.cisco.com/warp/public/729/gigabit/gigbt_tc.htm
8. http://www.hp.com/rnd/techlib/techlib_gigabit.htm