

FASTER ETHERNET AND THE ATM MARKET BOUNDARY

G. Kent Webb, San Jose State University, webb_k@cob.sjsu.edu

ABSTRACT

As a network technology, ethernet flourished in low-cost, low-end markets. Simple to make and with open standards, many companies created products. The resulting improvement in price, performance, and market acceptance resulted in ethernet replacing the more established and sophisticated token-ring technology that dominated early large corporate LANs. As ethernet gets faster, accelerating from the original 10 Mbps into Gigabit speeds, the technology is poised to challenge the dominant backbone and WAN standard, ATM. A discussion of new ethernet developments is formalized with a decision model used to define a market boundary with data illustrating why and where a technology may dominate.

Keywords : Ethernet, ATM, Network Design, Technology Market

INTRODUCTION

Two lines of technical history merge with the advent of Gigabit ethernet: local and wide area networking. The first part of this introduction examines the development from the roots of ethernet, as a local area networking technology. The second part of the introduction looks at the markets challenged by the advances in ethernet having to do with backbones, wide area networks, and the Internet

Local Area Networks (LANs)

The data link layer defines the standard for how data travels on the circuits connecting computers in the network. Equipment based on data link layer standards include: the network interface cards that connect the computers to the network, and the hubs or switches that combine the circuits in the network [7]. While a number of incompatible technologies have been presented to the market over the last 20 years, the major choice for LAN network design has been between ethernet and token-ring.

Ethernet was originally developed beginning in about 1973 by Robert Metcalfe at Xerox PARC in Palo Alto, California, in an effort to drive the demand for high-speed printers that could be used on corporate local area networks [1]. Ethernet was described in the patent application as a "multipoint data communication system with collision detection", the simplifying and key design element of the technology. In the ethernet model, the network interface card that connects an individual computer to the network listens for activity on the network, if none is detected, a transmission is sent. The transmission may successfully reach its destination, but if another network interface card on the network makes the same decision at about the same time, a collision occurs. Collisions in some ethernet networks are about as common as collisions in the demolition derby. In order to deal with this problem, ethernet relies on CSMA/CD, **C**arrier **S**ense **M**ultiple **A**ccess / **C**ollision **D**etection. When a collision is detected, the sending devices

wait for a random amount of time then tries to send again, and again, as many times as necessary to complete the transmission. In simple ethernet networks this results in rapid degradation of response time as traffic on the network builds up.

Robert Metcalfe left Xerox to promote ethernet and convinced DEC, Intel, and Xerox to support the standard. Although Dec and Xerox are no longer active in the market, Intel has been joined by numerous companies including 3Com that have developed and advanced the technology. In 1983, the Institute of Electrical and Electronics Engineers (IEEE) devised IEEE 802.3 the common standard for 10 Mbps ethernet.

IBM developed a more sophisticated technology in response, token-ring. It avoids collisions by the use of a traffic manager, a token that circulates on the network. Computers ready to transmit onto the network must wait until the token arrives; so only one computer can transmit at a time. By the early 1990s, token ring dominated the corporate LAN market because it was easier to manage, more efficient, and network administrators felt they could rely on IBM to support the product into the future.

Under normal network traffic conditions basic ethernet will typically run at only about 50 percent of theoretical capacity, about 5 Mbps for a 10 Mbps ethernet network, because collisions start to occur at such a rate that throughput and response time rapidly degrade. A typical office user will generate about 25 Kbps of traffic, limiting business LAN sizes to about 20 users. Token-ring, with its better traffic management, will run at 80 to 90 percent of theoretical capacity [2] or about 25 to 30 percent faster.

In order to address this problem in the early 1990s, ethernet manufacturers developed a low cost switch that could be used to replace the hub in the LAN environment. While a hub allows traffic to propagate throughout the network, a switch restricts traffic to just a few circuits, greatly improving the effective throughput of the network. Ethernet switches were embraced by the market because of their effectiveness and relatively low cost. This success resulted in higher manufacturing volumes and even lower cost and ease of use. Development of ethernet switching is one of the two key technology trends that will propel ethernet into higher ends of the network market.

The other major technical development important to the acceptance of ethernet has been rapid increases in speed. In 1995 IEEE 802.3u defined 100 Mbps ethernet, the same technology but with faster signaling. The faster speed helped ethernet gain acceptance at the higher end of the LAN market so that now about 85 to 90 percent of LANs are ethernet and virtually all new LANs are ethernet. In 1998 IEEE 802.3z defined Gigabit ethernet [3]. At that speed, ethernet became a serious contender for the technology of the Enterprise Backbone, the middle one-third of Figure 1.

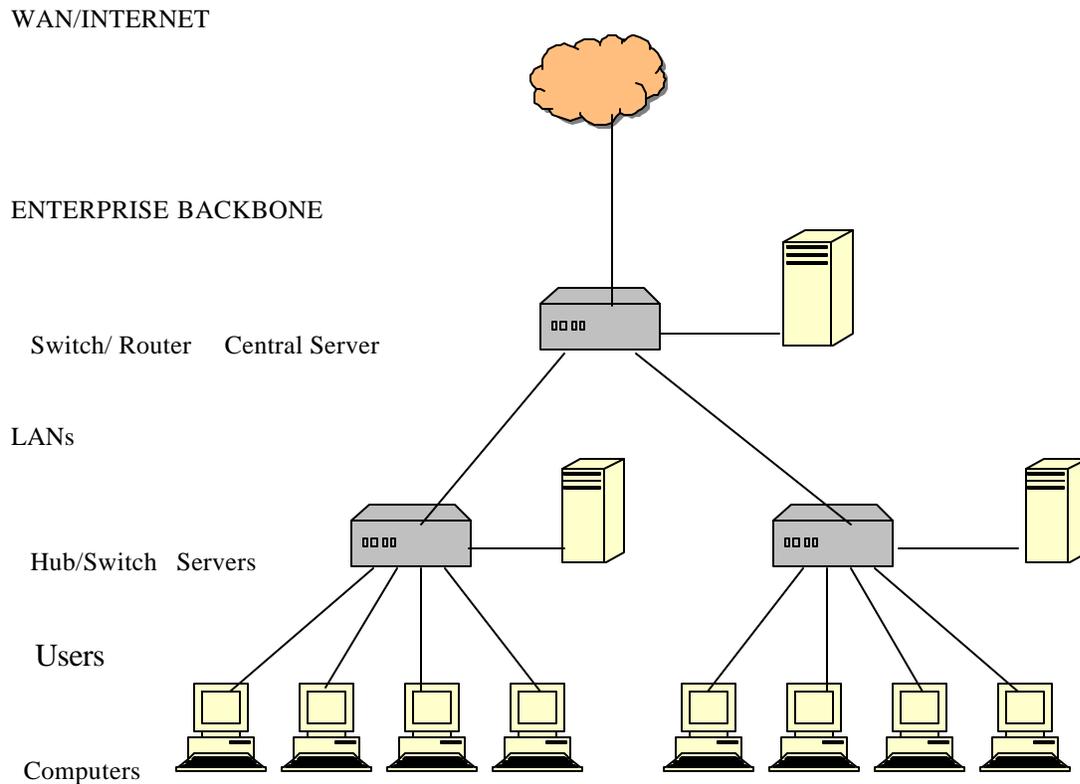
In 2002, the IEEE 10 Gigabit Ethernet Task Force ratified the 802.3ae standard. New products have already been introduced based on this standard [5] that will increase the acceptance of ethernet as a backbone technology and make it competitive in Metropolitan Area Networks (MANs) and for some Wide Area Network (WAN) applications. Only a few years ago, some network analysts expressed skepticism that 10 Gbps ethernet could be implemented, but the

technologists supporting ethernet now have their sights set on 100 Gbps and are looking forward to ethernet at terabit rates. This prospect has some wondering how far up the network, represented by Figure 1, that ethernet might move.

Wide Area Networks (WANs)

WANs emerged from the voice wide area network otherwise known as the telephone network. Early WAN technologies and protocols were built to exploit the existing analog network. X.25, a WAN standard developed at about the same time as ethernet, was designed to compensate for noisy analog telephone lines. Built on X.25 and ISDN standards, frame relay became the dominant private WAN technology during the 1990s [6]. IBM entered the WAN market with a technology that, like token-ring, offered a significant technological improvement over a more public standard. More scalable and better able to handle mixed transmission types, IBM's Asynchronous Transfer Mode (ATM) is the current data-link layer for most of the Internet and is edging frame relay out of the private WAN market.

Figure 1: Representative Network Layout



The Internet also emerged from the voice wide area network. During the 1960s when the Department of Defense decided to tie together a collection of incompatible computer systems it

specified two major goals in its development of the precursor of the Internet, the Advanced Research Projects Agency Network (ARPANET). Michael Hauben provides a detailed summary of the process of creating ARPANET and TCP/IP as a protocol independent of the underlying physical and data link network layers. The first of two major goals defined by the Department of Defense was: "To construct a 'subnetwork' consisting of telephone circuits and switching nodes whose reliability, delay characteristics, capacity, and cost would facilitate resource sharing among computers on the network [4]." ATM, with its ability to set up virtual circuits and behave like a telephone circuit for real-time data and with the backing of IBM as a guarantee of the product's success, became a clear choice for the Internet backbone.

Another goal in the development of ARPANET was the development of a network protocol independent of the underlying physical and data link network layers, resulting in TCP/IP [4]. The backbone of the current Internet relies on a mix of copper and optical cable for the physical layer with ATM (Asynchronous Transfer Mode) at the data link layer. Some see the introduction of low-cost, switched ethernet into these high-end network markets as a logical next move [8].

Ethernet is clearly established in the LAN market, the bottom third of Figure 1. ATM currently dominates top third, WAN and WAN access. The current market boundary between the two technologies lies in the middle third of the figure, the Enterprise Backbone (this could also include Metropolitan Area Networks) but the boundary is not well defined.

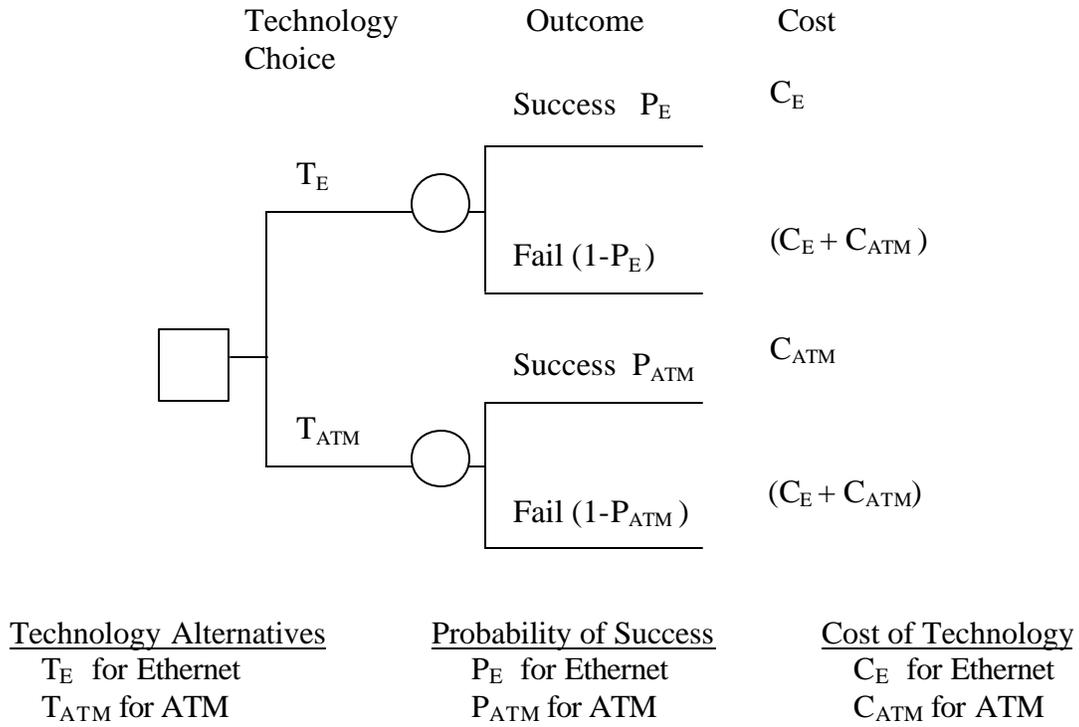
In terms of new systems, there are both new ATM and ethernet backbones under development. For WAN access, some Internet Service Providers have recently introduced Gigabit ethernet connections at lower prices than ATM. Most high-speed access, however, remains ATM. Among new cable television systems that provide high-speed Internet access, some designers have begun to advocate ethernet to replace the ATM standard for their market. At the extreme, some analysts maintain that ethernet will completely replace ATM.

A MODEL DEFINING THE ETHERNET AND ATM MARKET BOUNDARY

This section presents a more formal analysis in the form of a decision model illustrated in Figure 2 that assumes network designers make their decisions by looking at the cost of two roughly similar technologies and the risk associated with choosing a technology that is rejected by the market, leaving them with a network that will have to be rebuilt in the future using accepted technology.

Designers in this example face a choice between the two technologies of ethernet and ATM, T_E and T_{ATM} , with costs C_E and C_{ATM} . Designers who choose a technology that fails in the market will have to rebuild their networks with the other technology. Their total cost will be $(C_E + C_{ATM})$, their initial construction cost and the cost of building a new network with the alternate technology.

Figure 2: Technology Decision Analysis



The probabilities of success for each technology, P_E and P_{ATM} , represent the subjective probability formed by individual network designers. These values do not need to sum to one since sometimes two technologies are accepted by the market. The breakeven point for the decision tree, where an individual network designer would be indifferent between the two technologies, is

$$P_E C_E + (1 - P_E)(C_E + C_{ATM}) = P_{ATM} C_{ATM} + (1 - P_{ATM})(C_E + C_{ATM}) \tag{1}$$

Which can be reduced to: $P_E/C_E = P_{ATM}/C_{ATM}$ (2)

That is, the network designer will be indifferent between the two technologies when the ratio of the product success probabilities to the product costs is the same. Designers will chose one technology over the other when the ratio of the technology cost is less than the ratio of their subjective probability of product success, or:

If $P_E/C_E > P_{ATM}/C_{ATM}$ choose ethernet (3)

If $P_E/C_E < P_{ATM}/C_{ATM}$ choose ATM (4)

A market boundary for the total percentage of new systems built with either technology depends on the costs of each technology and the distribution of the success probabilities. Given a consensus estimate for the success of ATM, a market boundary can be calculated for ethernet as

equation 5. An individual designer with a higher estimate than P_E^* for the success of ethernet will choose ethernet, while any designer having a subjective estimate lower than P_E^* for the success of ethernet will choose ATM.

$$P_E^* = (C_E/C_{ATM})P_{ATM} \tag{5}$$

For example, as the cost of ethernet declines with respect to ATM, designers need a lower estimate of success for ethernet in order to choose it in their network designs. As the probability that ATM will succeed goes down, designers also need a lower estimate of success for ethernet in order to choose it.

Historical Data

Table 1 presents recent historical data on the cost for the two technologies, C_E and C_{ATM} . Taken from industry sources, the costs represent the average price per port for a switch plus the average price of a network interface card. The table also includes a guess about what the average network designer thinks is the probability of success for ATM. The data from Table 1 are used to estimate a probability market share boundary, P_E^* based on Equation 5 and reported in Table 2 below.

Table 1: Historical Cost Data and an Estimated P_{ATM}

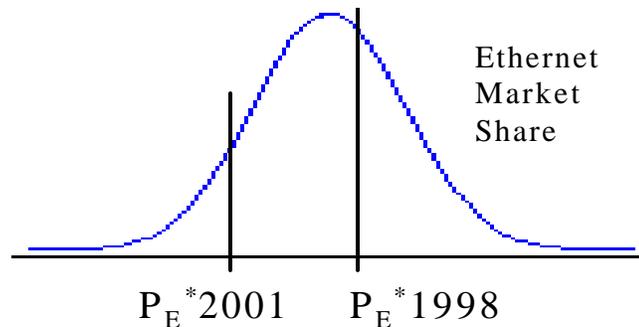
Year	C_E	C_{ATM}	P_{ATM}
1998	1500	2000	.75
1999	1000	1500	.73
2000	700	1100	.71
2001	250	800	.70

Table 2: Estimated Success Probability Boundary for Ethernet, P_E^*

Year	P_E^*
1998	.56
1999	.49
2000	.45
2001	.22

Figure 3 provides an interpretation of the data in Table 2. Assuming a normal distribution for an individual designers subjective probability estimate of ethernet success P_E , designers whose individual P_E is higher than P_E^* will choose ethernet. Other designers will chose ATM. In 1998 about 44 percent of designers were in the area to the right of P_E^* and so chose ethernet for new systems.

Figure 3: Ethernet Market Share



(Designers with a higher estimated probability of success that P_E^* will choose ethernet. The area to the right of P_E^* defines the ethernet market share.)

As ethernet costs moved lower faster than ATM, P_E^* in Figure 2 is shifted to the left resulting in larger market share for ethernet. This pattern approximates actual design behavior.

CONCLUSIONS

The best technology does not always win in the market. Ethernet pushed the more sophisticated token-ring out of the LAN market and is in the process of pushing ATM out of the Enterprise Backbone market because it is cheaper and it can do the job. Ethernet is already being used as a WAN access or WAN technology in some markets. As ATM continues to lose market share in the overall network market, it will be difficult for manufacturers to compete against steadily declining ethernet prices. Since the Internet network protocol, TCP/IP, was specifically designed to support various underlying network standards, ethernet can readily replace ATM throughout the network. The end result may be an ethernet internet.

REFERENCES

- [1] Bellis, Mary. "The Invention of the Ethernet – Local Area Networks"
<http://inventors.about.com/library/weekly/aa111598.htm>
- [2] Caruso, Jeff, "What to do about token ring?" Network World, April, 28, 1999
- [3] Frazier, H., and Johnson, H. "Gigabit Ethernet: From 100 to 1,000 Mbps." IEEE Internet Computing, January/February 1999.
- [4] Hauben, Michael "Behind the Net - The untold history of the ARPANET, Or - The "Open" History of the ARPANET/Internet", <http://www.dei.isep.ipp.pt/docs/arpa.html>
- [5] Phil Hochmuth, "10 Gigabit Ethernet to step into the N+I spotlight". Network World, May 7, 2001.
- [6] Miller, M. A. (1991). Internetworking: A Guide to Network Communications LAN to LAN; LAN to WAN. New York, NY: M&T Books
- [7] Stallings, W. "Local Networks: An Introduction" (Prentice Hall, 1997).
- [8] Mark Sebastyn, Senior Product Marketing Manager, Appian Communications "Optimizing Revenue with Optical Packet Services"
http://www.internetindustry.com/mag/01_03fa/06opt/index.shtml