

Ethernet: An Engineering Paradigm

Mark Huang

Eden Miller

Peter Sun

Charles Oji

6.933J/STS.420J

Structure, Practice, and Innovation in EE/CS

Fall 1998

1.0	Acknowledgements	1
2.0	A Model for Engineering	1
2.1	The Engineering Paradigm	3
2.1.1	Concept	5
2.1.2	Standard	6
2.1.3	Implementation	6
3.0	Phase I: Conceptualization and Early Implementation	7
3.1	Historical Framework: Definition of the Old Paradigm	7
3.1.1	Time-sharing	8
3.1.2	WANs: ARPAnet and ALOHAnet	8
3.2	Anomalies: Definition of the Crisis	10
3.2.1	From Mainframes to Minicomputers: A Parallel Paradigm Shift	10
3.2.2	From WAN to LAN	11
3.2.3	Xerox: From Xerography to Office Automation	11
3.2.4	Metcalfe and Boggs: Professional Crisis	12
3.3	Ethernet: The New Paradigm	13
3.3.1	Invention Background	14
3.3.2	Basic Technical Description	15
3.3.3	How Ethernet Addresses the Crisis	15
4.0	Phase II: Standardization	17
4.1	Crisis II: Building Vendor Support (1978-1983)	17
4.1.1	Forming the DIX Consortium	18
4.1.2	Within DEC	19
4.1.3	Within Intel	22
4.1.4	The Marketplace	23
4.2	Crisis III: Establishing Widespread Compatibility (1979-1984)	25
4.3	The Committee	26
5.0	Implementation and the Crisis of Domination	28
5.1	The Game of Growth	28
5.2	The Grindley Effect in Action	28
5.3	The Rise of 3Com, a Networking Giant	29
6.0	Conclusion	30
A.0	References	A-1

Ethernet: An Engineering Paradigm

Mark Huang

Eden Miller

Charles Oji

Peter Sun

6.933J/STS.420J

Structure, Practice, and Innovation in EE/CS

Fall 1998

1.0 Acknowledgements

The authors would like to thank the following individuals for contributing to this project. Without their invaluable assistance, much of the information gathered would not otherwise have been obtainable.

Gordon Bell

Sam Fuller

Rob Ryan

Bob Metcalfe

David Boggs

Maris Graube

The Mayfield Fund

Ken Morse

Arne Hessenbruch of the Dibner Institute

Jennifer Jones of Jennifer Jones Consulting

Lisa Alfke of Xerox PARC

Scott Kirsner of *Wired* magazine

2.0 A Model for Engineering

In 1972, *ether* was just an archaic word coined in the nineteenth century to describe the invisible medium through which light supposedly traversed. A year later, in 1973, the word became a joke between two friends at Xerox PARC—David Boggs, Robert Metcalfe—to describe the prototype networking technology that Metcalfe had developed.

Ethernet as we know it today—the industry standard for local-area networking—was born in May of that year (PBS 1998). Twenty-five years later, Ethernet has been neither disproven, as the theory of ether was by Michelson and Morley, nor discarded, as the PARC research center was by Xerox. Instead, it is a \$10 billion dollar industry that shows no signs of slowing down even after a quarter of a century—an eternity to most in the computer industry.

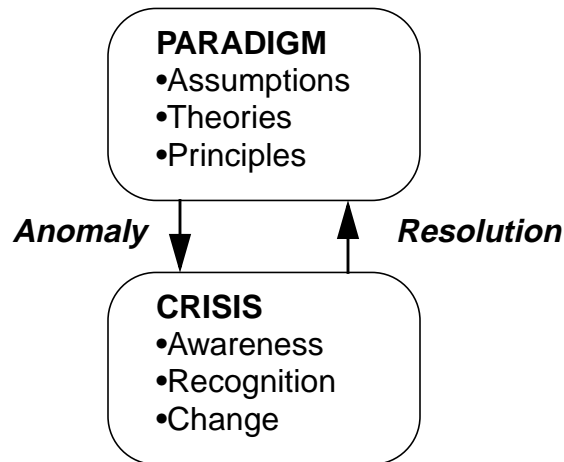
Ethernet's success has been nothing short of phenomenal. For the past three years, Ethernet and its successor, Fast Ethernet, have retained greater than 85% market share in the network interface card (NIC) industry—the most visible component of Ethernet hardware (Comm Serv 1998). The next closest competitor—IBM's Token Ring—has not been a major threat for almost a decade. Ethernet usage continues to grow at rates of 20-25% a year, and the most well-known networking equipment provider, 3Com Corporation, is fast on its way to becoming the largest, with revenues in excess of \$5.4 billion (3Com 1998).

That something so old is still in wide use today, and only now appears to be reaching technological maturity, strikes many not only as phenomenal, but also as quite anomalous. A few explanations could be given to explain why computers today are using a technology invented before the PC itself. First, it could be that Ethernet was not technically or commercially feasible until recently, and that the concept lay dormant for many years. However, large distributed networks were using Ethernet technology almost a decade ago; even before Ethernet was standardized, it was seen as a valid option for local area networks (LANs). Perhaps the need for LANs themselves only became apparent a few years ago. This explanation, too, is questionable: early network implementations like ARPANET and ALOHAnet were being explored, as well as implemented, as far back as the '60s, before Ethernet was even invented. Maybe Ethernet is only a fad, the latest in a series of Silicon Valley technologies that come and go. The revolutions in network technology that have occurred every two to three years, however, suggest otherwise. The network transformation from shared to switched, coaxial to twisted-pair, and megabit to gigabit are rebirths, rather than replacements, of the pre-existing Ethernet model.

Yet none of these explanations satisfactorily answer the question of why the standard has remained relatively unchallenged. Clearly, Ethernet possesses some unique characteristics that distinguish it from other standards. For one, it is not and has never been the winner by default, yet it has dominated the networking industry for a decade. Several alternatives—FDDI (fiber optics), Token Ring, and, now, ATM—still exist today and have their own followers, but none of them match up to Ethernet in terms of cost to performance ratios. Second, Ethernet is built upon a solid, elegant concept that is so basic that it has remained largely unchanged since its invention. Third, and perhaps most importantly, it is an open standard. Implementation is entirely up to the hardware manufacturer—or manufacturers, as the case turned out to be. Ethernet broke off from the proprietary standard track from the very beginning, and virtually all commentators attribute its success to this fact.

FIGURE 1.

Thomas Kuhn's scientific paradigm is composed of a body of knowledge—assumptions, theories, and principles, in roughly increasing order of faith—and the people who believe in that knowledge. In a time of crisis, the validity of this knowledge is reconsidered if the anomaly that spawned the crisis is recognized. When a new paradigm is developed, the crisis is resolved.



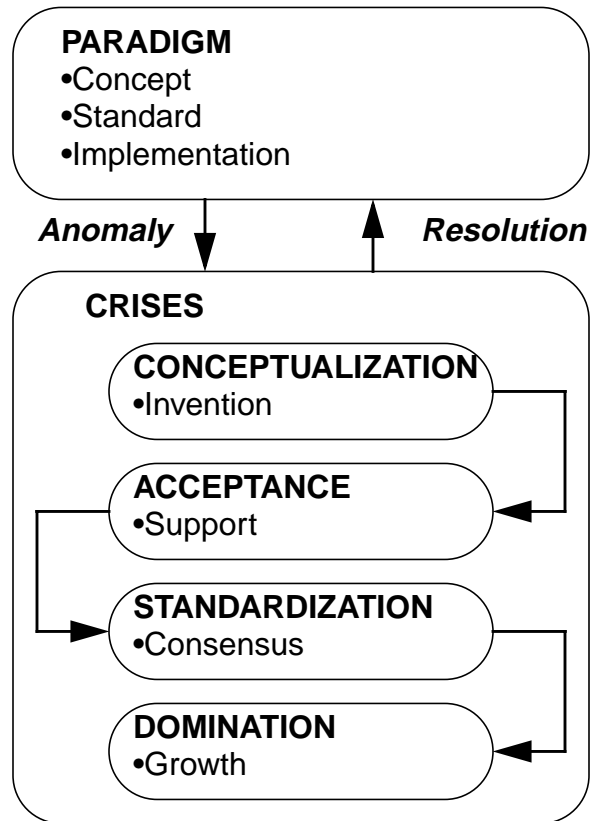
This three-phase development of an idea from concept to standard to implementation characterizes the engineering paradigm, of which Ethernet is an almost perfect example. Similar to Kuhn's idea of the scientific paradigm shift—outlined in Figure 1 on page 3—an engineering paradigm shift occurs when a new technology is found that addresses a need. This need is analogous to Kuhn's anomaly; suddenly, perhaps even because of a shift in scientific interest, a new technology becomes required. The period during which a technology is conceptualized is analogous to Kuhn's period of scientific desperation: many different theories are proposed, and each hopes to resolve the anomaly, or to meet the need in the case of engineering. When a concept becomes a standard marks the turning point in the engineering paradigm shift, much like when a new scientific theory is finally adopted. When a standard starts to be implemented is when the technology is generally considered accepted. The parallel in the Kuhn model for implementation is “normal science”—relatively straightforward, deductive work that fleshes out the details of the new paradigm. For engineers, however, implementation cannot be relegated to the realm of “normal engineering”—it is a crucial part of the engineering paradigm.

2.1 The Engineering Paradigm

The successful engineering paradigm lasts. It is characterized by innovation that refines, rather than replaces, the underlying technology over a number of years. At the same time, however, engineering paradigms tend to exist on a much shorter time scale than do scientific paradigms. Concept, standard, and implementation phases all occur within five to ten years, rather than within decades or even centuries as the case often is for scientific paradigms. When examined in sufficient detail, however, the individual phases—and the crises which come between them—are quite discernible for technologies that fit the paradigm model in question, such as Ethernet.

FIGURE 2.


The engineering paradigm is composed of three phases: concept, standard, and implementation. During these three phases, three (or four if standardization becomes a crisis) crises emerge. These crises do not have to occur in a linear fashion, but it is helpful to think of them as usually occurring in some chronological order.



Using this framework, the development of Ethernet can be seen as composed of three phases and four crises. The timeline in Figure 3 on page 5 gives a rough sketch of the chronological order of the events in its history, and will prove useful as a reference when discussing each individual phase and its associated crisis.

FIGURE 3.

Phases of Ethernet development. Along the top are the years involved. The three broad phases of the Ethernet paradigm are labeled “concept,” “standard,” and “implementation.” The four crises of Ethernet development span the three phases and overlap during certain pivotal years.



2.1.1 Concept

An engineering paradigm is based upon a fundamental concept which is necessary but not sufficient for its success. The more fundamental the concept, the more likely the paradigm is to last, but also the harder it is to nurture into a marketable product.

The concept is the base or core technology behind the paradigm. In the case of the transistor, for example, it was the application of solid-state physics to the problem of electronic amplification. In the case of Ethernet, it was the idea of packet-switching and collision detection on a local area network. Engineering concepts may come about because of a social or economic need—radar, for instance, was developed during the war to provide early warning of bomb attacks—but, more often, they simply identify a problem as solvable or a solution as feasible. Ethernet is an example of the latter: it supported the idea of the local area network (LAN) as a valid solution to the problem of interconnecting many computers.

This was the “first crisis” of Ethernet: the need to validate the concept and introduce a new paradigm. The old paradigm was dominated by time-sharing, mainframes, and the wide-area network (WAN). The shift from WAN to LAN marched in step with the shift from mainframes to distributed computing and, eventually, PCs. The Ethernet concept was before its time in that it was a mature technology by the time a networking market developed—an unusual occurrence—but it logically followed the larger shift in computing ideas happening around it.

To contrast the idea of the engineering concept with that of the scientific concept, it is helpful to think of the revolutionary scientific concept as an open-ended question looking for an answer: a theorem requiring rigorous proof. An engineering concept, on the other hand, is most often an answer—a product or a technology—looking for the right question such as an application or a market. For Metcalfe and Xerox in the late 1970s, Ethernet was the answer—as Boggs would later prove¹, it was almost the *ideal* answer for a LAN.

2.1.2 Standard

One of the key differences between the engineering and the scientific paradigms is how and by whom each is accepted. To the scientific community, a *model* accepted by the vast majority of its members summarizes the prevailing paradigm. The layman need not even be involved for the model to be useful. For engineers, the problem of acceptance is more complex. The standard of the engineering paradigm is a precise *specification* of the technology whose intended audience is composed of both industrial peers and paying customers. The specification must strike a careful balance between gaining widespread acceptance among customers and refraining from alienating industry counterparts.

Standardization is the usual compromise that addresses this issue. Ideally, many competing specifications of the same basic technology converge to a single standard. The difficulty, of course, lies in making the often arbitrary and political decisions that accompany the process of defining that single standard. Any one of these decisions can lead to a crisis in the development of the technology as a whole. In addition, once a standard is agreed upon, complying with it often stifles further innovation if the standard is not readily extensible.

Ethernet's standardization took well over a decade, and the crisis lasted almost long enough for IBM to enter the fray. IBM tended to avoid the entire standardization crisis by ensuring "*de facto* standard" status for its products, through marketing or even through design, but its Token Ring technology emerged three years too late. These were the second and third crises of Ethernet: the need to gain acceptance and technological momentum, and the need to standardize quickly enough to make implementation feasible before IBM entered the market. To overcome the second crisis, a culture of support—a combination of industrial consortiums (DIX), professional organizations (IEEE), and internal supporters (primarily within DEC and within PARC)—was mobilized to provide the "critical mass" necessary to start Ethernet out as a serious product. To overcome the third crisis, a standard was reached (the IEEE 802.3 standard) that eventually became "like a religion" to all involved.

2.1.3 Implementation

Another way that engineering and scientific paradigms differ is that the concepts and ideas that are born out of an engineering paradigm need to be concretely implemented to be successful. While it is true that scientists spend much of their lives "fleshing out" the ideas of the prevailing paradigm, it is done in the context of "normal science." For engineers, "normal engineering" does *not* include proof of concept or initial implementation: these are necessary elements of the paradigm shift. Without timely implementation, the validity of a technology can be cast in serious doubt and "normal engineering"—the refinement or enhancement of the initial implementation—may never get a chance to emerge.

-
1. In response to claims that Ethernet could not stand up to heavy loads, Boggs, Mogul, and Kent conducted a multitude of experiments at DEC that simulated various realistic load conditions. They concluded that in all but the most pathological of cases, Ethernet's performance was almost ideal. As they write in "Measured Capacity of an Ethernet: Myths and Reality," "The most well-known myth is that Ethernets saturate at an offered load of 37%. This is a fair summary of what happens for certain worst-case assumptions, but has very little to do with reality."

In the case of Ethernet, the major obstacle to its implementation was the same company that had stood in the way of its immediate acceptance: IBM. By 1985, IBM's Token Ring had become a reality. The crisis which Ethernet faced in the late 1980s was one of growth: the only proven way to avoid being squeezed out of the market by IBM was to grow faster than it—to fill the market as quickly as possible. DEC, Data General, and Xerox all had first-hand experience with this type of competition, in the field of mini- and micro-computers. Ethernet manufacturers overcame the crisis and eventually came to dominate the market because the increased quality and decreased costs that resulted from the fierce competition between them enabled the technology as a whole to fill the LAN market.

3.0 Phase I: Conceptualization and Early Implementation

The story of Ethernet's conceptualization begins in the 1960s with the networking research fostered by the United States government under the Defense Advanced Research Projects Agency (DARPA). DARPA's interest in connecting computing systems for military communications resulted in the development of a research community excited about the promise of networking and eager to explore the fundamental requirements for designing viable networks. In the early 1970s, a few key players and ideas from the DARPA effort collided at Xerox's Palo Alto Research Center (PARC) and the Ethernet revolution was born.

This section establishes the setting and the rationale for the conceptualization and early implementation of Ethernet. First, it provides the historical framework and a brief description the preceding paradigm—the computing and networking environment that predated Ethernet. Second, it defines the crisis by describing the technological and sociological anomalies facing the computing industry, Xerox, and the engineers. Lastly, it describes early Ethernet technology and how it resolved the crisis but maintained the needs of the users of the old paradigm.

3.1 Historical Framework: Definition of the Old Paradigm

The invention of a technology cannot be understood apart from the historical framework that led to the innovation. A historical framework, in this context, consists of the people, events and technologies that act as the foundation for the concepts promoted in an old paradigm. Therefore, historical frameworks define past paradigms and identify the associated engineering and social norms. Like scientific paradigms, the goal of a new engineering paradigm is to find a “better fit” solution that resolves the crisis without neglecting the desired features of the previous model. Therefore, the analysis of an old engineering paradigm often aids in the understanding of the requirements of the new paradigm by defining existing user expectations as well as new user demands (anomalies). In the case of the Ethernet, two key developments influenced the way computer users viewed inter-computer communication: first, the popularity and practicality of time-sharing; second, wide area networks (WANs) and the ARPANET.

3.1.1 Time-sharing

Time-sharing is an operating system feature that allows several users to run several tasks concurrently on one processor. Usually, each user has a dedicated terminal for input and output. Developed by MIT professor Fernando Corbató, a member of the Whirlwind project, in 1961 and extended by the SAGE military defense system in 1963, time-sharing evolved from the dissatisfaction of computer users with batch processing (Computer 1998). The expense of computers in the 1950s and 1960s prevented an institution from owning more than one or two machines. Therefore, users in the pre-time-sharing era processed information by submitting programs to the computer center and waiting for a response. Time-sharing solved this crisis by giving each user direct access to the machine and a piece of the processor time.

How does this relate to the development of the Ethernet? As a side effect, time-sharing fostered communication between users and programs and this interaction led to the initial notions of a “computer network.” In fact, in the early 1960s, a computer network was defined to be a single computer center having a multiplicity of remote terminals. Therefore, computer users developed the expectation that user-to-user and process-to-process communication was the norm. New networking paradigms, like Ethernet, were expected to provide a structure to simulate this feature of time-sharing.

3.1.2 WANs: ARPAnet and ALOHAnet

Although time-sharing provided users with the first taste of a “networked” environment, it still had many functional limits. For example, since few computer manufactures built compatible machines, the capability to transfer files between two time-sharing systems was either limited or downright impossible. This limitation also prevented communication. Messages created by users on one system had no means of being passed to users of another. Therefore, on-line interaction was limited to the local system. Communication between computer centers relied on traditional means (telephone and postal mail).

DARPA decided to tackle the development of a wide area networking (WAN) that would solve the functional limits of existing networked systems. The defense department concluded that information sharing was a key component of military operations and sought to build a nation-wide network with decentralized control and high tolerances for failure. After some initial conceptual feasibility investigations, Robert Taylor, the head of the Information Processing Technology Office, ordered full-scale development in 1966. He brought on specialists like Lawrence Roberts to supervise the creation of the Advanced Research Projects Agency Network (ARPAnet), DARPA’s WAN, and contracts were given to many academic institutions to investigate linking time-sharing systems. In October of 1968, members from each contract site, primarily composed of graduate students, formed the Network Working Group (NWG). Taylor, Roberts and the people in the NWG formed the nucleus of a small community of network specialists (Norberg 1996, 155).

From the ARPAnet research and the NWG came several key developments that influenced Ethernet development. First, DARPA adopted a communication model based on packet-switching, a concept conceived by Paul Baran at the Rand Corporation in 1960. Existing long-distance communications systems like the telegraph and telephone relied on either message-switching or circuit-switching. Both systems were inadequate for the ARPAnet for differing reasons. Message-switching is a store-and-forward technique with long delay times between transmission and reception; therefore, message-switch-

ing is ill-suited for dealing with real-time communication. Used by telephone networks, circuit-switching rely on switching stations to establish a direct, dedicated connection. Circuit-switching has the undesirable property of requiring a highly connected network since most source/destination pairs have their own unique routes. Not only is this expensive, but circuit-switched networks also exhibit little tolerance for failure. The destruction of a switch station permanently severs communication between a source/destination pair (Norberg 1996, 159). Since Ethernet was a localized extension of the WAN technology and mimicked many of the capabilities of the ARPAnet, packet-switching was chosen as Ethernet's transmission protocol.

Second, the linking of nation-wide time-sharing systems fostered further user dependence on inter-system communication. However, user expectation in the late 1960s was no longer rooted in simple social convenience (as it was earlier in the decade), but grounded in economic reality. The 1960s marked a period of rising software costs due to hardware incompatibilities. A company or government office would often purchase different computers for different tasks, and each of these systems required its own set of software. The ARPAnet eliminated the need for duplicate programming efforts within the network. In a report to congress in 1969, the economic gain of program reuse through the use of the ARPAnet network was cited as making a factor of 10 to 100 difference in effective computer capacity per dollar (Norberg 1996, 172). ARPAnet made the need for communication between machines greater, and Ethernet needed to maintain the standards set by ARPAnet.

Third, DARPA research funding supported the research of Norman Abramson at the University of Hawaii. Since Hawaii could not use the traditional WAN implementations due to an unstable telephone service, Abramson wanted to investigate radio transmission as a possible alternative. His work resulted in the ALOHAnet, a packet-switched, radio network transmitting on a single channel to a central computer. Regulation of the radio spectrum meant that the ALOHA designers did not have the flexibility of using multiple frequencies. Furthermore, Abramson and his colleagues reasoned that dedicated channels connecting every remote location to the central machine was wasteful. They suggested that instances of two or more simultaneous transmissions, also known as a collision, occur rarely in normal usage. Therefore, they devised a single channel scheme where a collision is resolved by retransmission after a random delay. Although Abramson's assumptions were correct about average load transmissions, their implementation only demonstrated an utilization efficiency of 17% under heavy load (Metcalf 1984,82). Despite the performance flaws of ALOHA, the concepts of packet-collision and shared-medium transmission were revolutionary and incorporated into the design of Ethernet. However, as the design description below will show, Ethernet did solve the efficiency problem.

Lastly, the ARPAnet project gave birth to the first breed of network pioneers. The members of the NWG were the first (and at that time, the only) experts in interconnecting systems. This tight-knit group of academics freely shared their ideas with each other and interested parties. They became the most vocal proponents for networking. Not surprisingly, many of the names associated with Ethernet are intimately connected to the ARPAnet. When companies began to hire network engineers in the early 1970s, they sought out ARPAnet members. Among the recruits that relate to this story are Robert Taylor and Robert Metcalfe. Taylor headed the research laboratory that invented Ethernet and Metcalfe designed and co-implemented Ethernet.

3.2 Anomalies: Definition of the Crisis

With time-sharing at the local level and the ARPANet as the service for long-distance communication, what problems did the old computing paradigm face? What was the crisis that brought about the necessity of Ethernet? Four anomalies combined to form the crisis. First, the emergence of minicomputers and microcomputers rendered mainframes and time-sharing obsolete. Why share a computer when you can have your own machine? Therefore, existing terminal-based networks were severed and users lost the ability to communicate and share data. Second, network usage patterns of individual users differed from the behavior of nodes on a WAN. Local usage patterns tend to occur in sudden bursts. Therefore, the low throughput supported by existing WAN technologies made it unable to address the dilemma caused by the demise of mainframes. Therefore, the need arose for a technology specifically designed for local area networks (LANs). Third, Xerox, the photocopier giant, decided to pursue a new way of office automation that involved distributed computing and single-user computers. The company's plans for an "Office of the Future" included station-to-station data transfer as well as the sharing of peripherals such as laser printers. Lastly, professional crisis on the part of the two Ethernet inventors drove them to advance the invention.

3.2.1 From Mainframes to Minicomputers: A Parallel Paradigm Shift

The disdain for multi-user computing, which began with the rejection of batch processing in the 1950s, continued in the 1960s, and led to the decline and demise of mainframes in the 1970s. Users accepted time-sharing as solution early on because of the extreme cost and size of all computers. However, a paradigm shift in transistor technology with the advent of the integrated circuit in 1958 prompted some computer manufacturers to move away from constructing bulky and expensive multi-user systems (Computer 1998). Not surprisingly, once engineers realized that they could build small, powerful machines (classed as minicomputers), many startups formed to tackle this feat. Among the first was Digital Equipment Corporation (DEC).

DEC jumped on the minicomputer bandwagon quickly with their first product, the PDP-1 in November of 1960. Although it was little more than a glorified terminal connected to a line printer, it sufficed in doing basic calculations and running simple programs. Soon, DEC and other minicomputer startups carved a niche for themselves in the commercial computing market. By 1965, the mass-production of minicomputers kicked the revolution into high gear. The gap in performance and technical sophistication also began to shrink. For example, the DEC PDP-6, built in 1964, used a 36-bit architecture and incorporated applications and a timeshared operating system. By 1970, machines like the PDP-11 series filled a large portion of the market and mainframes were "niched" into the high-end computing arena (Digital 1998).

FIGURE 4.

From left to right, a DEC PDP-1 and a DEC PDP-11.



The downfall of multi-user mainframes and the rise of single-user minicomputers led by DEC marked a manufacturing paradigm shift that influenced the development of Ethernet and acted in parallel to the paradigm shift occurring in network technologies. Since single-user machines generally operated independently of one another, minicomputers enhanced personal productivity at the expense of group communication. Users lost the local and global community brought about by time-shared systems connected to the ARPAnet. More importantly, corporations who bought into the fantastic savings of minicomputer hardware felt the pinch in buying software they needed to share. A solution was needed, and fast!

3.2.2 From WAN to LAN

Could the problems created by minicomputers be solved by the existing WAN technology? After all, connecting computers for communication and data transfer was the goal of the ARPAnet. Anyone observing the ARPAnet project in 1970 would witness a robust infrastructure of ideas, technology and experts. Unfortunately, researchers attempting to apply WAN principles to local area networks (LANs) met with a startling discovery. Users expectation of local communication mechanisms conflicted with the established goals of WANS. The usage patterns at the local level favored multiple high bandwidth transactions occurring in small time intervals rather than a consistent stream of low to moderate transmissions. This, of course, conformed to the experience that these users had with the time-sharing systems. Communication and data sharing on a mainframe was bounded by time-slices and the throughput rates of various transport elements, such as the input/output devices and the bus. So, thus arose a second anomaly, users wanted a high-bandwidth, “bursty” technology and the existing technology provided a poor match.

3.2.3 Xerox: From Xerography to Office Automation

The severing of local interaction caused by the widespread acceptance of minicomputers and the inadequacies of existing systems in addressing the transmission behavior of

LANs both illustrate the breakdown in the existing networking paradigm caused by expectations of the user base. However, a different sort of problem was brewing at Xerox Corporation that equally affected the way people think about and use computers. The 1960s was a decade of unmatched success for Xerox. They had converted their proprietary duplication process into a multi-billion dollar monopoly. However, by the late 1960s, Peter McColough, CEO from 1968 to 1981, envisioned a new direction for his company. He argued that Xerox needed to diversify into other areas of office automation in order to maintain their breakneck growth. One of his predictions was the impact of computing and networks in the office environment. As a result, he ordered the creation of the Palo Alto Research Center (PARC), a center focused on exploring the future of office automation (Smith 1998). PARC fostered an academic environment conducive to research. It attracted some of the best minds in the field of computing research, and the company immediately reaped intellectual rewards. By the early 1970s, the company could already boast of multiple designs for distributed computing environments and laser printing. However, there was a missing piece, the network. They had blueprints for a personal workstation and a revolutionary printing device for output generation, but they had no means of interconnecting the two. This was a difficult predicament; an anomaly that prompted Xerox to search the ARPAnet community for someone with a solution.

3.2.4 Metcalfe and Boggs: Professional Crisis

Xerox's commitment to research and its lofty goal of a distributed computing system attracted first rate talent to PARC. For example, one major research arm of PARC, the Computer Science Laboratory (CSL), hired Robert Taylor, of ARPAnet fame, as the manager. Other big names like Butler Lampson and Alan Kay were assigned the monumental role of designing the first personal computer, the Alto. However, the inventors of Ethernet, Robert Metcalfe and David Boggs, had humble beginnings at Xerox and were motivated more by basic professional necessities rather than images of grandeur. Their personal struggles illustrate how social difficulties played a key role in the formulating the crisis of conception.

FIGURE 5.

The Xerox Alto.



When Robert Metcalfe joined the CSL in 1972, he was in a real bind. Having just failed his PhD defense at Harvard, he joined PARC with a great deal of uncertainty and much to prove. Fortunately, Metcalfe made some important friends while he was doing his PhD research. During his tenure at Harvard, he found a job working on the ARPAnet and MIT's Multiple Access Computer Project (Project MAC) (Metcalfe 1994, 81). He recognized the potential of the government project in terms of sources of funding (interview). The opportunity also opened the way for him to enter into the inner circles of the networking community. His membership in the NWG made him one of the few elite experts in networking. Therefore, his acceptance into PARC without a degree was hardly a surprise; especially since the head of the CSL, Taylor, was part of the brotherhood. However, the shock and resentment from his rejected thesis drove him forward.

Unlike Metcalfe, David Boggs joined Xerox in 1972 with little fame. In fact, the starving graduate student hoped that his internship at PARC would help pay the bills for his PhD at Stanford. He learned about PARC through a weekly seminar hosted by local companies given at Stanford. The networking research being done at PARC intrigued him. As a teenager, Boggs was fascinated by radios and fell in love with communications ever since. When he arrived, he was assigned the tedious task of assembling Data General Nova computers. The job was not incredibly exciting; and after a while, he wanted out (interview).

3.3 Ethernet: The New Paradigm

Dissatisfaction by the community, the unique attributes of LANs, Xerox's need for interconnecting their personal computer to peripherals, and professional dissatisfaction by Metcalfe and Boggs, all combined to form a crisis begging for a new concept in networking. By 1972, there was an undeniable need for a LAN technology that could handle high-speed, "bursty" traffic, deal with a variety of device connections, and encourage communication and program sharing. Metcalfe proposed Ethernet as the

solution. Ethernet's design considered many of the problems it needed to solve and elegantly addressed all the anomalies.

3.3.1 Invention Background


During the first two months at Xerox, Taylor tasked Metcalfe with the design and the construction of ARPAnet interfaces for the PARC Multiple Access Xerox Computer (MAXC). Since this kept him intimately involved in the ARPAnet circles, he continued to attend ARPA events and promote the ARPAnet. During June of 1972, on an ARPA related visit to Washington. He stumbled onto the Aloha System paper written by Professor Abramson on a friend's table. Jet-lag caused Metcalfe a fit of sleeplessness and he began to read the paper. Being a packet networking engineer and an applied mathematician, he was immediately fascinated by the heavy-load randomized-retransmission problem and the inelegance of the existing solution that was based on Poisson queuing theory (Metcalfe 1994, 83). He started work that night to change the model. Four months later, he published the paper "Steady-State Analysis of a Slotted and Controlled Aloha System with Blocking." He demonstrated that traffic-based control of retransmission could increase utilization to 36%.

After the paper, Metcalfe began to design the entire system in earnest. Although Ethernet is loosely based on Aloha, many salient features were rejected or modified. For example, Ethernet does not rely on a central computer to monitor the signal. Instead, every machine will have its own carrier detection system and only attempt to transmit when the channel is free. Another difference is packet length. Ethernet uses variable packet sizes which also improved packet efficiency. Most importantly, Metcalfe designed a sophisticated traffic back-off control technique for retransmission after conflicts that increased system stability to almost the ideal level (Metcalfe 1984, 82). He submitted this algorithm as an additional chapter of his PhD thesis. The revised document was eventually accepted in June of 1973.

Between May and June of 1973, Metcalfe typed up his ideas about Ethernet as a series of memos to the Alto group (Metcalfe 1994, 84). Eventually, the legal department of Xerox suggested that he file an invention record and then a patent. However, his work was only half done. Ethernet looked fine on paper, but Xerox wanted a product for the Alto. Therefore, he set out to build a prototype. Metcalfe used the basement of the PARC building to lay out the coaxial cabling and clobbered together a crude transceiver. Unfortunately, not being very experienced at building electrical devices (including how to solder), he struggled at every step. Meanwhile, working near Metcalfe in the basement was Boggs. After sharing a few weeks in the basement, they became friends. Since Boggs graduated with an electrical engineering degree from Princeton and was being underutilized at his existing position, he offered to help Metcalfe with the prototype. Eventually, Metcalfe approached Boggs' boss and convinced him to let David work on the Ethernet project for the summer of 1973 (interview). By 1974, the two engineers had built and debugged a 100-node experimental Ethernet. After building the prototype, Boggs continued to work on the hardware designing and building the first transceivers and repeaters. Metcalfe, however, shifted to more academic pursuits and wrote a paper on Ethernet in 1974 as well as filing a patent on it on March 31, 1975 (Metcalfe 1994, 84-87).

FIGURE 6.

Early sketch by Metcalfe of the Ethernet idea.



3.3.2 Basic Technical Description

How exactly does Ethernet work? Figure 6 on page 15 shows an early sketch of Ethernet drawn by Metcalfe. The technology consists of a single coaxial cable connected to various devices. Each connection begins with a physical tap embedded in the wire by a transceiver. Each transceiver includes, in addition to the usual transmitter and receiver sections, a gate that compares the data from the interface stage with the data on the cable. An unequal result means a collision has occurred. The transceiver is then attached to an interface stage that does encoding/decoding and buffering. The encoding is done using clock-driven shift registers that convert the buffered data into a serial stream, feeds the data into a phase encoder, that then connects to the driver section of the transceiver. Decoding is accomplished in a similar way in reverse order. There is an address decoder attached to the input shift register to identify packs with correct/incorrect addresses. A random number generator and a counter that counts the number of collisions may be embedded into the device or the interface section to schedule retransmission.


3.3.3 How Ethernet Addresses the Crisis

Ethernet resolved the crisis of conception by providing an interesting academic problem for Metcalfe and Boggs to solve; a shared, high-speed communication system that replaced the terminal network of the mainframe era, and a product that was specifically integrated into the Alto and the laser printing devices built by Xerox. In fact, the intended use of the Ethernet was clearly outlined in *United States Patent 4,063220: Multipoint Data Communication System with Collision Detection*. In the background section of the document, Metcalfe addressed issues of applying the technology to interconnecting a variety of devices including a laser printer. Figure 7 on page 16 duplicates the generalized system diagram. Notice how the modules are designed so that some functionality can be internalized in the using device (Metcalfe 1977). This permits software implementations of the collision detection functionality or even buffering and encoding/decoding. He also specified the high-speed capabilities of a short wire as a transmission medium and cited a 3Mbit per second throughput rate. In 1975, that was approximately 60 times faster than more WANs. Finally, the sophisticated retransmis-

sion protocol coupled with signal detection and variable packet size made Ethernet ideal for “bursty” traffic. Unlike telephone lines and a circuit-switched network, the line is never completely occupied by one user and everyone gets a share of the action.

FIGURE 7.

Figure from the original Ethernet patent.



By 1979, the crisis of concept had been resolved. Metcalfe and Boggs continued to bug fix and expand the prototype network at Xerox so that it encompassed multiple sites and thousands of nodes. By then, the technology that began on a sleepless night in Washington was proven. However, it was not yet an engineering paradigm. In order for that to happen, it had to leave Xerox and dominate the market.

4.0 Phase II: Standardization


4.1 Crisis II: Building Vendor Support (1978-1983)

Successful evolution from a research technology to a commercial product mandates the creation of a market. This implies a degree of both vendor and customer acceptance. In order for Ethernet to emerge as a viable LAN technology, a bridge had to form between Xerox PARC and a number of outside computing manufacturers powerful enough to create a market niche. These external industries, in turn, faced many of the same issues of confronted by Xerox in the previous section. This time, however, Ethernet had to emerge victorious over in-house solutions as well as other externally developed LAN specifications. In many ways, the second crisis of our model exhibits the same symptoms as the first crisis. The difference occurs after the crisis is resolved—instead of the one convert, Xerox, Ethernet succeeded in proselytizing small industrial following. Proponents of Ethernet aimed to push compliant products onto their customer base in hopes of encouraging other manufacturers to build Ethernet compliant products. Through this strategy, they hoped to establish *critical mass* in the LAN marketplace.

Bob Metcalfe explained critical mass in the following manner. If you are designing a network, the cost varies linearly with the number of nodes attached to the network. Value, on the other hand, is a quadratic function of the number of nodes. Critical mass occurs at the intersection of value and cost (interview). In essence, it is the moment when the value of a network exceeds its cost and this drives market acceptance (see Figure 8 on page 17).

FIGURE 8.

Metcalfe's Law. Value grows quadratically, while cost grows linearly with the number of installed nodes.



This relationship is often referred to as Metcalfe's Law in deference to its creator. It seems overly simplistic, however, to assume that critical mass depends on the size of the network alone. For this reason, I prefer a different definition of critical mass, one that

encompasses a broader range of variables such as the number of products being developed, the strength of corporate backing, the cost of implementation, and the size of a customer base. When these factors converge, a technology can establish a presence sufficient enough to create a market; this enables future adoption, development and propagation of the technology in a self-sustaining manner. In effect, critical mass constitutes a crucial moment in creating an industry standard because it demonstrates both customer and vendor acceptance. The process of attaining critical mass can lead to an industry crisis, in which both users and vendors must decide which competing LAN specification to adopt. Since computer networking demands compatibility almost by definition, the quest for widespread acceptance assumes a particularly salient role. Resolving this crisis oftentimes mandates the creation of an industry standard.

As an open standard, Ethernet had a distinct advantage in rallying vendor support. Peter Grindley's model for open source versus proprietary standards illuminates how Xerox's decision affected Ethernet's future. Grindley, a denizen of the London Business School, treats the proprietary/open source debate as a management decision. If a manager decides to make his product an open standard, the market will grow. His company, however, will receive a smaller percentage of market revenue. On the other hand, if the manager opts for a proprietary standard his company will acquire a full share of the profits, but product demand could dwindle in a smaller market. The IBM vs. Macintosh story illustrates the tension between these two opposing forces. IBM's decision to make their product an open source standard restricted the company to a significantly smaller profit margin. Yet, the resultant size of the PC market almost drove the competing Macintosh out of existence (Grindley 1995).

Companies concerned with widespread LAN compatibility were attracted to Xerox's open standard because they recognized the potential for market growth. In fact, the popular narrative of how DEC and Intel came to endorse Ethernet centers on their conviction that "proprietary network schemes will be less successful than public-type networks that enable users to select the equipment they prefer and whose communication costs are low" (Elec News May 1980). The open standards issue played a crucial role in forming the DIX consortium, but it does not fully explain why DEC and Intel decided to embrace Ethernet over competing LAN technologies. The remainder of this section will explore this process of acceptance within DEC and Intel, analyzing the complex web of motivation and tension that evolved along the way.

4.1.1 Forming the DIX Consortium

On May 13, 1980 DEC and Intel announced their intention to join Xerox in pushing Ethernet as an industry standard for LAN architecture. This allegiance, eventually referred to as the DIX consortium, marked the first instance of Ethernet bridging the gap between Xerox PARC and the outside corporate world. The formation of this triumvirate signified an acceptance of Ethernet's technical validity and an investment in its future development. As Robert Brannon, the marketing manager of Intel, stated, "Intel is going for broke on Ethernet" (Comp Dec Jan. 1982). Moreover, the set of expertise brought together by this consortium provided a powerful alliance of diverse technical skill and resources capable of integrating Ethernet into a wide range of computers and peripheral devices. DEC planned to use their knowledge of micro- and mainframe computer networks to develop transceiver hardware compliant with the Ethernet specification. Intel hoped to imbed Ethernet controllers in microprocessor boards, providing users with a cost-effective means of linking computing machinery. Xerox held both the

Ethernet patents and the tacit knowledge resulting from years of designing and maintaining Ethernet-based LANs. Moreover, Xerox's decision to make Ethernet an open source specification unofficially sanctioned the formation of the DIX consortium by removing all suspicion of antitrust violations. All three parties agreed that compatibility was key to the success of LAN architecture; managers within DEC and Intel were convinced Ethernet's open source specification would prove more fruitful in the long run than proprietary solutions. Or so the story goes. In reality, the managerial conviction portrayed by the trade journals veiled the role played by circumstance and the uncertainty associated with a corporate gamble.

4.1.2 Within DEC

In 1978 Sam Fuller, then manager of the VAX architecture project, found himself faced with the predicament of selecting a local area network technology to use as an addition to VAX. At the time, an official commercial market for LAN technology had not emerged. Fuller readily admits DEC did not consciously examine their customers and discern a "need" for local connectivity. Yet, DEC employees had heard stories of customers wiring DEC PDP-11 computers together in an ad hoc fashion in attempts to create a network interface. Other customers had made isolated requests for a means of linking PDP-11 computers together using some type of multi-megabyte interconnect. Therefore, while DEC product managers did not actively establish a pre-existing market for LAN technology, they had enough knowledge to suspect LAN technology would provide a profitable addition to VAX. DEC had already started developing a proprietary approach to the LAN problem, known as the DEC Dataway. The Dataway, however, was designed with the intention of interconnecting factory computer systems and did not allow multiple access. Fuller had to decide if the Dataway provided the best networking solution for the office. Recognizing the size of his task, he began looking for outside advice. Shortly thereafter in 1979, Metcalfe left Xerox PARC with the intention of founding 3COM. Fuller, who remembered Metcalfe from his days in graduate school, convinced him to work for DEC as an outside networking consultant until 3COM became a viable commercial venture.

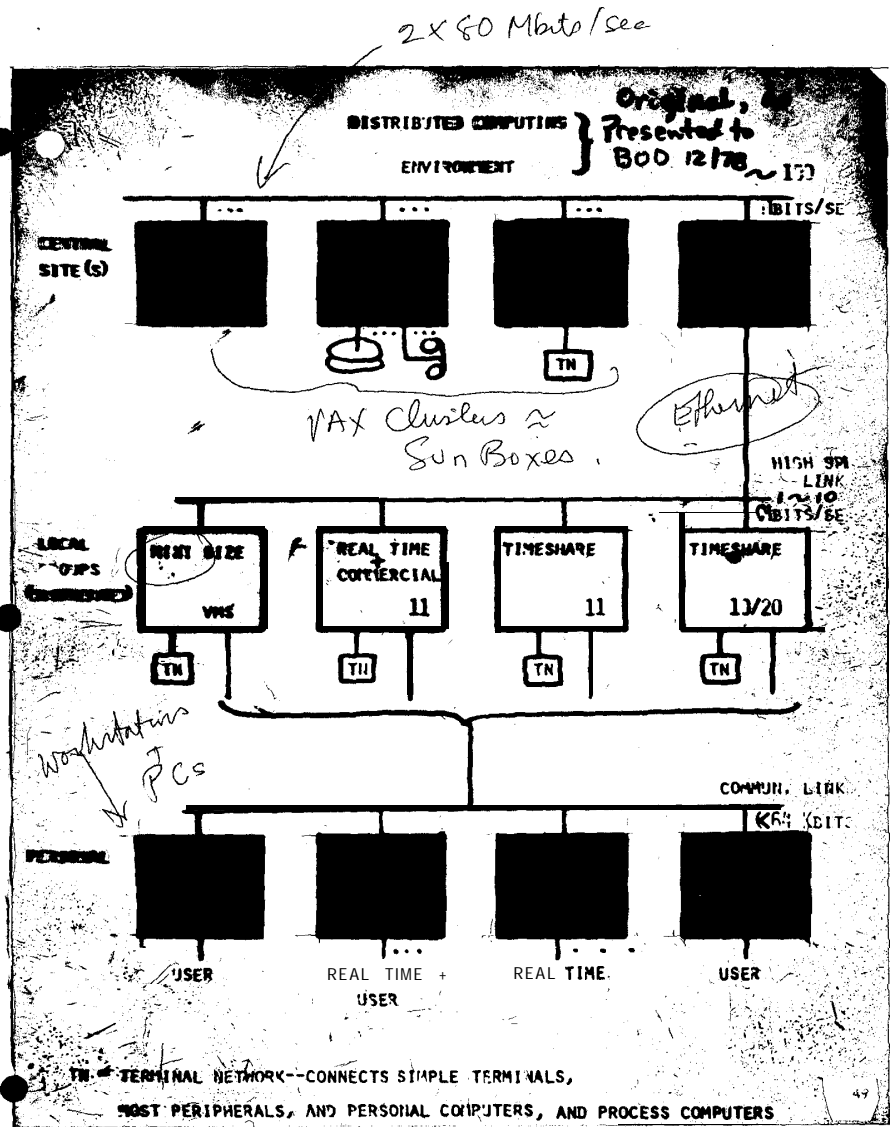
If Fuller expected an objective opinion, he certainly did not receive one from Metcalfe. From the moment Metcalfe arrived at DEC, he began extolling the benefits of an open standard put together by industry. Fuller proved to be an easy audience, recalling that "within a few days of talking about it, it became fairly obvious that the most effective networking effort would be an open standard. It would be pretty arrogant of any company to assume that they could make all of the products" (interview). Under Metcalfe's guidance, Fuller also became convinced of Ethernet's potential to fill the role of open source LAN panacea. His personal certainty, however, was not enough to convince DEC's upper level management of Ethernet's technical validity or the merit of open source standards in general. Recent corporate decisions indicated a rather inconsistent approach to the open source vs. proprietary standards debate—a practice illustrated by DEC's decision to replace the highly successful Unibus and Q-bus I/O adapter architecture (an open source standard) with the proprietary BI architecture. Moreover, DEC had a vested financial interest in their proprietary LAN solution, the Dataway. Consequently, Fuller had to confront these bureaucratic, political, and economic biases successfully if DEC was to adopt Ethernet.

Three major factors emerged that contributed to DEC's eventual acceptance of Ethernet: Gordon Bell, Xerox PARC documentation, and the structure of the engineering divi-

sions within DEC. Bell, the Vice President of Engineering at DEC, recognized the merits of Ethernet early on and became a staunch advocate of Ethernet within the six months following Metcalfe's arrival. Bell's personal notes on the VAX Basic Product Strategy (Figure 9), written in 1978, suggests Bell considered Ethernet even before Metcalfe arrived at DEC. Having supervised a thesis on ring networking while still a Professor of Computer Science at Carnegie Mellon University, Bell "knew what they could do" and recognized their technical limitations (interview). He played a major role in Ethernet's advancement, both inside and outside of DEC; Metcalfe later named him as a co-inventor of the Ethernet industry standard (Metcalfe 1994). In fact, Bell initiated the discussions with Dave Liddle of Xerox that led to Ethernet's transformation from a proprietary solution to an open source standard. As VP of engineering, Bell held a large amount of sway within DEC, enabling him to persuade CEO Ken Olsen as well as with other members of the DEC management chain. Furthermore, Bell was able to convince the engineers who designed the VAX processors of Ethernet's technical merit—a crucial step, considering these engineers would eventually have to build VAX hardware in accordance to Ethernet specifications. Once Bell converted the engineers, specifically the head engineer of the VAX group Bill Demmer, Fuller and Bell knew they had strong ground to stand on when arguing for DEC's support of Ethernet. Bell continued to champion Ethernet up until the day he left DEC in 1983, as illustrated by a speech he gave to the Xerox/Intel/Digital Seminar in 1982. In it he states, "I believe that Ethernet is one of the keys to the development of the Fifth Generation just as the Digital Unibus was one of the keys to the development of the minicomputer generation" (Bell 1982). Considering the previous success of DEC's Unibus, Bell's comparison is both powerful and optimistic.

FIGURE 9.

Page from the original VAX product strategy document. Note Bell's handwritten reference to "Ethernet" in the upper right, and the date (12/78), months before Metcalfe's arrival at DEC.



The Blue and White Series reports distributed by Xerox provided another outlet for PARC engineers to assert Ethernet's superiority. Fuller recalled, "The reports from Xerox PARC were great marketing and PR. I was looking for what ought to be our next generation and the first thing that came to mind was Ethernet" (interview). A brief glimpse though these reports reveals how an outsider could perceive Ethernet as a supe-

rior LAN technology. First, they project Ethernet as a mature technology. References to doctoral dissertations on Ethernet design, implementation, and performance evaluation suggested that Ethernet technology had been examined considerably and subsequently refined. Moreover, the PARC reports are laced with debugging tips ranging from the practical (the most effective way to locate a network short) to the mundane (what to do if a lightning storm brings down the network). The sheer number and resourceful nature of these debugging strategies illustrated the longevity of Ethernet's operation within PARC and the depth of Xerox's experience maintaining Ethernet LANs. Second, the reports demonstrate a well-established record of performance measurements, such as utilization and performance statistics. This enabled executives and engineers to predict Ethernet's product functionality without resorting to conjecture. Third, the reports emphasized simplicity both in the existing design and in their suggestions for future improvements ("The Ethernet Local Network" 1980, Boggs 1979, Metcalfe 1975).

Finally, the structure of the engineering divisions within DEC helped Bell and Fuller overcome the social inertia that led many to loyally cling to the DEC Dataway. As it so happened, the Dataway project was located in DEC's Industrial Products Research Division—a sector divided between engineering and management jurisdiction. Due to this split, the principle advocates of the Dataway did not report to Bill Demmer and he, in turn, could assume a more dispassionate and constructive view of their work. After some consideration, he decided to support Ethernet's adoption and suggested abandoning the Dataway project. As head engineer of the VAX project, Demmer had considerable influence in deciding which LAN technology DEC would incorporate in VAX. With the demise of the Dataway, DEC management quickly realized that Xerox held the patents for a product with great potential and DEC officially joined forces with Xerox shortly thereafter. This realization, however, also resulted in contention. Bell recounted that once Ethernet's value became apparent "the questions were 'why can't we make our own and have it proprietary?'" Ken [Olsen] hassled everyone on the cable and eventually all the key guys left" (interview). Thus, while DEC publicized the need for an open LAN standard, this history reveals the ambiguity of their commitment. Moreover, it demonstrates the circumstantial and political influences that contributed to Ethernet's acceptance within DEC.

4.1.3 Within Intel

In the late 1970's, Rob Ryan accepted a position as a Project Manager at Intel. Before this, he worked in an architecture group at DEC developing routing and congestion control algorithms for DECnet Phase I and Phase II. Having witnessed the push for networked computers within DEC, Ryan was shocked to learn that Intel had no plans to develop a communications chip. He recalled, "I was interested in getting them involved in what I thought was going to be an explosive market" (interview). Ryan reasoned that what drove computer sales would also drive chip sales; Intel had the potential to make a considerable profit by diversifying into network communications. He approached Phil Kaufman, then Head of Architecture at Intel, with his idea. Together they formed Intel's Systems Division. This group, headed by Ryan, acted as the initial instigators for LAN technology within Intel.

Both Ryan and Kaufman had contacts at DEC and Xerox and they immediately began talking with representatives from both companies—Kaufman with the management hierarchy and Ryan with the engineering "worker bees." Shortly thereafter, the Systems

Division started designing the first Ethernet chip, better known as the 82586. This was the first time in Intel's history that the design of a complex chip occurred outside of the Chip Division and, naturally, this caused internal friction. Ken Fine, the head of the Chip Division, wanted to back IBM's Synchronous Data Link Control (SDLC) and argued that Intel would be wasting their time on new technology with an uncertain future. Since Intel has a reputation for fostering a confrontational environment, it comes as no surprise that members of the Chip Division openly clashed with members of the Systems Division over which networking specification to adopt. This debate traveled up the management hierarchy and eventually reached the desk of Andy Grove, then the President of Intel. After hearing both sides of the argument, Grove sided with Ethernet.

We can only speculate on the rationale behind Grove's decision. Ryan believes that he and Kaufman simply presented a stronger case. It seems doubtful, however, that this alone would determine the course of Intel's future in the networking market. Certainly, the corporate momentum generated by Xerox and DEC contributed to Ethernet's appeal. Moreover, Ryan had already secured two of Intel's best engineers, Bob Beach and George Marshall, to work on the 82586 project. In fact, Intel had just ranked Beach as the top engineer in the company with Marshall following behind in the number two slot. The technical prowess of the Systems Division undoubtedly helped Ethernet gain acceptance from both engineering and management. Yet, I believe one of the most persuasive factors was Grove's conviction that Intel could persuade IBM to adopt Ethernet. In this sense, Grove's acquiescence highlights an underlying corporate gamble, in which Ethernet's success was far from certain.

In the spring of 1980, five representatives from DEC met with five representatives from Intel via AT&T's PicturePhone Meeting Service. After a two-hour meeting, both companies "shook hands" and jointly agreed to push the Ethernet specification. Gordon Bell then drafted a letter to Dave Liddle of Xerox, outlining the agreement reached by DEC and Intel and officially invited Xerox to join the consortium. Thus, the DIX consortium was born.

Although this brief glimpse into the internal workings of DEC and Intel has helped illustrate the complex process of accepting a new technology within a bureaucratic setting, one must remember that the story is still far from complete. The narrative presented here has been heavily influenced by individual testimony and thus reflects a composite of the motivation and subjectivity exhibited by each interviewed source. At times, different accounts of the same events have implied both ambiguity and contradiction. While linear narratives, such as the one presented here, help convey the crucial points in Ethernet's acceptance, we must remember that the actual embrace of this technology occurred in a more organic fashion—perhaps over a series of informal phone calls and meetings held between different points in the corporate strata. This revised model, albeit too complex to explore in this paper, can help us account for discrepancies in credit assignment and event portrayal.

4.1.4 The Marketplace

Commercial interest played a large role in pushing local area network technology into the computing marketplace, as revealed by the stories of DEC and Intel. This does not answer the question, however, of whether public enthusiasm stemmed from a pre-existing user need or from the projected need advocated by LAN proponents. An article

from 1980 cites Boeing Computer Services Company (BCS) as one of the earliest LAN customers. They purchased “Hyperchannel” from Network Systems Corporation in 1977 with the intention of creating a heterogeneous computer network immune to issues of compatibility in each of their three main processing centers (EDP Analyzer June 1980). Similarly, the story of DEC customers independently wiring PDP-11’s together suggests that this desire was not isolated to Boeing. These early examples, coupled with the Association of Data Communication Users (ADCU) 1979 conference theme on the “philosophy and the implementation of integrated networks”, suggests that the motivation behind LAN development originated both within and without the computer industry (Infosystems March 1979). Or, as MacKenzie suggests, “needs are created simultaneously with a means of fulfilling them” (93).

Although at this stage, Ethernet’s survival depended on establishing critical mass within industry rather than catering to an abstract concept of need. The DIX consortium aimed to achieve critical mass in the LAN marketplace by collectively pushing a technology they believed would emerge as a de facto industry standard. Building the desired vendor support and increasing Ethernet’s presence in commercial development required certain concessions. For example, when DIX released the first version of its Ethernet specification in October of 1980, Xerox announced they would charge a one-time \$1,000 licensing fee. This figure barely covered the processing cost of the license. Xerox, however, recognized the importance of enticing outside vendors to build Ethernet compliant products, and consequently sacrificed immediate profit in favor of allowing the market to grow.

Xerox demonstrated their endorsement of Ethernet as early as 1979 when they incorporated the networking capability into the new Model 860 word processing system. By November 1980, the company had developed a series of shared resource processors that allowed Ethernet compliant machines to share hard disk files and a low-end laser printer. This processor enabled Xerox customers to connect their older 800 and 850 word processors to Ethernet networks. The fact that Xerox felt a need to incorporate a relatively new technology, which lacked a solid customer base, into their older products demonstrates the company’s commitment to market penetration. If Xerox customers wanted to install networking capabilities in their office, they simply had to upgrade their processor and become Ethernet compliant. Although the list price of the processor was costly by today’s standards (\$20,000), when compared to the price of upgrading an entire word processing system, it provided a cost-effective option for many businesses (Elec News Nov. 1980). The act of making Ethernet backward compliant also illustrates Xerox’s desire to push Ethernet into business markets as quickly as possible.

Similarly, DEC incorporated Ethernet as one component of DECnet, their multi-faceted networking solution. DEC’s needs differed from those of Xerox: they had to interconnect micro-computers, not word processors and laser printers, and their primary concern was compatibility. They did not share Xerox’s exclusive loyalty to Ethernet. An advertisement for DECnet published in 1983 reads, “No single universal standard has emerged. Nor is one likely to... That’s why Digital is committed to supporting and, in fact, promoting the more important standards now surfacing in the various environments” (Data Comm 82). Ethernet, however, served as DEC’s preferred hardware implementation, mainly because it conformed to the framework defined by the Digital Network Architecture. DEC customers could easily incorporate Ethernet capabilities into DECnet without having to change the higher-level software protocols. Fuller

recounts, “We sort of from day zero assumed we would probably run DECnet over the Ethernet cables, using the Ethernet datalink for high performance local area networking” (interview). The flexibility offered by DECnet enticed customers who were wary of the confusing and rapidly changing network market. Therefore, DEC’s commitment to a varied networking toolbox helped create a market for Ethernet, where it evolved into an accepted LAN standard.

Intel, however, provided the greatest contribution to Ethernet’s success by making the technology affordable. In the days before VLSI technology, designing integrated circuits (ICs) from scratch constituted a rather expensive task. In 1980, Intel started building the first Ethernet chip, the 82586. By 1982, they marketed the first LSI chips that incorporated the Ethernet specification as well as a two board Multi-bus controller (ISBC 550). Although the ISBC started as a prototype, the demand for the board eventually created an economy of scale and drove the price down from \$4000 to \$2000 (Mini-Micro May 1982). In 1983, they successfully used VLSI technology to place the functional capabilities of four boards onto a single board—an accomplishment that again lowered the price of Ethernet hardware.

By pushing Ethernet compliant products into the marketplace, the DIX consortium tried to raise consumer and vendor awareness of the potential for Ethernet-based LANs, as well as create a de facto industry standard. When the DIX consortium released the “Blue Book” standard for Ethernet I in 1980, the specification had already undergone a massive transformation. In contrast to PARC’s “seat of the pants” specification that “read like a novel”, the new specification took pains to “prove correctness” through mathematics (Ryan interview). If you look at the original Blue Book, you immediately notice that the authors carefully outlined the Ethernet specification in three ways: through a narrative, through mathematics, and through an algorithm. In a literal sense, this revised framework facilitated the construction of Ethernet compliant products by making the standard more rigorous. However in a metaphorical sense, the shift in emphasis from conceptual narrative to technical detail provides an illustrative manifestation of our Phase I to Phase II progression in engineering revolutions.

4.2 Crisis III: Establishing Widespread Compatibility (1979-1984)

Ethernet’s proponents tried to create a commercial market capable of withstanding competing solutions proposed by companies such as Wang. The efforts put forth by the DIX consortium helped Ethernet build vendor support and acquire a preliminary customer base. This approval, both from within and without the DIX consortium, signified the resolution of the second crisis brought about by a need for acceptance.

Despite this initial momentum, Ethernet lacked a firm commitment from the commercial market. Vendors such as Hewlett Packard hesitated to adopt a technology embraced by their competitors (DIX). Other vendors refused to support a particular implementation before IBM announced their intended LAN strategy. Proponents of competing technologies, such as token passing or slotted time division networks, emphasized Ethernet’s technical limitations; these actions simply augmented market indecisiveness. The demonstrated lack of vendor commitment to a single technology brought about the third crisis in our model for engineering revolutions: the need for widespread compatibility. In the case of Ethernet, the formation of the IEEE 802.3 provided the desired resolution.

Note that this third crisis does not pertain to all engineering revolutions. Oftentimes, an industry standard emerges without outside intervention. For example, competing technologies may not be able to secure a foothold in the marketplace if a company with substantial market power endorses a particular implementation. High entry costs may also bar competing firms from developing alternative approaches. In these circumstances, proprietary or closed technologies flourish.

Ethernet, however, did not enjoy these advantages and consequently had to rely on a subsequent standardizing effort—this one backed by the IEEE. Although the IEEE by definition serves as a professional organization for electrical engineers, they also act as a standardizing body operating under the American National Standards Institute (ANSI). Their efforts to create a LAN standard not only refined DIX's version of the Ethernet specification, but also forced vendors to decide which LAN technology to adopt. The creation of the IEEE 802.3 standard, a revised version of DIX's Ethernet, and the subsequent adoption of this standard by ECMA (European Computer Manufacturers Association) and the ISO helped convince a broader array of companies to build Ethernet compliant products. To put it another way, if the DIX consortium gave Ethernet access to the playing field, then the IEEE let them stay in the game.

4.3 The Committee

In 1976, Maris Graube accepted a position at Tektronix working on the IEEE 488 standard, an interface for programmable instrumentation. This early interconnect functioned only over distances of less than 20 meters, causing Graube to look for a means of linking “computer-like things over longer distances” (interview). In search of a solution, he joined forces with the Perdue Workshop for Industrial Process Control and began developing an early LAN capable of connecting computers in an oil refinery or chemical plant. Graube first envisioned constructing and standardizing a less demanding LAN specification for the office during this time. He recounts:

I happened to talk to a fellow named Bob Stewart who was involved with the IEEE's microprocessor committee and he said, 'why don't you do a standard under the auspices of IEEE', I said, 'How do I go about doing that,' and he said, 'well you fill out this piece of paper called a project authorization request and I'll help you get it approved.' (interview).

The IEEE approved Graube's request in the fall of 1979. He chaired the first IEEE 802 committee meeting in February 1980.

The circumstances surrounding the 802 committee's inception highlight the need to account for parallel development when discussing engineering paradigms—meaning that the four stages of crisis need not occur in a set order or in serial fashion. The industrial standard and the IEEE standard originated independently of one another during the same period. In fact, Graube never expected 70 people, most of whom were representing industry heavyweights such as DEC, Xerox, and IBM, to attend the first 802 meeting.

The 802 committee began as a single attempt to create a universal LAN standard. In 1981, the committee drafted a compendium of every LAN technology under consideration for the 802 standard. The resulting document exceeded 500 pages in length. As discussions progressed, two distinct technologies dominated committee meetings:

Ethernet and Token Ring. DIX supported the former implementation and IBM the later. During meetings, these dueling factions frequently drove technical debates to a stalemate and eventually hampered all progress. Even more detrimental, these corporate feuds added an element of arbitrariness to committee decisions. Oftentimes the outcome of a technical debate depended solely on the size of the contingency sent by each party. Fuller recalled, “we had a lot of people going to the IEEE standards meeting and it was for them quite frustrating because there was always the politics of what now became a large industry pushing their own agenda” (interview). Bell was more blunt. “We wanted [the IEEE standard] imprimatur, but were fairly pissed about IBM’s attempts to derail it with their ring” (interview).

Finally in 1982, Graube decided to split the 802 committee into several subcommittees, each of which were more or less autonomous. This action served a dual purpose, only one of which was intended. First, it freed the 802 committee from the previous political impasse. Second, and perhaps more important, the formation of the subcommittees delayed the confrontation between Ethernet and IBM. In fact, it gave the appearance that the IEEE endorsed both technologies equally. The finalized version of the IEEE 802.3 standard, released in 1983, not only helped Ethernet emerge victorious over less popular LAN technologies, but also set the stage for the future market battle between Ethernet and Token Ring.

However, Ethernet did not emerge from standardization process unscathed. The committee aimed to improve the DIX Blue Book Standard by extending the frame length, adding error correction capabilities, augmenting the bandwidth, increasing the data rate, and reversing the packet order. The 802.3 committee tried to create a technically superior version of the DIX Ethernet standard, yet the acceptance of a proposed revision oftentimes depended on factors external to technical merit. In the early 1980’s, for example, designing integrated circuits from scratch entailed a complex and expensive process. Therefore, the existing functionality built into Intel’s Ethernet chip provided a limiting framework for future revisions. Graube referred to a conversation with Intel representative Phil Arst where he asked if Intel could change their chip to have a better delimiter. Arst replied, “We can’t do that, we’re too far along.” Due to this silicon inertia, a technically superior revision took a back seat to corporate interest, even though it later “screwed up things royally” (interview).

Regardless of how or why they occurred, the changes proposed by the IEEE created a compatibility rift between 802.3 and the DIX blue book standard. Trade magazines published circa 1983 alluded to a market tension between the two standards and speculated over which would survive. In reality, the perceived friction between the two implementations appears over-hyped. From its inception, the DIX Consortium pushed for a widely accepted LAN standard based on Ethernet technology. When the IEEE 802.3 specification fulfilled this goal, DEC, Intel, and Xerox found little use for the consortium and the 802.3 standard quickly supplanted the DIX standard. Thus, as Metcalfe states, “The IEEE monster did not kill Ethernet, rather it became Ethernet” (Shotwell 55).

5.0 Implementation and the Crisis of Domination

5.1 The Game of Growth

By the late 1980s, it was not certain to any extent that Ethernet would come to dominate the LAN market. IBM's Token Ring technology—a vague threat in 1983, but a reality by 1985—was a major adversary. Technically, by 1985, Token Ring's proven capabilities were on par with Ethernet: it could provide adequate bandwidth (4 Mbps vs. Ethernet's theoretical 10 Mbps), at comparable prices (\$830 per node at its introduction vs. Ethernet's 1982 introductory price of \$950 per adapter) (Allocca 1998). The real fear that Ethernet manufacturers harbored was that with the support of IBM's incredible marketing power, Token Ring would grow to fill the entire market for LAN technology. To combat IBM, the manufacturers, paradoxically, fought amongst themselves: fierce internal competition within the Ethernet ranks improved the quality and affordability of the technology in a short period of time. The competition did not end without the usual casualties: mergers, failures, and takeovers. Three of the largest companies that started out in the early '80s—Bridge Communications, InterLan Corporation, and Ungermann-Bass—all ceased to exist independently by 1991, swallowed by the industry's explosive growth.¹

5.2 The Grindley Effect in Action

The vendor-neutral nature of Ethernet 802.3 enabled this growth. The distinction between vendor-neutral and vendor-independent is an important one to consider. Had Ethernet been vendor-independent—if vendors had each been able to implement their own versions of Ethernet in the late 1980s—this growth phase would never have arisen. It was only because the technology was standardized and (as far as any of the small startups were concerned) immutable that Ethernet moved ahead because of the internal competition. No one vendor's product had an inherent, technical advantage over those of the others; they were all producing the same product.² In other words, Ethernet as a standard was vendor-neutral. Where a vendor did have leeway was in the manufacturing of its units. Through the introduction of Very Large Scale Integration (VLSI) technology and Application Specific Integrated Circuits (ASICs) into production lines, both cost reductions and performance gains were realized by a large number of vendors. Both of these technologies caught on very quickly: as soon as one vendor made an incremental improvement in the cost or yield of its production, others rapidly followed suit.

At the other end of the spectrum was IBM. Texas Instruments won exclusive rights to the production of IBM's Token Ring chipsets very early on. Though vendors could license Token Ring for use in their products—3Com itself hedged its bets in 1986 and began producing Token Ring adapters—they could not design their own chips. Since it was not under any pressure from competing manufacturers, TI soon became the bottle-

1. Bridge Communications was acquired by 3Com in 1987. InterLan was acquired by Racal in 1991. Ungermann-Bass was acquired by Tandem Computers, now a division of Compaq, in 1987.

2. This statement isn't entirely accurate. Some older implementations that were still being manufactured after 1985 were not strictly 802.3-compliant. In almost all cases, however, they would work on an 802.3 network, because of the backward compatibility of the standard.

neck in IBM's attempts to lower the cost of Token Ring. The result was a price tag that was, at its peak, double that of comparable Ethernet adapters (Allocca 1985).

The Grindley effect, first introduced in Section 4.1 on page 17, was clearly apparent. 3Com, InterLan, Bridge, and the rest divided Ethernet's share of the LAN market, and thus also profits, between them: lean times came and went in the early 1980s as each company struggled to stay afloat. IBM retained firm and solitary control of its share and absorbed most of the profits from Token Ring. As Ethernet's total market share grew by tremendous proportions, however, so did the divided shares of the participating companies, eventually to the point in 1987 when 3Com, the upstart startup, became the largest independent networking manufacturer in the world. IBM never matched the growth of Ethernet and learned the hard way that 100% of nothing is nothing.

5.3 The Rise of 3Com, a Networking Giant

3Com Corporation, while not the only one of the early Ethernet manufacturers, was certainly representative of the bunch. Co-founded by Metcalfe himself in 1979, it lay dormant for almost three years before an adequate market for LAN emerged. In 1982, the Intel 8086-based IBM PC was introduced, almost a decade after Xerox PARC's prototypes of the Office of the Future. Had Ethernet not been invented when it was—if a mature LAN technology had not been available in 1982—it seems very likely that a crisis would have shortly emerged as enterprises purchased large quantities of PCs and began to ask how they could be networked. Instead, the answer to the crisis had been waiting for years for precisely this problem—these questions—to emerge. Ethernet was designed around the idea of connecting together a large number of disparate systems (computers, printers, mainframes, servers, and more). When owning large numbers of these systems became possible with the PC revolution, Ethernet and 3Com suddenly, but not unexpectedly, found themselves with a market.

By 1984, when 3Com went public, the company's sales had grown from \$4.7 million to \$16.7 million. The 802.3 standard was approved by the IEEE in 1983, but it was not until subsequent approvals by the American National Standards Institute (ANSI) and the International Standards Organization (ISO) in 1984 that 3Com truly benefited from standardization. 3Com continued to grow at astronomical rates, posting revenues of \$64 million in 1985. That same year also saw the introduction of IBM's Token Ring, however. Though revenues were good, 3Com's position in the LAN market was far from secure: by 1986, they had captured only 8 percent of the LAN market, while IBM remained smug with 28 percent.

Rather than continuing to emphasize their strength in LAN connectivity through their flagship product, the Etherlink transceiver, 3Com began to diversify early under Metcalfe's leadership in the marketing department. In 1986, 3Com's software division began testing the waters of the fledgling network operating system (NOS) business with a product called 3+. 3+ and 3Server, a network server computer and another 3Com attempt at diversification, formed their first complete network system. The system sold briskly at first, but a new competitor, Novell Corporation, soon began taking over the market with its incredibly popular NetWare system. By 1987, 3Com had for the most part given up on the market and returned its focus to expansion. In September 1987, 3Com acquired Bridge Communications and for the time being became the largest manufacturer in the industry.

3Com would later make another failed attempt at diversification, this time in conjunction with Microsoft. The LAN Manager "fiasco," as Metcalfe recalls it, involved a complicated deal that soon turned sour. The software was supposed to provide powerful capabilities to network administrators. LAN Manager, like 3+, eventually lost to a competing product from Novell, but the part of the deal that most hurt 3Com was a convoluted term of the contract. Under the term, it was forced to pay royalties to Microsoft for sales that it did not even gain revenue from. In addition, it was also bound to provide product support for these sales. Eventually, 3Com severed its ties with the project after taking heavy losses (Kepos 1995).

Throughout these years, however, 3Com's position in the NIC market had steadily solidified. Sales of its Etherlink transceivers had provided a steady source of revenue for many years but, more importantly, they had also branded the market with 3Com's name. In 1989—the same year that 3Com abandoned LAN Manager—the one millionth 3Com Ethernet adapter was sold, just a year and half after the half million mark was reached (3Com 1998). In 1990, Metcalfe was passed over for the position of CEO in favor of Eric Benhamou. Immediately, Benhamou, like Metcalfe and former CEO Bill Krause before him, embarked on a crusade to diversify 3Com. This time, however, 3Com's efforts were largely successful. The game had been won: Ethernet had grown to fill the LAN market just in time for the '90s. Only after its core long-term goal had been fulfilled was 3Com able to possess the security of a dominated market that was necessary to diversify. Today, 3Com produces a wide range of products including NICs, routers, switches, broadband, and, just recently, even personal organizers such as the wildly popular PalmPilot.

6.0 Conclusion

The timeline presented in Figure 3 on page 5 stops at 1990, primarily because so much has happened in the past eight years that discussing recent events would comprise an entirely new project. Ethernet adapters, once close to \$1,000 per node, are now available for less than \$30 at retail outlets. The Ethernet 802.3z standard now defines Gigabit Ethernet, a protocol one hundred times faster than the original. The World Wide Web continues to transform the needs of the average computer user, and companies like 3Com have moved quickly to meet those needs.

The concept, standard, and implementation that define Ethernet have not been fundamentally altered, however. Despite the fact that hundreds of incremental improvements have been made to Metcalfe and Boggs' original design, a 1998 Ethernet adapter differs little in basic functionality from a 1988 adapter. The concept has remained the same.

Similarly, the standard and the implementation have changed in detail but not in intent. IEEE standards 802.3A (10Base-2), 802.3C (repeaters), and 802.3I (10Base-T) were all introduced after the original, but the intent of the 802.3 standard—to emphasize that vendors manufacture the same *technology* as much as they make different products—remains the same. 3Com's Parallel-Tasking Architecture marked one of the earliest major changes in Ethernet adapter design, but it, too, was simply a refinement of the original implementation.

Conclusion

We used these constants—concept, standard, and implementation—as the basis for our model for engineering. We defined crises as the periods of time during which new constants are calculated. Much like how Ethernet resolves conflicts among a chaotic mess of packets by imposing order, we were able to classify disparate, often contradictory accounts of Ethernet’s history by placing them into our framework. Our model for engineering is by no means the only way of looking at the development of a technology. In the case of Ethernet, however, it provides at least one way of telling the whole story.

Conclusion

A.0 References

- 3Com. *3Com Fast Facts*. <http://www.3com.com/new3com/glance.html>.
- Bell, Gordon. "Why Digital is Committed to Ethernet for the Fifth Generation", a summary of remarks made at the Xerox/Intel/ Digital Seminar, February 10, 1982. Report provided by Gordon Bell.
- Bell, Gordon, Former Vice President of Engineering, DEC, email interview conducted by Eden Miller, December 4, 1998, concerning his role promoting Ethernet's acceptance at DEC.
- Boggs D., J. Shoch, E. Taft, and R. Metcalfe, "Pup: An Internetwork Architecture," *Xerox PARC Blue and White Series Reports*, P7900031, July 1979.
- Boggs, David R., Mogul, Jeffrey C., & Kent, Christopher A. "Measured Capacity of an Ethernet: Myths and Reality." *WRL Research Report* 88/4 (1988): 1-31.
- Boggs, David R., Former engineer at PARC, Xerox, telephone interview conducted by Peter Sun, December 2, 1998, concering his role in the invention of Ethernet.
- Communications Service. *1998 Networking*. <http://research.instat.com/netwrk/netwrk98.htm>.
- Computer. *Timeline of Computer History*. <http://www.computer.org/computer/timeline>.
- Dahod, Ashraf M., "Local network standards: No utopia," *Data Communications*, 12 (March 1983), 173-180.
- Davis, Dwight B. "Ethernet spec spawns diverse group of interface products," *Mini-Micro Systems*, XV, No. 5 (May 1982), 143-151.
- Davis, Dwight B.. "Xerox moves to bolster embattled Ethernet," *Mini-Micro Systems*, XV, No. 2 (February 1982), 17-18.
- "DEC, Intel, Xerox Form Local Data Com Net Team," *Electronic News*, 26 (May 19, 1980), 1.
- "Desks of America Unite," Advertisement for Xerox 860 Information Processing System, *Computer Decisions*, 14, No.1 (January 1982), 20.
- Digital. *Digital Celebrates 40th Anniversary*. <http://www.digital.com/flash/b40/>.
- "Ethernet", Poster created by the Mayfield Fund commemorating Ethernet's 25th Anniversary. Donated by Jennifer Jones Consulting.
- "Ethernet criticized," *Computer Decisions*, 14, No. 1 (January 1982), 12-14+.

References

- Fuller, Sam, Former Head of the VAX Architecture Project, DEC, telephone interview conducted by Eden Miller, November 24, 1998, concerning his role in establishing Ethernet's acceptance within DEC.
- Graube, Maris, Former Chair of the IEEE 802 committee, Tektronix, telephone interview conducted by Eden Miller, November 30, 1998, concerning his role in the IEEE 802 committee.
- Grindley, Peter. *Standards, Strategy, and Policy: Cases and Stories*, Oxford: Oxford University Press, 1995.
- Hessenbruch, Arne, Postdoctoral Fellow in the History of Science, Dibner Institute, interview in Cambridge, MA conducted by Eden Miller and Peter Sun, November 20, 1998, concerning the role of standards in history.
- "In Your Future: Local Computer Networks," *EDP Analyzer*, 18, No. 6 (June 1980), 1-13.
- Kirsner, Scott. "The Legend of Bob Metcalfe," *Wired*, 6.11 (November 1998), 182-186+.
- Kuhn, Thomas. *The Structure of Scientific Revolutions*, Chicago: The University of Chicago Press, 1962.
- Lettieri, Larry. "Joint Effort turns out Ethernet chips, boards," *Mini-Micro Systems*, XV, No. 11 (November 1982), 71-74.
- MacKenzie, Donald. *Inventing Accuracy*, Cambridge: The MIT Press, 1993.
- "MDS Operators Rap Xerox Net Proposal," *Electronic News*, 25 (January 15, 1979), 25.
- Metcalfe, R.M. and D. R. Boggs, "Ethernet: Distributed Packet Switching for Local Computer Networks," *Xerox PARC Blue and White Series Reports*, P7600002, November 1975.
- Metcalfe, Robert M., "How Ethernet Was Invented", *IEEE Annals of the History of Computing*, 16, No. 4, (1994), 81-88.
- Metcalfe, Robert M., Former research scientist at PARC, Xerox, telephone interview conducted by Peter Sun, December 8, 1998, concerning his role in the invention of Ethernet.
- Metcalfe, Robert M., David R. Boggs, Charles P. Thacker, and Butler W. Lampson. *United States Patent 4,063,220: Multipoint Data Communication System with Collision Detection*. December 13, 1977.
- "Networking standards. Nobody keeps your options open like Digital," Advertisement for DEC, *Data Communications*, 12 (March 1983), 82.

References

- Norberg, Arthur L. and Judy E. O'Neill. *Transforming Computer Technology: Information Processing for the Pentagon, 1962-1986*. Baltimore: The Johns Hopkins University Press, 1996.
- "The office of the future gets a common voice," *Business Week*, No. 2638 (May 26, 1980), 57-58.
- Public Broadcasting Service. *Nerds 2.0.1*. http://www.pbs.org/opb/nerds2.0.1/serving_suits/3com.html.
- Rosenberg, Ronald. "Xerox Regards Ethernet as an AT&T Alternative," *Mini-Micro Systems*, XIII, No. 2 (February 1980), 53-54.
- Saxton, W.A. and Morris Edwards. "The integrated network movement," *Infosystems*, 26, No. 3 (March 1979), 78-80.
- Shotwell, Robyn, ed., *The Ethernet Sourcebook*. New York: Elsevier Science Publishing Company, Inc., 1985.
- Smith, Douglas K. and Robert C. Alexander. *Fumbling the Future: How Xerox Invented, Then Ignored, the First Personal Computer*. New York: William Morrow and Company, Inc., 1988.
- "The Ethernet Local Network: Three Reports," *Xerox PARC Blue and White Series*, P8000013, February 1980.
- "We've got what makes the difference in Local Area Networking," Advertisement for ACC Exchange System, *Data Communications*, 12 (March 1983), 9.
- "Would Cook' Litton Ovens," *Electronic News*, 25 (January 15, 1979), 25.
- "Xerox, Intel, DEC Release Specs for Ethernet System," *Electronic News*, 26 (October 6, 1980), 16+.
- "Xerox Processors Debut for Ethernet," *Electronic News*, 26 (November 4, 1980), 30.
- "Xerox's bid to unlock the office of the future," *Business Week*, No. 2617 (December 24, 1979), 47.