

**ETHERNET PRESS SEMINAR**  
**WORLD TRADE CENTER—NEW YORK CITY**  
**FEBRUARY 10, 1982**

**Schedule of Events**

- 8:15–9:00 Coffee & Danish
- 9:00–9:40 Introduction: **The Impact of Very Large Scale Integration (VLSI) Technology on Communications**  
*Dr. Robert Noyce, Vice Chairman, Intel Corporation*
- 9:40–10:20 **The Evolution of Distributed Processing and Local Area Networks**  
*C. Gordon Bell, Vice President, Engineering, Digital Equipment Corporation*
- 10:20–10:30 Intermission
- 10:30–11:15 **Productivity in the Office Environment**  
*Mr. David E. Liddle, Vice President and General Manager, Office Systems, Office Products Division, Xerox Corporation*
- 11:15–12:00 Panel Question and Answer Session
- 12:00–1:15 Luncheon

**ETHERNET EDITORIAL SEMINAR**  
**WORLD TRADE CENTER — NEW YORK CITY — FEBRUARY 10, 1982**

**Schedule of Events**

- 1:30—1:50 **Introduction: The Impact of Very Large Scale Integration (VLSI) Technology on Communications**  
*Dr. Robert Noyce, Vice Chairman, Intel Corporation*
- 1:50—2:10 **The Evolution of Distributed Processing and Local Area Networks**  
*C. Gordon Bell, Vice President, Engineering, Digital Equipment Corporation*
- 2:10—2:30 **Productivity in the Office Environment**  
*Dr. David E. Liddle, Vice President and General Manager, Office Systems, Office Products Division, Xerox Corporation*
- 2:30—2:45 Intermission
- 2:45—3:05 **The Changing Economics for Computer and Terminal Interconnection**  
*Phillip Arst, Intel Corporation*
- 3:05—3:25 **Ethernet Performance Based on a Simulated User Environment**  
*William Hawe, Principal Engineer, Systems Performance Analysis, Digital Equipment Corporation*
- 3:25—3:45 **Dispelling Ethernet Myths**  
*Bob Prentis, Manager of Network Standards*
- 3:45—4:15 **Panel Question and Answer Session**  
*Moderator: Ronald T. Yara, Strategic Marketing Manager, Intel Corporation*
- 4:15—5:00 Hospitality Reception

## **ROBERT N. NOYCE**

Intel Corporation  
Santa Clara, California

Robert N. Noyce is Vice Chairman of the board of directors of Intel Corporation, Santa Clara, California. A co-founder of Intel Corporation in 1968, Dr. Noyce was President until 1975 and chairman of the board from 1975 to 1979.

Dr. Noyce is co-inventor of the integrated circuit with Jack Kilby. They have jointly received the Ballantine medal of the Franklin Institute, and the Cleo Brunetti Award of the IEEE for this work. With Gordon Moore he has received the AFIPS Harry Goode award for leadership in computer science. Dr. Noyce was awarded the National Medal of Science and the I.E.E. Faraday Medal in 1979, and the IEEE Medal of Honor in 1978. He is a member of the National Academy of Science, the National Academy of Engineering, the American Academy of Arts and Sciences, and is a Fellow of the IEEE.

Dr. Noyce was born in Iowa in 1927. He received a B.A. degree and membership in Phi Beta Kappa at Grinnell College (Iowa) in 1949, and a Ph.D. in physical electronics at the Massachusetts Institute of Technology in 1953. He did research at Philco Corporation until 1956 when he joined Shockley Semiconductor Laboratory, Palo Alto, California, shortly after its founding, to work on transistor technology. (The lab was founded by William Shockley, co-inventor of the transistor at Bell Telephone Laboratories.)

In 1957, Dr. Noyce co-founded Fairchild Semiconductor Corporation, Mountain View, California. He was research director until early 1959, when he became vice president and general manager. By 1968, the sales for Fairchild Semiconductor had risen to over \$100 million.

As research director of Fairchild Semiconductor, Dr. Noyce was responsible for initial development of the firm's silicon mesa and planar transistor product lines. Also, his inventions in the integrated circuit field enabled Fairchild to produce the first commercial integrated circuit.

In July, 1968, Dr. Noyce co-founded Intel Corporation with Gordon E. Moore, who had also been a co-founder of Fairchild Semiconductor and a member of the Shockley laboratory staff. (Dr. Moore succeeded Dr. Noyce as President and then Chairman of Intel.)

Their goal was to make LSI technology a practical reality. At the time, LSI was still in its early stages of development and used primarily to produce custom circuits. Intel developed the Schottky barrier bipolar and the silicon gate metal-oxide-semiconductor technologies which allowed several thousand transistors to be integrated on a single chip of silicon with a relatively high production yield. Intel used the silicon gate MOS technology to produce the first high density memory components and the first microprocessor. It now produces most of its LSI products with advanced versions of this technology.

Dr. Noyce holds 16 patents for semiconductor devices, methods and structures.

Intel has grown to approximately 16,000 employees. In 1979 revenues totaled \$663 million and net income \$77.8 million. Intel manufactures and markets large scale integration (LSI) and VLSI semiconductor devices, such as microprocessors and memory components, and systems built with LSI devices.



## Impact of VLSI on Communications

We are on the eve of major developments in worldwide data communications on all fronts—within the factory and office, between buildings and cities, and between countries. Indeed, we now hear of the “Second Industrial Revolution,” the “Paperless Society,” the “Information Age,” and the “Knowledge Revolution” from popular writers or news reports, from conferences of industry and labor leaders, and studies of governments and learned societies. All of these global forecasts point to an increased need to communicate and to expand those communications facilities to reach a much broader community of users.

The semiconductor industry is, in large part, responsible for the enormous demand for increased data communications. In its first decade, the microprocessor has been designed into more than 100,000 products. The development of standard VLSI building blocks has allowed manufacturers to introduce microcomputer-based products at an unprecedented rate. The fact that these new systems are becoming increasingly interdependent will result in data communications networks (to interconnect those systems) becoming as pervasive as the microprocessor is today. It is important, therefore, that the same orientation toward global optimization, which resulted in the development of standard microcomputer building blocks, be continued by the semiconductor industry, equipment manufacturers, and end users in the defining and implementing of advanced data communications capabilities.

The impact that VLSI will have on communications must be viewed in the context of the impact the semiconductor industry has had on computing. Through standard building blocks, manufacturers were able to drive costs down, while increasing capabilities, to change the economics of computing from “one for many” to “one for one.” Whereas the large mainframe and expensive system resources imposed a “one for many” environment, and the lower-cost minicomputer a “one for few” relationship, VLSI (microprocessors memories and software) is now making possible the era of the personal work station, a “one for one” relationship. The impact of VLSI on communication, then, can be seen as an opportunity to provide cost-effective interconnection of those personal work stations and the centralized capabilities supported by minicomputers and mainframes, providing uniform access to information, resources, and services.

Global optimization in communications will lead to interconnection that achieves:

- A. Location-independent access to
  - information
  - resources
  - services
- B. Media-independent access
  - telephone wire
  - coaxial cable (TV cable)
  - fiber optic cable
  - others
- C. Interoperability—different equipment from different manufacturers communicating with each other.

The benefits of such a solution are:

- A. Timely access to, and distribution of, information, independent of where the user is or the transmission medium used.
- B. Cost-effective sharing of both distributed and centralized resources and services.
- C. Optimized end user solutions that can include equipment from multiple manufacturers.

How will such a solution come about? As with other established markets, it will be an evolving process. Because of the growing demand for data communications, however, the development of capabilities and architectures by various manufacturers will be rapid, running well ahead of the actual installation of such equipment. It is important, therefore, to start now in the definition of standardized interfaces that will lead to both location and media-independent access and interoperability. This will also allow standard VLSI communications building blocks to be developed, resulting in an impact on communications analogous to the impact that microprocessors had on computing—that of driving the economics of communications from “few to one” to “any to any.”

The three companies; Digital Equipment Corporation, Intel Corporation, and Xerox Corporation, have made that first step in the area of standardized, high-speed communications within the building: Ethernet. Starting with the basic Ethernet technology that had been under development and testing at Xerox since 1975, the three companies entered into a cooperative agreement. That agreement involved the development of a high-speed, Local Area Network, and publishing the specification to encourage general, widespread implementation. The goal of the three companies was to achieve interoperability within a building through a standard, high-speed, Local Area Network.

Intel's contribution to the cooperative effort has been focused in two areas. We currently have an extensive family of LSI communications peripherals that support existing protocols, including HDLC/SDLC and Bisynch. Using that base of knowledge, we are utilizing our VLSI expertise to develop a high-speed Local Area Network Controller which will support the Ethernet specification, which we expect to sample before the end of the year. We have also implemented Ethernet in several systems products, the first of which is a distributed microcomputer development system that begins customer shipment this quarter.

The efforts of the three companies, I believe, reflect the industry "sense" of the problem to be solved and the need for cooperation. In many respects, we are sitting in the same position the railroad industry was in when they saw the opportunity to provide freight and passenger service throughout the country—the time was at hand to agree on the width of the railroad tracks. Ethernet provides that "standard width" for *integrated* solutions within a building. Similar efforts are required to allow the strengths of VLSI solutions to be properly focused to provide "cost effective one for one" computing *and* "any to any" communications, thereby achieving the "Wired (World) Community."

intel®

**THE IMPACT OF VLSI  
ON COMMUNICATIONS**

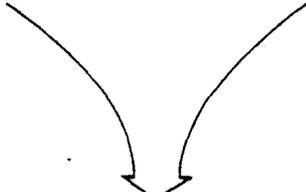
intel®

**WE ARE ON THE EVE OF  
MAJOR DEVELOPMENTS IN  
WORLDWIDE DATA COMMUNICATIONS**

- Within the Factory and Office
- Between Buildings and Cities
- Between Countries

intel®

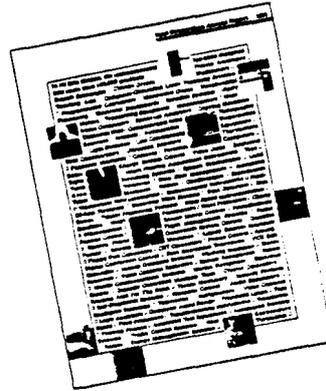
- “Paperless Society”
- “Second Industrial Revolution”
- “Information Age”
- “The Knowledge Revolution”



**UNPRECEDENTED DEMAND FOR  
COMMUNICATIONS**

intel®

**IN ITS FIRST DECADE, THE  
MICROPROCESSOR HAS BEEN DESIGNED  
INTO MORE THAN 100,000 PRODUCTS**



intel®

**DATA COMMUNICATIONS WILL BECOME  
AS PERVASIVE AS MICROPROCESSORS  
ARE TODAY**

- Orientation Toward Global Optimization
  - Semiconductor Industry
  - Manufacturers
  - End Users

intel®

**THE SUCCESS OF THE SEMICONDUCTOR  
INDUSTRY**

- A Healthy Mix of Competition and Cooperation

intel®

**"ALL COMPETITORS WHO BOTH COEXIST AND  
ARE PROFITABLE OVER TIME ARE  
SIGNIFICANTLY DIFFERENT."**

**—BRUCE HENDERSON**

□ **THE SEMICONDUCTOR INDUSTRY PLAYERS ARE  
DIFFERENT IN:**

- Objectives**
- Strategic Approach**
- Management Style**
- Products and/or Markets**

intel®

**AGAINST THIS BACKDROP OF INTENSE  
COMPETITION, A HISTORY OF COOPERATION**

**—Conscious and Deliberate**

**AN ORIENTATION TOWARD GLOBAL OPTIMIZATION**

intel®

**CONSCIOUS AND DELIBERATE  
COOPERATION**

- **Cross Licensing and Second Sourcing**
- **Joint Ventures**
- **Standards for Packages, Functions, etc.**

intel®

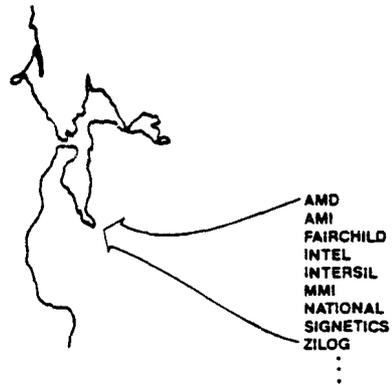
**AGAINST THIS BACKDROP OF INTENSE  
COMPETITION, A HISTORY OF COOPERATION**

- Conscious and Deliberate
- Unconscious

AN ORIENTATION TOWARD GLOBAL OPTIMIZATION

intel®

**SILICON VALLEY  
GEOGRAPHIC CLUSTERING**



intel®

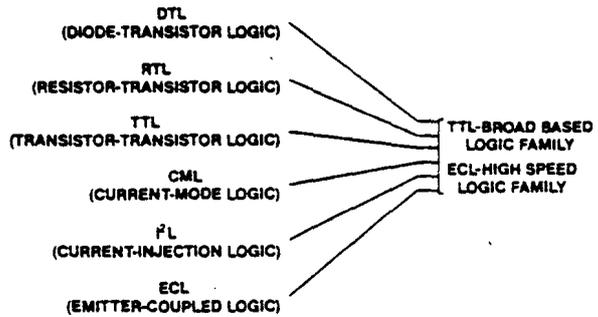
**AGAINST THIS BACKDROP OF INTENSE  
COMPETITION, A HISTORY OF COOPERATION**

- Conscious and Deliberate
- Unconscious
- Market Consolidation—Defacto Standards

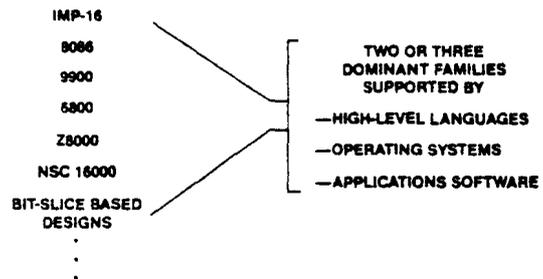
AN ORIENTATION TOWARD GLOBAL OPTIMIZATION



## MARKET/TECHNOLOGY CONSOLIDATION — BIPOLAR TECHNOLOGIES



## MARKET/TECHNOLOGY CONSOLIDATION — 16-BIT MICROPROCESSORS



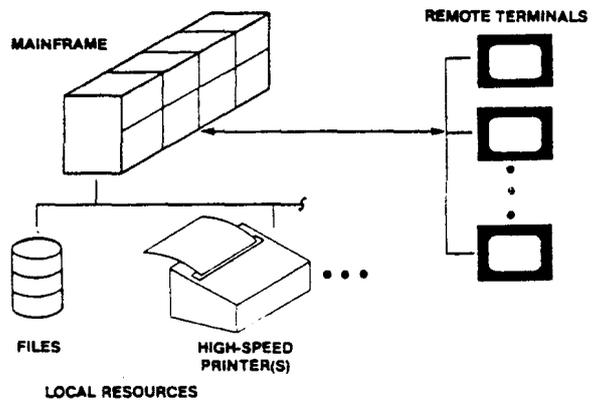
## BENEFITS OF GLOBAL OPTIMIZATION

- Semiconductor Industry
  - Better Use of Available Resources
  - More Competitive, on a Broader Front
- Our Customers
  - Standard Building Blocks with Wider Application
  - Develop Better Solutions, Sooner and at Lower Cost

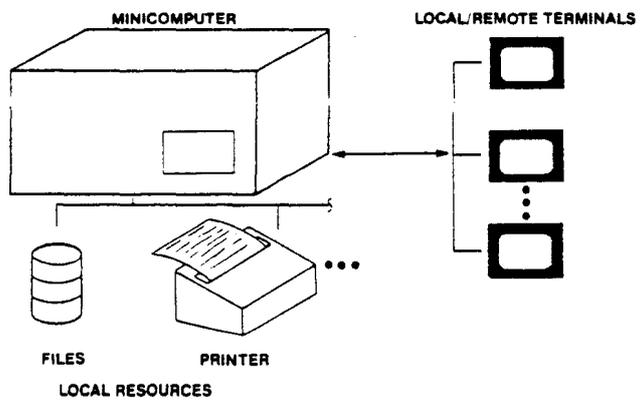
### IMPACT OF THE SEMICONDUCTOR INDUSTRY ON COMPUTING:

Changing the Economics of Computing from  
"One for Many" to "One for One"

#### THE MAINFRAME: ONE FOR MANY



#### THE MINICOMPUTER: ONE FOR FEW



intel®

### THE PERSONAL WORKSTATION: ONE FOR ONE



- Local Resources and Local Computer
- Files                      —Keyboard
- Printer                    —Display

intel®

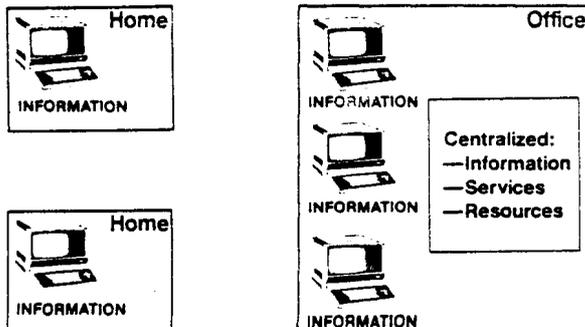
### THE OPPORTUNITY FOR COMMUNICATIONS IS TO PROVIDE COST-EFFECTIVE INTERCONNECTION

“ONE FOR ONE” AND “ONE FOR MANY”  
(Distributed)                      (Centralized)

- Information
- Resources
- Services

intel®

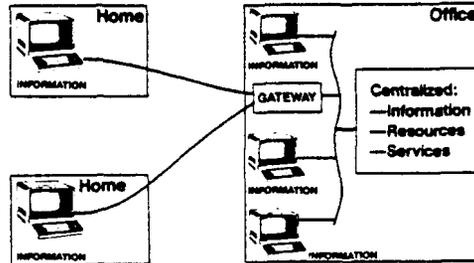
### A PREREQUISITE TO MANAGING INFORMATION IS THE ABILITY TO ACCESS IT



intel®

### LOCATION AND MEDIA INDEPENDENT ACCESS TO:

- Information
- Resources
- Services



intel®

### INTEROPERABILITY—ANY TO ANY COMMUNICATIONS USING EQUIPMENT FROM DIFFERENT MANUFACTURERS

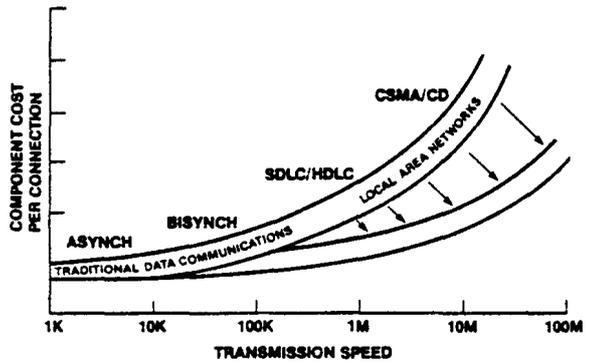
- Providing Access to
  - Information
  - Resources
  - Services

intel®

### THE KEY—STANDARDIZED INTERFACES

- Orientation Toward Global Optimization
- Develop Standards that Lead to Location and Media Independence, e.g., Transmission Standards for Color T.V.
- Standard VLSI Communications Controllers: Driving the Economics of Communications from "Few to One" to "Any to Any"

## VLSI IMPACT ON COMMUNICATIONS



## ETHERNET: A FIRST STEP TOWARD STANDARDIZED HIGH-SPEED COMMUNICATIONS WITHIN A BUILDING

### DIGITAL/INTEL/XEROX COOPERATION

- Develop a High-Speed Local Area Network  
—Thorough Correctness Proof
- Publish the Specification for General Widespread  
Implementation

**GOAL: Interoperability Within a Local Area Network**

intel®

### INTEL CONTRIBUTION

- Current LSI Communications Peripherals Support Standard Protocols—HDLC/SDLC, BISYNCH, etc.
- Developing a VLSI LAN Controller Which Supports the ETHERNET Specification
- Develop System Level Products Utilizing ETHERNET

intel®

### SUMMARY

- The Impact of VLSI on Computing has been Significant
- There will be a Similar Impact of VLSI on Communications
- An Orientation Toward Global Optimization will be Central to Achieving the "Wired Community"

intel®

**Philip L. Arst**  
Intel Corporation  
Santa Clara, California

Philip L. Arst is manager of the Data Communications Product Line at Intel Corporation, Santa Clara, California. He is responsible for setting corporate direction for the company's data communications activities, as well as, engineering and marketing responsibilities for a line of data communications systems products. Mr. Arst is also responsible for managing Intel's Ethernet program and has led this program since its inception. This responsibility includes work on the Ethernet specification, interfacing to standards bodies and bringing a family of products to market.

Mr. Arst has an extensive background in the data communications field, including work on a series of projects which resulted in one of the earliest commercial implementations of a local area network (LAN).

As the data communications product manager in the Data Systems Division of Xerox Corporation, El Segundo, California, Mr. Arst developed a front-end communications processor for their mainframe computer product line. And, as the data communications product planner at the company's Integrated Office Systems Division, El Segundo, California, he was responsible for the formulation of strategies and products incorporating Ethernet and global data communications capabilities for integrated office systems.

While associated with the Collins Radio Company, Newport Beach, California, Mr. Arst was active in the system and software designs of early message switching systems. And, this is where he was involved in one of the earliest commercial implementations of an LAN, the Collins C-System TDM Loop.

Mr. Arst has received a BSEE degree from the U.S. Naval Academy and an MBA from the University of Chicago.

## **The Changing Economics for Computer and Terminal Interconnection**

The combination of VLSI components and the inherent systems cost reductions provided by the Ethernet architecture will materially lower the costs of data interconnection and switching within localized geographic areas such as buildings, factories or laboratories. This approach offers such significant advantages over today's telecommunications based techniques, that we foresee Local Area Networks being installed in all business establishments and the LAN interface component becoming a standard part within personal computers and workstations destined for use in the business establishment.

The VLSI design process is typically a three year program. Current estimates are that VLSI devices for Ethernet, the simplest of today's crop of LAN protocols, will be approximately 50% more complex than Intel's 8086 16-bit microprocessor. Bringing this device to market therefore represents a formidable design and product challenge for the semiconductor manufacturer. The stability and simplicity of the Ethernet protocol makes this practical.

Current estimates are that this controller component will implement the full Ethernet protocol (i.e., the entire Bluebook) with the exception of the physical link (transceivers and cabling). In this manner, not only will 80-100 ic's and a full circuit board of today's implementation be replaced, but also the user will be freed from any programming at the Ethernet data-link level as the component manufacturer will have done it all for him/her. We therefore foresee the electronics cost of the Ethernet data-link dropping to a \$30-\$40 level by 1985 if LAN demand provides for the production volumes we believe they will.

However, the electronics interface is only the tip of the cost iceberg. While it is incorrect to compare today's telecommunications based solutions to local networks (because the local network provides needed high bandwidth data services for computer to computer communications which are beyond the capability of the digital PABX) it is still useful to examine the relative cost components of each.

In today's telecommunications based systems, the Electronics Interface is cheap, but modems, dedicated ports on PABX equipment (at \$500-\$1000 per port) and front end or message switching computers are also required to transmit, route and distribute data between distributed and centralized data processors and user workstations. Rewiring and reconfiguration are also an important portion of today's cost equation as they are often required to accommodate change and growth.

The Ethernet bus architecture eliminates the cost of the switching function provided by the centralized PABX unit or the front end/message switching processor by building a distributed switching capability into the controller electronics of each workstation. This is accomplished by interconnecting all processors and workstations on a single shared channel. In this manner, each receives the traffic of all other stations on the net and selects only traffic which bears its address.

An additional unique capability of the Ethernet Architecture is its transceiver design which permits easy reconfiguration. This permits it to avoid expensive rewiring and switching equipment reconfiguration when needs change or equipment is relocated. However, the transceiver design of Ethernet has its drawbacks as these devices are currently expensive (approximately \$300 in small lots). Fortunately, the transceiver is also susceptible to considerable cost reduction. The first step will be the integration of its electronics into a single or a few chips. But since the major cost of the transceiver is in its mechanical parts (i.e., housing, connectors, separate circuit board and power supply), a systems approach can be taken to lower this class of costs.

These systems approaches are typically based upon sharing a single transceiver between many stations. Products of this category consist of:

- Transceiver multiplexers which permit the sharing of a single transceiver by 4, 8 or more stations (or its elimination entirely in small systems).
- RS-232c interfaces for multiple "dumb" terminals which share a single transceiver and set of electronics (such as the Ungerman-Bass Network Interface Unit).
- Packaging the transceiver electronics (i.e., chip) within the workstations and bringing a flexible version of the Ethernet cable to a tap on the cabinet. By clustering these "cables to the cabinet" terminals and then interfacing them to the main Ethernet cable via a simple repeater, significant cost reduction can again be achieved. Through utilization of these techniques, we foresee an Ethernet interface consisting of a VLSI controller and a separate transceiver selling in volume OEM quantities in the \$120-\$150 per node range in 1985 and in the cable to the cabinet configuration of \$30-\$50 per node.

A further cost reduction of the Ethernet VLSI component will be obtained by applying its basic CSMA/CD (Carrier Sense Multiple Access with Collision Detection) technology to other applications. For example, we foresee CSMA/CD LANs being built within cabinets of electronics, such as a personal workstation. The CSMA/CD LAN would interconnect the station processor, its floppy disk, printer and other devices. Intel products will support these non-Ethernet applications, thereby further building product volumes and lowering Ethernet costs.

These cost levels, plus the higher functionality provided by the Ethernet architecture, will, in our opinion, make the Ethernet controller the computer terminal interface of the 1980s.

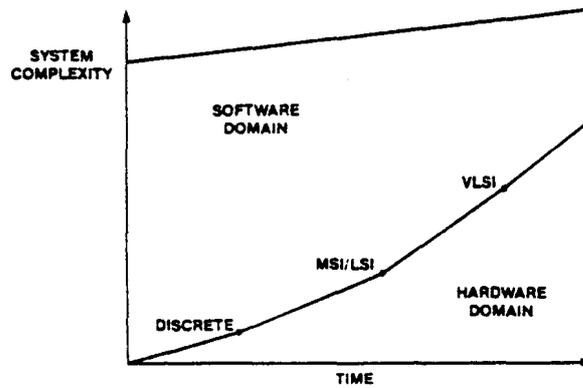
intel®

## "THE CHANGING ECONOMICS FOR COMPUTER AND TERMINAL INTERCONNECTION"

Philip L. Arst  
Intel Corporation

intel®

## THE DATA COMMUNICATIONS CHALLENGE A Systems Solution



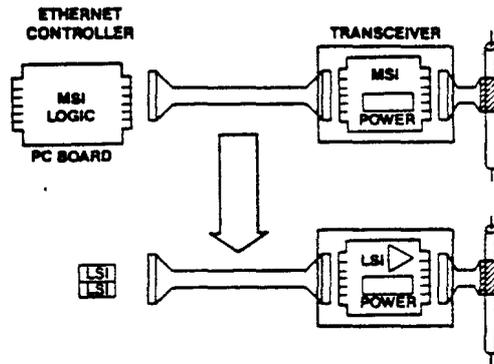
intel®

## DESIGN TRENDS

Date	Device	Part Type	Device Complexity	Design Cycle-time (Years)
1976	UART	8251	2,800	1.5
1979	USART	8274	24,000	2
1979	16-Bit Microcomputer	8086	29,000	3
1981	Ethernet Local Network	—	43,000	3

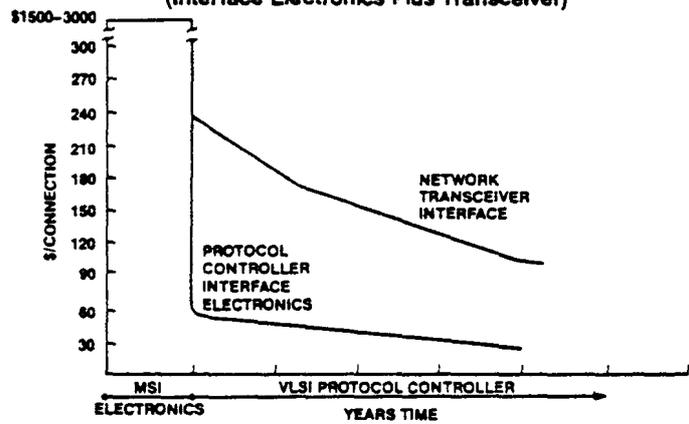
intel®

### VLSI COST REDUCTION



intel®

### ETHERNET INTERFACE COST TRENDS (Interface Electronics Plus Transceiver)



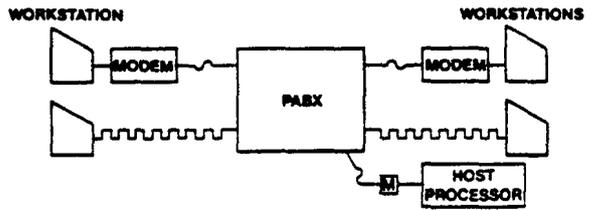
intel®

### TELECOMMUNICATIONS-BASED WIRING AND SWITCHING (Con't)

- Today's Approach
  - Full Voice Service (Blocking)
  - \$500 to \$1000 Per Port
  - High Initial Cost
  - Modems Required for Data
- Voice Data Approach
  - Full Voice/Medium Speed Data (Non-blocking)
  - Generally Requires Rewiring
  - High Initial Cost
  - Cost Per Port?

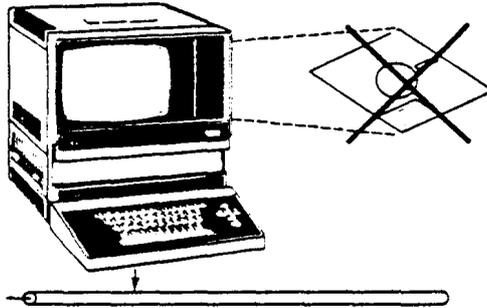
intel®

### TELECOMMUNICATIONS-BASED WIRING AND SWITCHING



intel®

### TOMORROW'S PERSONAL WORKSTATION

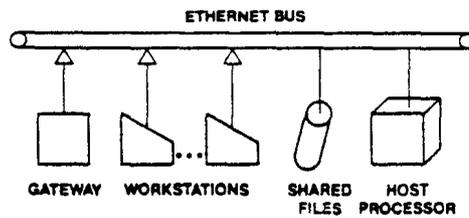


All Electronic Workstation  
With a High-speed  
Ethernet Interface

- ☐ Shared Resources
  - Files
  - Data Output
  - Data Input

intel®

### LOCAL AREA NETWORK-BASED WIRING AND SWITCHING



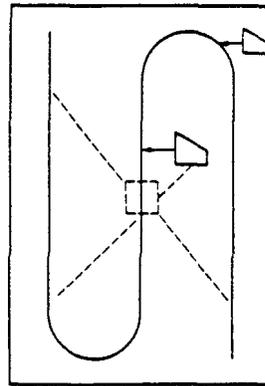
intel®

### ETHERNET APPROACH

- Switching on Bus
- Modem Function in Transceiver
- High Bandwidth
- Short Wiring Runs
- Relocatable Bus Connectors
- Coexistence With Existing Voice Systems

intel®

### ETHERNET CONFIGURABILITY



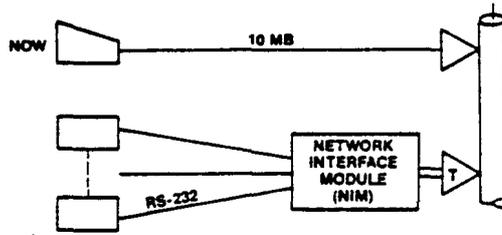
intel®

### ETHERNET INTERFACE COST ASSUMPTIONS

- Estimated OEM Average Selling Price
- Interface Built into Terminal
- Shares Microprocessor and Memory

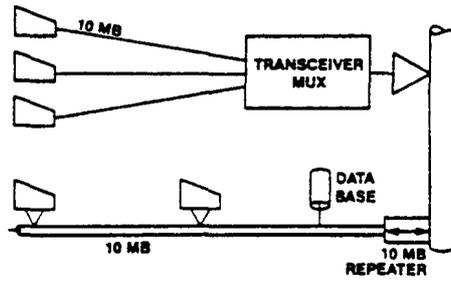
intel®

### SYSTEMS CONFIGURATION ALTERNATIVES



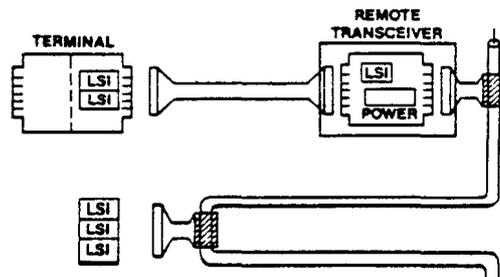
intel®

### SYSTEMS CONFIGURATION ALTERNATIVES (Con't)



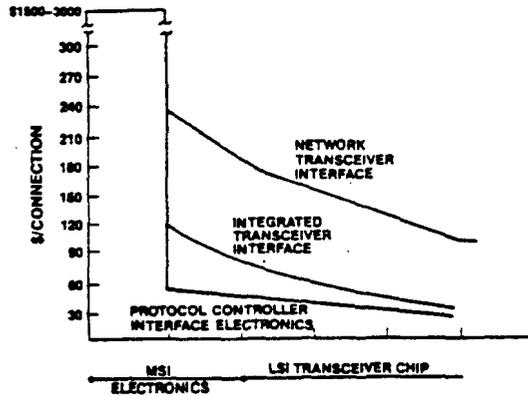
intel®

### CONFIGURATION ALTERNATIVES YIELDS COST REDUCTION



intel®

### ETHERNET INTERFACE COST TRENDS (Interface Electronics Plus Transceiver)



intel®

### ETHERNET—THE COSTS TUMBLE

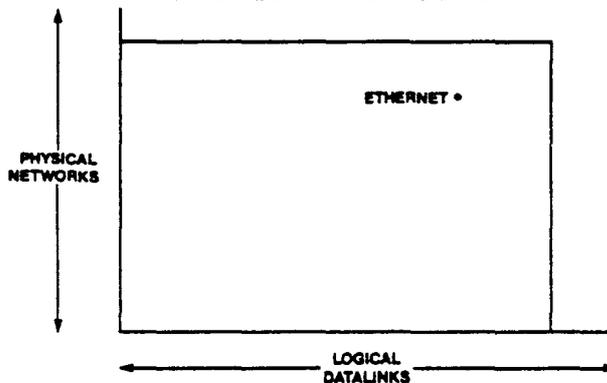
- Simple Field-proven Protocol—Suitable for VLSI
- High Bandwidth and Configurability Enable a New Generation of Applications
- Ethernet Cost Projections
  - Improves Upon Today's Telecommunications-based Costs
  - VLSI to Drop Costs Further
  - Systems Configuration Tradeoff's Used to Go to Rock Bottom

intel®

CHANGED  
"THE ~~CHANGING~~ ECONOMICS FOR COMPUTER  
AND TERMINAL INTERCONNECTION"

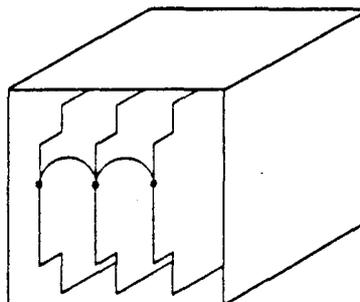
intel®

### CARRIER SENSE MULTIPLE ACCESS WITH COLLISION DETECTION



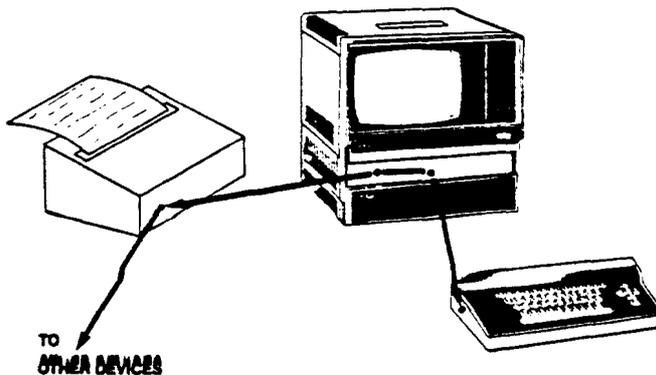
intel®

### SERIAL BACKPLANE WITH CSMA/CD



intel®

### INTRACABINET INTERCONNECT WITH CSMA/CD



intel®

### **CARRIER SENSE MULTIPLE ACCESS WITH COLLISION DETECTION**

- **Just Like Verbal Conversations**
  - Shared Channel
  - Listen Before Transmit
  - Transmit Only When No Carrier Present
  - If Collision, Stop and Wait a Random Time
  - Retransmit if No Carrier Present

intel®

### **INTEL'S VLSI EFFORT IN LANs**

#### **INTEREQUIPMENT COMMUNICATIONS**

- **<\$50 per ETHERNET Connection**
  - Interoperability
  - High Performance

#### **INTRACABINET COMMUNICATIONS**

- **<\$25 per ETHERNET Subnet**
  - Same VLSI
  - Performance/Cost Tradeoffs

# **DIGITAL AND ETHERNET**

**Gordon Bell  
William R. Hawe**



## BIOGRAPHICAL INFORMATION

### C. GORDON BELL

Vice President, Engineering

C. GORDON BELL, 47, is Vice President, Engineering for Digital Equipment Corporation.

In this position, Bell has responsibility for the company's research, design and development activities in computer hardware, software, and systems and is a member of the Operations Committee, Digital's 13-member senior management team.

Bell joined Digital in 1960 as Manager of Computer Design, a position he held for six years. He took a leave of absence from Digital in 1966 to join the faculty of Carnegie Mellon University in Pittsburgh. He rejoined the company in 1972 as Vice President of Engineering.

Prior to joining Digital, Bell held several engineering positions including that of research engineer at the MIT Speech Communications and Electronic Systems Laboratories.

Bell earned his B.S. and M.S. degrees in Electrical Engineering at Massachusetts Institute of Technology in 1956 and 1957 respectively.

He is a widely published author on computer architecture, and computer design. His books include the McGraw-Hill book, "Computer Structure," co-authored with Allen Newell; and the

Biographical Information  
C. GORDON BELL  
Page 2

Digital Press books, "Designing Computers and Digital Systems, Using PDP-16 Register Transfer Modules," with John Grason and Allen Newell; and "Computer Engineering: A DEC View of Hardware Systems Design," co-authored with J. Craig Mudge and John McNamara.

Holder of several patents in the computer and logical design areas, Bell has also served the U.S. Government as a member of COSINE Committees of the National Academy of Science for computer engineering education, and the National Science Foundation, Office of Computer Activities and Computer Science and Engineering Researching Study (COSERS). Bell has also served on the Council for International Exchange of Scholars.

Among his professional affiliations, Bell includes the National Academy of Engineering, Fellow of the Institute of Electrical and Electronic Engineers, member of the National Research Council's Computer Science and Technology Board, fellow of the American Association for the Advancement of Science, the Association for Computing Machinery, and is listed in "American Men of Science" and "Who's Who."

He is a resident of Lincoln, Massachusetts.

#####

January 1982

# **WHY DIGITAL IS COMMITTED TO ETHERNET FOR THE FIFTH GENERATION**

**A summary of remarks made by Gordon Bell, Vice President—  
Engineering, Digital Equipment Corporation, at the  
Xerox/Intel/Digital Seminar.**

**New York, February 10, 1982**

# WHY DIGITAL IS COMMITTED TO ETHERNET FOR THE FIFTH GENERATION

## ABSTRACT

What the Digital Unibus did for minicomputers, Ethernet will do for the Fifth Generation.

Nearly all recent computers are organized around a single, high speed bus (Unibus-type structure) which provides communications among its processors, memory, disks, and interfaces to the external environment. This simple structure has been one factor in the rapid evolution and proliferation of computers. Unfortunately, a bus for interconnecting computer components within a cabinet, is not suitable for interconnecting a network of computers within a building.

Ethernet is a high speed, 10 megabits per second, standard bus providing the first two levels of the ISO Open Systems Architecture. It permits the dynamic connection of computers at a site to form a local-area network (LAN) in an open-ended fashion without the need of centralized equipment or planning and control. In the Fifth Generation, the network becomes the system and Ethernet is a key prerequisite of the generation.

Ethernet will be used initially, in an evolutionary fashion, to interconnect networks of today's computers to each other and to terminals and personal computers. Since Ethernet is a factor of 1000 higher speed standard than today's network links, and easily used to form networks, we expect a rapid transition to a tightly integrated network, where the network is the system. In this generation, separate function computers (eg. personal workstation, file server, print server, real time, timeshared) will be tightly integrated, interchanging many types of messages, such as, files, computed graphics, pictures, and voice. This kind of network will permit a radically different use of computers, and only then can we be certain that this is the Fifth Computer Generation.

Because Ethernet is so important to the Fifth Generation, Digital is committed to it as a standard. We use these networks and will be providing products in the near future.

**Ethernet Is The Unibus  
Of The Fifth Generation**

**“You have to look at  
Ethernet as a standard...”**

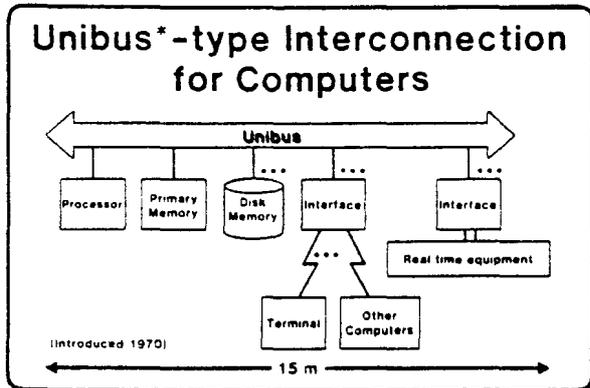
Ethernet is one of the keys to the development of the Fifth Generation because it provides a standard for the interconnection of all sizes and types of computers in a passive, local-area network.

Up until now, interconnection has been a very difficult task simply because there has been no standard.

A standard is a blueprint that shows you how to build the components that will go into a system or onto a network.

**“... system components are connected by a single high-speed bus in an open-ended fashion.”**

If you look at current computer architecture you will find standards. One such standard is the Digital Unibus that defines the architecture used in the largest selling series of minicomputers ever built—the PDP-11 series. The Unibus standard made it possible for our users and a number of different manufacturers to build memory boards, communications interfaces, and other components that can be plugged directly in a PDP-11 system in an open-ended fashion.

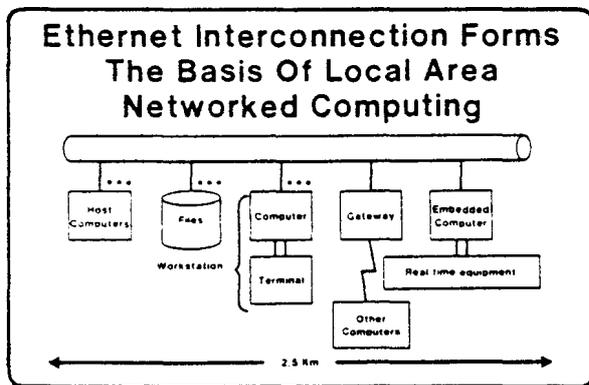


If you look at the Unibus-type architecture, or any competitive implementations of the Unibus idea such as Intel's Multibus or Motorola's Versabus, you will find that all system components—processors, system memory, data storage, and data communications interfaces—are connected by a single, high-speed data path or bus.

This bus enables the computer to move data within the system at very high speeds. Unfortunately there has been no standard bus to move data between systems at the similar speeds. Ethernet communications won't replace Unibus or any competitive busses but Ethernet will solve the local-area networking problem.

**“Think of Ethernet as an extended bus ...”**

Ethernet is an extended bus. Up until now busses have provided high-speed computer communications within a very limited area—a single cabinet or room. Ethernet provides an extended bus that will link information processing nodes throughout a building, campus, or industrial complex.



The system components don't change. You have the same components in an Ethernet as you have in a single system. The only difference is that you now have more components and they're dispersed over a wider area. Where a Unibus system has a single processor, an Ethernet can have many.

Where a Unibus system has local data storage, an Ethernet will support databases distributed throughout the network.

Where a Unibus system interfaces to other computers, an Ethernet interfaces to other networks through gateways.

**“The network becomes the system....”**

In other words, with Ethernet, the network becomes the system. And when this happens, we will have a whole new computer generation—The Fifth Generation.

We—that is Digital Equipment Corporation—want to be a leader in the development of this new generation just as we were the leader in interactive computing and the development of the minicomputer generation that made distributed data processing possible.

Let me take a minute to define what I mean by a computer generation.

A new generation of computers comes about when there is a convergence of Technology and Need that forms a new Structure that is then followed by general Use.

With Ethernet and VLSI—Very Large Scale Integrated Circuits—we have the technology. That technology is needed to build and network an ever-growing number of computers, terminals, intelligent workstations, and personal computers that are being bought to solve many of the productivity problems facing business today.

There is also a new structure, the local-area network. Just as minicomputers and distributed processing changed the way computers were used in the 70s, local-area networks and personal workstations will change the way computers are used in the 80s.

The final requirement for a new computer generation is customer acceptance. Will the new technology and the new structure come into general use? In this particular case, I am convinced it will. Just as I'm convinced that Ethernet is the technology that will make this happen.

Ethernet provides the simplicity, speed, and universality needed in local-area networking.

Unlike other local-area networks, Ethernet is open-ended. It allows the user to build a local-area network from the bottom up without making a large capital investment or developing an inflexible long-range plan.

As I mentioned earlier, Ethernet is a passive communications medium. An Ethernet is really nothing more than a coaxial cable and standard protocols that define the way data is transmitted. For example, the Ethernet protocol defines packet size. It defines the way packets are addressed. It's really very simple. And it's been tested for 10 years and it works.

Ethernet can carry a great deal of information at very high speeds. But you don't have to take my word for it. I'd like to read you part of an advertisement written by another computer manufacturer who adopted the Ethernet standard.

**"Instead of taking . . . 44 seconds to transmit 10 pages of data, the transfer takes place in .042 second. In the 4.4 seconds it would take a conventional network to send one page of War and Peace . . . you could send the entire thousand page novel."**

It is not difficult to see the benefits. You can transmit entire files from a computer to a personal workstation almost instantaneously. You can transmit photographs, data sheets, engineering drawings, or even voice messages.

The key is universality. Any manufacturer who follows the Ethernet standard can build equipment to go onto the network.

But it is important that we realize that a network is more than just lines and nodes. Higher level protocols are needed to support the interconnection of dissimilar computers; to implement complex network functions such as file transfers and terminal-to-terminal communications; and to provide network management capabilities.

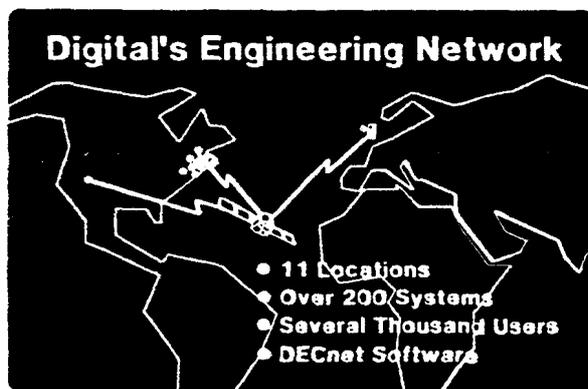
These protocols are complex. But they are a prerequisite for building a network such as the one that serves Central Engineering at Digital. One of the reasons we are committed to Ethernet is that it fits into the framework defined by Digital Network Architecture. We don't have to change the higher level protocols that are being used to support tens of thousands of DECnet nodes around the world. We can make Ethernet part of DECnet. We have a fit. And we have the range of capabilities required to implement complex computer networks.

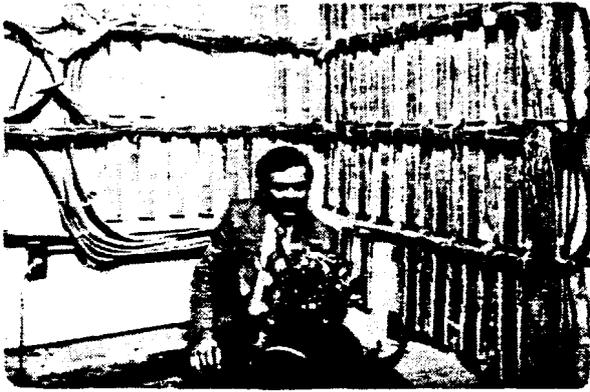
Let's look at an example.

Digital's Engineering Network is made up of over 200 different systems serving about eight thousand terminal users. But interestingly enough, 80% of the traffic on this network is local traffic—only 20% of the traffic is between locations.

Local-area networking addresses the local problem. It provides high speed interconnection among computers within the same building or complex, and it simplifies the interconnection of terminals and processors to host computers.

**"Higher-level protocols are needed . . ."**





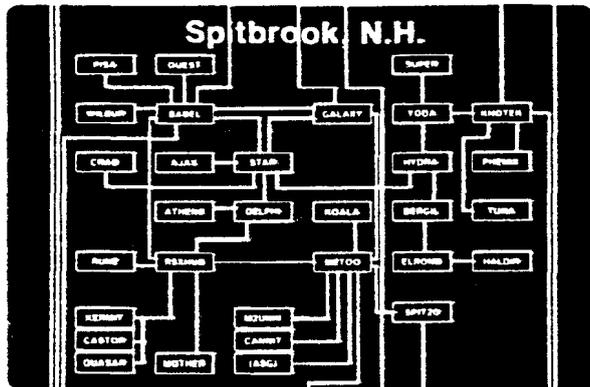
**“Ethernet can eliminate this complexity while providing the flexibility needed for future growth.”**

Let's look at the computer-to-computer and terminal-to-computer interconnection problem in a little detail.

The problem is a wiring problem. It is one thing to connect A to B; quite another thing to connect A to B through Z. Before you know it you have a very complex maze of wires and switches. This is the wire room in our Spitbrook, New Hampshire facility. As you can see, interconnecting a large number of devices is—at best—a very messy and, I might add, very expensive, business. You have fixed wires running all over the place. It's difficult to add systems or make changes.

Ethernet can eliminate this mess and provide needed flexibility. Ethernet will let us replace all this wiring with a single coaxial cable that will run throughout the building. When we want to add a terminal we'll just tap into the cable. It won't be necessary to run wires back to a central location. And we'll be able to add terminals to the network without interrupting network operations.

But we—like most other large organizations—are starting to provide individual users with intelligent workstations or personal computers rather than simple terminals. A simple terminal is usually a low speed device that can operate over telephone-type wiring. After all I can only read and write just so fast. I can type 50 words a minute. I can read about 200 words a minute. 9,600 bit per second transmission is more than fast enough for me as long as I only have a simple terminal. But when I have an intelligent terminal that can deal with information a lot faster than I can, I need to be able to communicate at computer speeds. Ethernet provides the speed needed to support intelligent user devices. The speed needed to transfer entire files or complex graphic images in a fraction of a second. I need Ethernet communications.

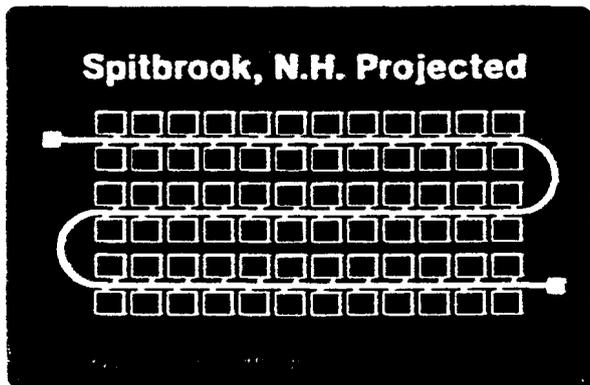


**“This is how our computers are connected today.”**

At the same time Ethernet solves the problem of interconnecting computer systems.

This is how the computers at Digital’s Spitbrook, New Hampshire facility are connected today. As you can see messages have to be routed through the network. This creates computer overhead. Many systems spend much of their time switching and forwarding messages. And as more and more systems are added to the network this overhead just keeps growing and growing.

Fortunately, Ethernet can eliminate the overhead problem because it eliminates message switching and forwarding. This is how Spitbrook will look when we install an Ethernet.



**“Ethernet eliminates message switching.”**

As you see each system is connected directly to the Ethernet. There is no message switching. No routing. No forwarding. No computer overhead.

Instead of a maze of wires you have a high-speed, high-capacity extended bus that serves the entire complex. As you see Ethernet is changing the very definition of a system. With Ethernet, the network becomes the system.

We have a new technology. A pressing user need. And, a new structure. Three of the four prerequisites for a new computer generation. The fourth requirement is use. There are currently about 100 Ethernets in operation. There are going to be thousands. We’ve already talked to our customers. We know what they want and we know that many of them are going to install Ethernets. That’s why I believe that we’re looking at a new computer generation.

**“Within the next few months  
we will be introducing our  
first Ethernet products.”**

We're going to build that generation. That's why we joined with Xerox and Intel to develop the Ethernet Specification. That specification conforms to both The Open Systems Architecture proposed by the International Standards Organization and Digital Network Architecture used in thousands of networks around the world. Right now we are implementing Ethernet as a part of Digital Network Architecture and within the next few months we will be announcing our Ethernet program and introducing our first Ethernet products.

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.

# **PERFORMANCE OF A SIMULATED ETHERNET ENVIRONMENT**

**A summary of remarks made by William R. Hawe, Principal Engineer  
in the Systems Performance Analysis Group, Digital Equipment  
Corporation, at the Xerox/Intel/Digital Seminar.**

**New York, February 10, 1982**



## BIOGRAPHICAL INFORMATION

### WILLIAM R. HAWE

William R. Hawe is a Principal Engineer in the Systems Performance Analysis Group at Digital Equipment Corporation. He is involved in the modeling and performance analysis of networks and distributed systems. His interests are in local area network architectures and the performance of higher level protocols. He is a member of the Local Area Networks Performance Working Group of the IEEE Project 802 Local Area Networks Standards Committee. Prior to joining Digital he was a member of the faculty at Southeastern Massachusetts University where he performed research in X.25 packet switching for personal computers.

He received a BS and a MS in Electrical Engineering from Southeastern Massachusetts University. There he was president of the local chapter of IEEE and vice-president of Eta Kappa Nu. He received SMU's IEEE Outstanding Electrical Engineer award while he was research director of the Murail project. This research involved the development of distributed collision avoidance and dynamic routing systems for automated train systems.

SLIDE 1

In this study we investigate the performance of a simulated Ethernet environment. The goal is to predict the capacity of the channel in terms of the number of active users that it can support simultaneously. This provides an understanding of the loading one could expect in a particular environment. It also establishes the capacity in the system for future growth.

SLIDE 2

The goals were to establish the traffic patterns in the existing system and to estimate the excess capacity that would allow growth. The traffic patterns were established through measurements performed on operational systems that were interconnected with conventional point-to-point connections. We wish to see how heavily loaded an Ethernet would be if installed as an interconnect mechanism for the hosts, terminals, etc. We were also interested in understanding the additional loading that would take place because of new devices and their use (print and file servers, etc.) along with increased load due to growth in the user population.

The behavior of users during various periods (such as a busy period) were monitored. The resultant data was then analyzed to produce a profile of the "typical" operations a user performs. From this, a workload which specifies the operations performed (and their frequency) was developed. This includes items such as the rates and sizes of commands, data, etc. that are exchanged between the user and the system.

To predict the growth capability present in the system, we simulated the Ethernet using a distributed architecture model and the user workload as the source of traffic. The number of users was then increased until the idle time on the Ethernet channel went to zero.

SLIDE 3

The results indicate that the Ethernet has sufficient bandwidth to support a large number of users of the type characterized in this environment. The delays in the Ethernet level of the architecture are small compared to other delays such as disk seeks, application program execution, etc. We also see that there are few collisions, even under heavy load.



SLIDE 4

In systems such as backbone networks, the delays in transferring information from node to node are usually dominated by the transmission and propagation delays. Processing time per message at the nodes is small compared to these factors. With the advent of local area networks we see a different relationship. Local area networks are generally built using interconnection mechanisms that have speeds of around 1 to 10 Mbps. They are generally confined to a limited geographic area such as few buildings. This means that now the transmission and propagation delays are much smaller in relation to the disk and CPU delays. For this reason, it becomes important to consider all levels in the system when evaluating the performance.

SLIDES 5,6,7

There are three parts to the study. First, measurements were performed to characterize the behavior of users in a program development environment. From this, a user profile (or workload) was developed. Second, the user workload is used as input to a model of the distributed architecture that is used in the Ethernet network. This results in a traffic load placed on the Ethernet. Finally, this load serves as input to a detailed Ethernet simulation. The number of users using the system in the simulation is then increased to observe the effects of increased load. It is assumed that enough hosts, terminals, etc. will be added to the system to support those additional users.

SLIDE 8

Here we are interested in the capacity of the system. There are many ways that one can investigate this aspect of the performance. Often the capacity of a channel is expressed in bits per second or percentage of the bandwidth used on the channel. Metrics such as this are difficult to interpret when one is interested in estimating how many users the system can support.

SLIDE 9

Therefore, in order to understand the capacity of the system we focus on the number of users that it can support. This is especially important when one is interested in determining whether or not there is sufficient capacity in the channel to support the existing user population as well as reserve capacity for future expansion both in the number of users and the types of traffic they generate.





SLIDE 10

There are two ways in which the environment affects the number of users that the system supports. First, it dictates the higher level protocols to be used to transfer information between hosts, terminals, etc. This in turn affects the amount of traffic generated by each user. Second, it specifies the packet size distribution and arrival rate distribution. These play a significant role in determining the performance of the Ethernet.

SLIDE 11

Here we investigate the program development environment. Measurements were performed at several locations which were considered to be representative of this environment. As an example of such an environment we consider a large University. Users in this environment perform the obvious activities associated with the development of programs. This includes editing files, as well as compiling, linking, running and debugging the programs. They also communicate with other users by sending mail and using interactive message facilities such as "Talk". They copy, delete, print and perform other file manipulation operations. In addition to these functions, they also obtain information from the system. This includes help messages, queries about system status, etc.

SLIDE 12

To characterize the activities of a typical user various parts of the system must be monitored. Data was collected at several installations representing this environment. The data was collected at various times during the day so that busy periods could be investigated.

The amount and frequency of information transfer between the terminal and the host was monitored. In addition, the disk I/O that occurs as a result of operations performed by the user was also measured. This includes disk I/O that is for temporary work files such as those generated by programs such as linkers and compilers. Note that when we examine the impact of sending disk I/O over the Ethernet to a file server we do not include this type of traffic. This is because it is more efficient to generate and manipulate those temporary files at the location that the linker or compiler is running. However, the source and destination files can certainly be located on a file server. We also monitored other forms of traffic resulting from user operations. These included CPU usage, printing, network I/O, etc.



SLIDE 13

As we mentioned, the current environment uses conventional methods for interconnecting hosts, terminals and other devices. Terminals are connected directly to the hosts. The hosts are interconnected using point-to-point connections. The network is not always fully connected. However, the routing capabilities of the hosts assure that the network is logically fully connected.

SLIDE 14

We consider the impact of an Ethernet installation in this environment. The Ethernet will carry traffic between the hosts for remote file and data access, remote logins, printing, etc. It will also carry traffic to and from new devices such as file servers and print servers. Existing terminals which are connected directly to the hosts can access remote hosts, servers, etc. by going through their hosts. Other terminals can also be connected to the Ethernet either directly (with the appropriate interface) or through terminal concentrators. With this approach they are not dependent on any one host's availability for access to the network. Personal computer workstations can also be connected directly to the Ethernet. Their traffic will be somewhat different than the terminal traffic because of the increased intelligence in the workstation. It will appear more like the host to host and host to server traffic. Gateways, routers, and other devices which allow communication outside of the local area network may also be connected directly to the Ethernet. Often hosts implement these functions in addition to their normal duties. The traffic which flows through those devices can be of any of the types already described.

SLIDE 15

As time passes the network will expand in several ways. More devices will be added as the user population increases. This includes terminals, concentrators, hosts, servers, etc. Hosts without local terminals could be added and called computing servers. The other way in which the network will expand is in the traffic patterns. The availability of devices such as file and print servers will stimulate the growth in the traffic associated with those devices. For example, as more files are moved to file servers, so that sharing is easier, the devices will be used more often.



SLIDE 16

We have discussed the users and their environment. Now we discuss the distributed architecture. A distributed architecture is necessary to provide an effective local area network. There must be facilities for reliable, controlled communications between users and processes inside and outside the local area network. This means that we need mechanisms for a user on the local area network to access information not only on the local network but also at some location that is not local. This would be accomplished by going through a gateway or router. Therefore, Ethernet is only a part of the total network architecture. It represents the lowest layers and is thus the foundation on which the local network is built.

SLIDE 17

The Digital Network Architecture (DNA) is an example of a complete network architecture. Here we see the relationship between DNA and the ISO layered architecture. The Ethernet comprises essentially the lower two levels for the local area network. Parts of the system which interface to public data networks could use the X.25 services. Other point-to-point links could use the DDCMP facilities. Above the data link is a network wide routing service. This delivers packets to the appropriate destination - either locally or remotely. Above that is an end-to-end service which provides for reliable communications between two processes. The Session layer controls the end-to-end service. Above that we have the applications and special purpose protocols. The network management facility has access to most of the protocol levels. It is used to monitor as well as control and configure them.

It is very important that all these layers in the architecture be considered when examining the user perceived performance of the local network. This is because each layer will add some additional load to the components of the system. Most will add some amount of additional traffic to the Ethernet. They will also use resources such as CPU cycles and memory space.

SLIDE 18

As we said, in DNA the Ethernet implements the physical and data link layers of the network architecture for the local area network. It offers a datagram service with delivery of packets on a "best effort" basis. In that sense it is different than other data link protocols such as DDCMP. The channel is, in general, relatively



error free so this protocol is a good match. The Ethernet uses the CSMA/CD protocol to share the 10 Mbps channel. It uses a distributed algorithm called binary exponential backoff to resolve contention for the channel. This algorithm is executed independently by each station and is fair to all. The specifications allow a maximum of 1024 stations or "taps" on the Ethernet cable. However, as we shall see, there can be more users than taps. This is true of terminal concentrators where several user terminals may share a single tap. Hosts may also have a single tap as well as several users or processes that are generating Ethernet traffic.

SLIDE 19

The DNA Transport protocol implements the network wide routing layer of the network. (This includes the local network as well as components that are not connected locally to the Ethernet.) This layer corresponds to essentially the ISO Network layer. It provides end-to-end routing of datagrams and routes packets to a destination even if the node is not on the Ethernet. To do this, it supports a network wide node address space. A node's address can be the same as its Ethernet address if it is on the Ethernet. However, all nodes are not necessarily connected to an Ethernet. Therefore, we need this address space. This layer also prevents congestion within the network and provides dynamic routing to bypass sections of the network that may have failed for one reason or another.

SLIDE 20

In order to provide effective, error free, and reliable process-to-process communication an end-to-end service is required. This is implemented by DNA's Network Services Protocol (NSP). NSP uses a virtual circuit to provide these features. This assures that packets are delivered to the user in the order they were sent. NSP makes sure that none are lost in the network. This is done by retransmitting lost packets. The timers used to decide when to retransmit a packet are self adjusting. This means that they adjust to the delays in the channel. This has the advantage of limiting the amount of unnecessary retransmissions thus reducing the load on the channel. The protocol also provides various flow control options. This allows the characteristics of the circuit to be tailored to the application. For instance, some applications may require tight control on the rates at which information is exchanged. These data rates impact the amount of resources (buffers, etc.) that must be devoted to the circuit. Flow control is



especially important when the receiver is slower than the sender. An example is host to terminal output where the host can usually output data at a rate much faster than the terminal (or user) are capable (or willing) of accepting.

SLIDE 21

The DNA Session Control layer is used to control the virtual circuit service that NSP implements. It allows users to set up and terminate circuits. It validates incoming connect requests and activates the appropriate processes for those that are valid. It manages the interface between the user applications and the circuit. It also provides name to address mapping. For example, if the user requests that a circuit be established to a node having a particular name, this layer determines the address of that node so that the connect request packet can be sent to the proper destination.

SLIDE 22

Above the Session Control layer are the applications protocols. The DNA Data Access Protocol (DAP) is one such protocol. It provides remote file access services. This means that the user can use this facility to access files as if they were stored locally on his system. The operation of the network is completely transparent. Another example of an application protocol is a virtual terminal protocol. This allows the user to connect to remote hosts through the network. The user then appears to be connected locally to that remote system.

The network management part of the architecture is used to monitor and control the various protocol layers. It can be used by the network manager to monitor the traffic in the network and thus is useful for capacity planning. It is also used to tune the network for better performance.

SLIDE 23

We use the user workload as an input to the distributed architecture model. The output of this is a load on the Ethernet. This consists of the user information being transferred between points on the network as well as various control and data packets associated with the protocol layers in the distributed architecture. The Ethernet simulation simulates the transmission of these packets.



SLIDE 24

To understand the behavior of the channel there are several metrics one can examine. The delay experienced in transferring a packet between stations is of obvious interest. The number of retries necessary to accomplish that transfer is also important. Retries occur whenever there is a collision between two or more packets. The specifications indicate that after 15 retries (ie: 16 attempts) the packet will be aborted. At that point, the higher layer protocols must retransmit that packet. In this case it is NSP that will do the retransmission. The number of retries then gives us an indication of how the channel is behaving.

SLIDE 25

The values of the performance metrics such as delay, retries, etc. are determined by variables that come from two general sources. The first are those associated with the Ethernet itself. These are the transmission speed and the propagation delay. Here the transmission speed is 10 Mbps. The propagation delay depends on the size of the network. There is a maximum size that the network can have and therefore the worst case propagation delay is bounded. A transmitter must continue to transmit a packet long enough so that it can propagate to the farthest parts of the network. This way all stations can detect that a packet is being transmitted. However, another station may have started to transmit a packet before the signal from the first one reached it. In that case there is a collision. The collision must propagate back to the sender while it is still transmitting. This way it will know that its packet has been corrupted. The sender must therefore transmit a packet long enough so that it can propagate to the end of the network and any collision can propagate back. This time is called "the slot time" and it is about the round trip propagation delay for the largest network. (The slot time is 51.2 microseconds in the Ethernet specification.)

The other factor which determines the performance is the workload. This is the combination of the user workload and the traffic from the distributed architecture. The packet sizes and the rates at which they arrive for transmission over the Ethernet combine to present an given "offered load" to the Ethernet.

SLIDE 26

Here we see the mean waiting time on the Ethernet as a function of the number of users. The waiting time is the time from when a packet first becomes ready for

transmission until it starts a successful transmission. It includes any time used in deference or collisions. We show three curves based on three levels of remote file traffic. Notice that for up to around 2000 users with this workload, the average waiting time is small when compared to typical delays at disks or in executing application programs or in processing protocol messages.

It is important to remember that the "users" in these curves are active users. This means they are logged in and actively working. Generally, the number of users that are actually using a system at any given time is only a fraction of the total user population. This is true not only for this program development environment but for other environments as well. For example, capacity planning of telephone systems uses knowledge of the relationship between the number of active users and the total user population.

Also note that the system can support more than 1024 users. As mentioned previously, the Ethernet specifications indicate that a maximum of 1024 taps may be connected to the cable. However, we have noted that taps can be shared by several users.

SLIDE 27

The number of retries a packet experiences is another indicator of the channel performance. A retry occurs whenever a packet has been involved in a collision. Here we see the mean number of retries plotted versus the number of active users for the three levels of remote file traffic. Note that for large numbers of users the average number of retries is still close to zero.

SLIDE 28

At some point when the number of active users is increased to a large enough number, the idle time in the channel will go to zero. This happens when the resources are all used in successfully transmitting packets and in overhead (such as collisions). Here this is plotted for the three levels of remote file traffic. Generally, one chooses an operating point at a point that allows fluctuation in applied load as well additional growth. We see that the Ethernet has ample room for growth at this particular installation based on its operating point. In other studies, such as the measurements of the PARC Ethernets, it has also been observed that the loading on the Ethernet in this and other environments is low.



SLIDE 29

It is important that one keep the Ethernet performance data in the proper perspective. Consider a simple example of a file transfer from a file server over the Ethernet to a host or workstation. The "transmission component" includes the actual transmission time of all the packets in addition to the waiting time for each packet. There will be data and control packets from the various distributed architecture layers. The "CPU component" includes the processing time for each packet as well as any application overhead such as that due to the file system and application protocol. This also includes queueing for the CPU that will occur because there are multiple processes sharing that resource. The slower the CPU, the larger this component will be. The "disk component" includes the disk seek delays in addition to the rotational latency and transfer times for the data. It also includes queueing for the disk that occurs because it is shared. Comparing the CPU and disk components to the transmission component, it is not uncommon to observe that the ratio can easily be 4 to 1 or even 20 to 1 or higher - even when the Ethernet is heavily loaded which makes the waiting time longer.

Other scenarios such as terminal I/O have similar relationships. There the disk component may or may not be as large. This depends on how much disk traffic the user generates. Linking and compiling programs, for example, can generate large amounts of disk traffic. The application program overhead in the CPU component can also be large.

SLIDE 30

To summarize, we have seen that the Ethernet is capable of supporting a large number of users of the type characterized in this environment. We have also seen that the delays associated with the Ethernet are typically small when compared against delays other parts of the system. We also note that few collisions are experienced. Therefore, the Ethernet seems well suited for this environment. It has ample capacity and performs well.



# PREDICTED CAPACITY OF ETHERNET IN A UNIVERSITY ENVIRONMENT

Madhav Marathe and Bill Hawe

Digital Equipment Corporation  
Systems Performance Analysis  
1925 Andover Street  
Tewksbury, MA 01876

## ABSTRACT

Local area networks are becoming increasingly popular as mechanisms for interconnecting a broad variety of devices within a moderate geographical area. The Ethernet\* is one of the major access methods currently being used for this purpose. Terminals, hosts, personal computer workstations, gateways, and various types of servers have all found their way onto the Ether. The number of devices that one may attach to the channel is limited by several factors. Finite bandwidth, limitations of the contention resolution algorithm, physical constraints, etc. all impose certain limits. The number of users that may use those stations or "taps" for communication is also limited by these and other factors such as the layered protocol architecture, the physical system architecture, the user workload, etc. Here we examine the limits imposed on the number of users due to the finite bandwidth of the channel. This study is performed for users in a time-sharing environment. Measurements were performed to estimate the characteristics of that environment at a large University currently using conventional direct connections between hosts and terminals. We wish to estimate the limitations on the number of users when the system uses an Ethernet for the interconnection of hosts, terminals, etc. The characteristics of the user environment were coupled with a distributed architecture model and used as input to an Ethernet simulation. The results of the simulation give an upper bound on the number of users which can be supported in this environment. This of course assumes that there are a sufficient number of hosts, etc. so that those resources are not a bottleneck.

## Keywords & Phrases

Ethernet, Ethernet performance, Ethernet simulation, higher level protocols, layered architecture, user level workloads, time-sharing, interactive program development.

\* Ethernet is a trademark of the Xerox Corporation.

## OVERVIEW

### Local Area Networks

Local Area Network interconnection schemes such as the Ethernet provide the framework in which one can construct systems which provide sharing of resources in an effective manner. Two aspects of the Ethernet which help achieve this goal are its speed and the fully-connected nature of its configurations.

To date, no one has come up with a standard definition of local area networks. However, most Local Area Networks do exhibit some general characteristics. Generally, they span areas of up to a few square kilometers. They are often contained completely in one or a small number of buildings. They usually have data rates in the range of 1 to 10 megabits/second. One group or organization almost always has complete control over the operation of the network. Since users are generally from one organization, there is a strong desire to access shared devices such as print servers, file servers, gateways, hosts, databases, etc. As a result, full physical connectivity is desirable. Because of the technology employed and the restricted size of the network, one observes lower bit error rates compared to conventional long-haul networks.

Because of the Local Area Network's speed it usually gets used for not only the traditional network communication but also for handling I/O traffic for shared disks, printers, etc. The personal computer workstations of the future will introduce a new class of traffic on the network. However, in the near future, the traffic on the local area network will consist of host/terminal traffic, host to host file transfers, mail, etc., specialized device traffic (print servers, etc.) and gateway traffic. We make use of this fact in modelling the workload on these networks. More information on Local Area Network technology and architectures can be found in [COTT90] or [PRE90].

## PREDICTED CAPACITY OF ETHERNET IN A UNIVERSITY ENVIRONMENT

### Ethernet

In this paper we are concerned with a Local Area Network built using an Ethernet [DIGI80], [METC76]. Ethernet uses a broadcast mechanism (coaxial cable) and a distributed access procedure to allow for sharing of the channel. The procedure is called Carrier Sense, Multiple Access with Collision Detection (CSMA/CD). Nodes on the Ethernet can sense on-going transmissions and defer theirs until the channel is idle. They also have the ability to monitor the channel while transmitting to determine if any other stations are also attempting to transmit. Once an idle channel is sensed a station may transmit. Because of the propagation delay on the wire, two or more stations may sense an idle channel and attempt to transmit simultaneously. This results in a collision. In order that all stations (including the one transmitting the packet) can "hear" the collision it is required that all packets be greater than a certain minimum size. That size is determined by a parameter called the "slot time". The slot time is slightly greater than the round trip propagation delay. Any station involved in a collision must stop sending the packet and reschedule the transmission. The algorithm used to determine when the next attempt should be made is called the truncated binary exponential backoff algorithm. Basically, every time a station is involved in a collision it backs off (ie: waits) a random amount of time whose mean is doubled every time it experiences a collision. The backoff time is reset after a successful transmission. This algorithm has the advantage of being fair to all nodes on the Ethernet since it is executed by all. Ethernet performance is fairly robust. It degrades slowly and recovers well from momentary overloads [MARA80], [SHOC80].

The day to day operational performance of a 3 Mbps Ethernet is reported in [SHOC80]. It is interesting to note that the utilization of the channel was quite low. Less than 0.03% of the packets transmitted were involved in collisions while 99% acquired the channel with no latency.

One of the main reasons for Ethernet's popularity is because it uses a passive broadcast medium. This results in very reliable operation. Ethernet interfaces can be built using VLSI technology and thus made fairly inexpensive. Multi-vendor environments can be implemented by adhering to interface specifications at any of several levels. For instance, one may choose to provide compatibility at the wire tap, the transceiver cable, the port, higher level protocols, etc. Because of the heterogeneous environments in which Ethernets are used one can expect to see a great variety of traffic distributions. In this paper we study the traffic generated in a University environment and predict the performance of the Ethernet when used to satisfy the needs of that environment.

### Methodology

This study deals with the behavior of Ethernet in the interactive time-sharing and program development environments. There are many installations which fall in this category. Our analysis is based on the measurements at one such installation - a large University with a number of large hosts presently connected to each other by conventional direct connections. We asked the question: "What will the traffic on the Ethernet at this University look like if an Ethernet was installed today?". We hypothesized that for the near future, the university will still have the dumb terminals (asynchronous, character mode) that are being used today and that these will be connected through terminal concentrators to the Ethernet. Others will still have direct connections to hosts since it is not likely that existing hardware will be thrown away. However, the users of those terminals still will generate Ethernet traffic in transferring files, sending mail, etc. The hosts will continue to have local secondary storage which will be used for user files and temporary workfiles. We assumed some level of file transfers and mail messages between hosts. Since we could not extrapolate the current traffic of this type into the superior sharing environment of the Ethernet, we assumed three somewhat arbitrary levels for traffic of this type.

## PREDICTED CAPACITY OF ETHERNET IN A UNIVERSITY ENVIRONMENT

Our principal objective is to predict the maximum number of users supported when the limiting resource is the Ethernet. In other words, we wish to estimate the number of users that can be supported on the Ethernet when all other resources such as terminals, processors and secondary storage are available in sufficient quantities so as not to be bottlenecks. There are two ways in which the environment affects the number of users supported. First, it dictates the higher level protocols to be used while executing the commands given by the user. This in turn affects the amount of traffic generated by a user. Second, it specifies the packet size distribution which has a significant role in determining the performance of Ethernet.

In estimating the Ethernet traffic we assumed that typical layered network protocols would be used. We coupled the user level workload with this model of the distributed architecture to estimate the average number of packets per active user per second. The packet level Ethernet simulation is then executed while increasing the number of users until the idle time goes to zero. Since the existence of Ethernet will cause more sharing and thus more host to host file and mail traffic, this workload alone is not sufficient to predict the total Ethernet load. We therefore study the network behavior with three levels (low, medium and high) of host to host file and mail traffic.

Note that in estimating the number of users a system will support one must also examine the user perceived response time and determine if it meets the requirements for the applications, environment, etc. Other bottlenecks such as disk delays, host processing of protocol messages, application program contention for memory and CPUs may play a larger role than the Ethernet in determining the user perceived delay. Those other possible bottlenecks may limit the number of users able to be supported to a smaller number than predicted here,

### Performance Metrics

As mentioned above, here we concentrate on the performance at the Ethernet level. The delay through the Ethernet and the throughput as functions of offered load are two important performance metrics. The delay is often small compared to the delays in the higher levels. The main parameter controlling Ethernet performance is the ratio of the one way propagation delay (ie: half the slot time) to the average packet transmission time. This is called "alpha". The performance improves as this ratio is made smaller [MAR80], [SHOC80]. This is because packets are exposed to collisions only during the first slot time of their transmission. Once a packet has been on the wire for that length of time it should not experience a collision. Under heavy load the throughput will be better if alpha is smaller [SHOC80].

The number of collisions a packet experiences in attempts to transmit is another interesting metric. Each collision causes the backoff range to be doubled. One would hope that, on the average, a packet does not experience many collisions. Measurements [SHOC80] and simulations [MAR80] have shown that there are few collisions in typical systems.

One could devise other metrics relating to the higher level protocols such as number of packets transmitted for each user message, etc. However, here we examine worst case scenarios and do not pursue that topic. It should be noted that the higher layers often dictate the performance of the network and therefore they should be carefully studied [MQUI80]. They will produce extra packets for each user packet transmitted. These control packets contend with the data packets for the limited resources of the shared channel (Ethernet). They also contend with other applications for resources (CPU cycles and memory) at the transmitter and receiver. Here we only address the issues relating to the shared channel.

## PREDICTED CAPACITY OF ETHERNET IN A UNIVERSITY ENVIRONMENT

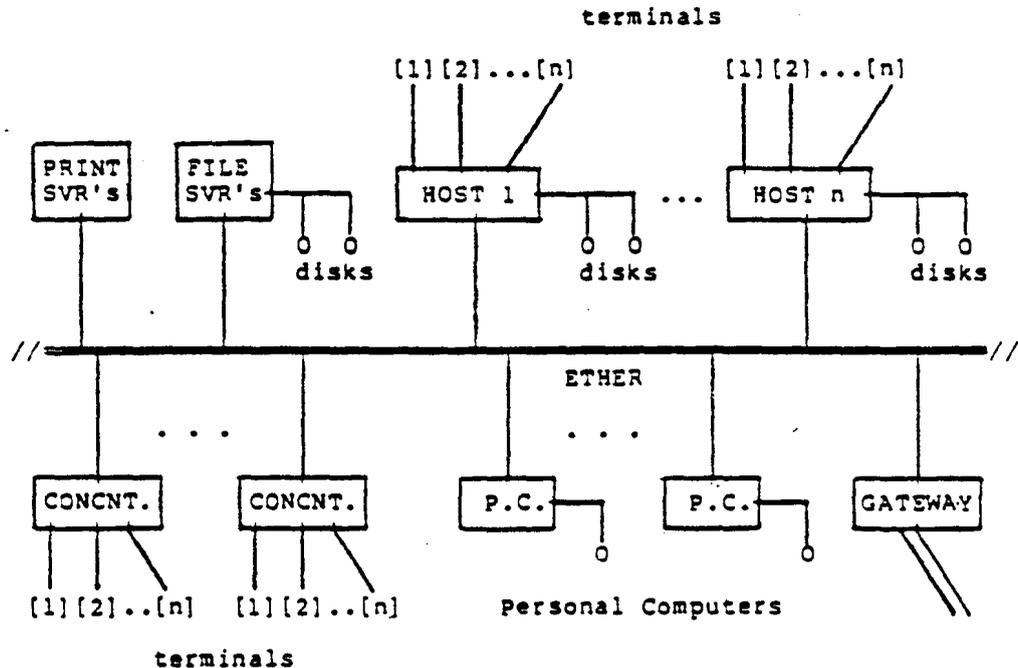


Figure 1. Local Area Network Components.

### PROGRAM DEVELOPMENT ENVIRONMENT

#### System Components

Figure 1 depicts a typical collection of components found in a Local Area Network. Users can be connected to the Ethernet through terminal concentrators, hosts, or through personal computer workstations. Disk requests made on the behalf of a user can be directed towards a local disk (on a personal computer or a host), or they may be directed towards a file server. Swapping and paging traffic is assumed not go over the Ethernet since the hosts have local disks for "system related" operations. In this study we assume that the disk requests generated by the users are, for the most part, satisfied at the host with which they are communicating. However, remote file access and transfer (for mail, etc.) does use the Ethernet. Initially, the Local Area Network will not contain all the devices depicted in Figure 1. However, as time passes file servers, etc. will be added to the system.

#### User Profile

The workload contains descriptions of the activities of the users. User perform operations such as file edits, links, compiles, executes, etc. They also perform typical "house keeping" operations such as directory listings, file copies and deletes, etc. They send and receive mail and communicate with other users using interactive message facilities. The characteristics of the users were measured during heavy usage periods for several days at the University. I/O as well as program image related data was collected. Table 1 summarizes some of the major points of interest in the user I/O characteristics. The table contains the mean value of several interesting statistics. It is important to note that many of these statistics had bimodal, trimodal, etc. distributions. This means that more than the mean is required to fully understand the data.

PREDICTED CAPACITY OF ETHERNET IN A UNIVERSITY ENVIRONMENT

PARAMETER	VALUE
1) Avg. Session Duration	1307 seconds
2) Avg. Input Size (Term -> Host)	10.7 bytes
3) Avg. Input Rate (Term -> Host)	0.16 inputs/sec
4) Avg. Output Size (Host -> Term)	26.5 bytes
5) Avg. Output Rate (Host -> Term)	0.34 outputs/sec
6) Avg. Printed Character Rate	2.91 chars/sec
7) Avg. Remote File Access Rate (Assumed Light Usage, See Text)	0.00567 accesses/sec
8) Avg. File Access Size (Directed Locally or Remotely)	3584 bytes/access

Table 1. "Per-User" Workload Summary

In deriving the total network traffic generated by each user, the data and control packets generated at each protocol layer as a result of a user transaction were totaled and used to drive the Ethernet simulation. The amount of disk traffic present on the Ethernet will change with time as more intelligent servers and workstations are added to the system and as usage patterns change due to those new capabilities. We therefore have varied the load due to disk traffic in the simulation. Various amounts of the user disk traffic were sent over the network. This traffic is normally channeled to/from the host's local disk and the host. Access rates of 0.00567, 0.0085 and 0.017 accesses/second/user were used. This corresponds to 3.3%, 5%, and 10% of the traffic a given user generates at the local disk on the host.

Figure 2 contains a histogram of the Ethernet packet sizes generated by the user interactions coupled with the protocol model. The packet size includes user data (if any) the preamble, CRC and all other protocol fields. The protocol model was based on the examples contained in architectural specifications. (See [DAP80], [DEC80], [DIGI80], [NSP80], and [SESS80] for details of the architecture. See

[WECK80] for an overview and description of its features and capabilities.) The model used assumes worst case examples. For instance, no acknowledgements are piggybacked. We also assume that each data packet transmitted requires its own acknowledgement and therefore there are no acknowledgements of multiple data packets. All of these assumptions are clearly worst case. They all increase the load on the Ethernet as well as the transmitter and receiver CPUs and memories.

RESULTS

Figure 2 contains a histogram of the Ethernet packet sizes generated by the user interactions coupled with the protocol model. The packet size includes user data (if any), the preamble, CRC and all other protocol fields from all protocol levels. The main contributor to the relatively large number of small packets (64 to 100 bytes) is the higher level protocol control packets. As mentioned previously, we have assumed the worst case for all protocol exchanges. This means that there are no piggybacked acknowledgements, etc. This imposes the heaviest load due to protocol control traffic. Since these are generally small packets, this distribution poses a demanding load on the Ethernet and

## PREDICTED CAPACITY OF ETHERNET IN A UNIVERSITY ENVIRONMENT

should produce conservative results for this user workload.

Figure 3 shows the Ethernet offered load versus the number of users for this workload. The Ethernet specifications indicate that a maximum of 1024 taps may be connected to an Ethernet. The simulation conforms to that rule. Note that several users can share a tap. This is the case with terminal concentrators and hosts that have local terminals generating Ethernet traffic. In the figures presented here, the "number of users" corresponds to actual users - not to physical transceiver taps (of which there is a maximum of 1024).

Figure 4 shows the mean waiting time versus the number of users. Figure 5 shows the 90th percentile of the waiting time. The waiting time is defined as the time from when the packet becomes ready for transmission until it begins successful transmission. It includes all time spent deferring, colliding and backing-off. As mentioned previously, three levels of remote file traffic were simulated. The "low level" corresponds to an access rate of 0.00567 accesses/user/second. The other two are for one and a half and three times the load due to that component. Note that with this time-sharing workload, the number of users supported is quite large.

Figure 6 shows the idle time on the Ethernet going to zero at the overload points. Again note that this occurs for an unusually large number of users. Figure 7 shows the number of attempts required to successfully acquire the channel as a function of the number of users. The number of attempts includes all collisions as well as the one successful attempt which acquires the channel. Note that even at an overload point with 2000 users, a given packet experiences an average of only one collision per successful transmission. Figure 8 shows the 90th percentile of the number of attempts.

### CONCLUSIONS

The results of the simulation indicate that the Ethernet has sufficient bandwidth to serve large numbers of users of the type characterized by the time-sharing workload. In practice, one generally does not operate the system

with the steady state load near the system limits. The finite rate at which the hosts, disks, users, etc. can generate and process information will prevent the steady state loading from achieving this level.

The waiting time experienced in attempting to gain access to the channel was shown to be within reasonable bounds. The number of collisions experienced by a packet attempting to acquire the channel was also shown to be quite low - even in the heavily loaded regions.

In summary, we can say that the Ethernet seems to be well qualified to carry the type of traffic experienced in the time-sharing environment. It has the capacity to support large numbers of users in this environment.

### Discussion

Here we have shown that the Ethernet is capable of handling the traffic generated in this time-sharing environment. To build an effective network, the operation of the higher level protocols must be examined. The delays encountered due to processing and queuing can result in poor user perceived performance if care is not taken in their implementation. One should also examine other environments to see how similar or different they might be and how this affects performance. For example, the office environment is very important.

### ACKNOWLEDGEMENTS

We would like to thank our colleagues in the Systems Performance Analysis Group, especially Rollins Turner, for obtaining the workload measurements as well as their help in analyzing the large amounts of data. We also wish to thank them, and others in Distributed Systems, for insights regarding the modelling of Ethernets in this environment. Finally, the people in Systems Performance Analysis and Distributed Systems Product Management who reviewed this paper deserve special thanks for their many useful comments and suggestions.

PREDICTED CAPACITY OF ETHERNET IN A UNIVERSITY ENVIRONMENT

FIGURES

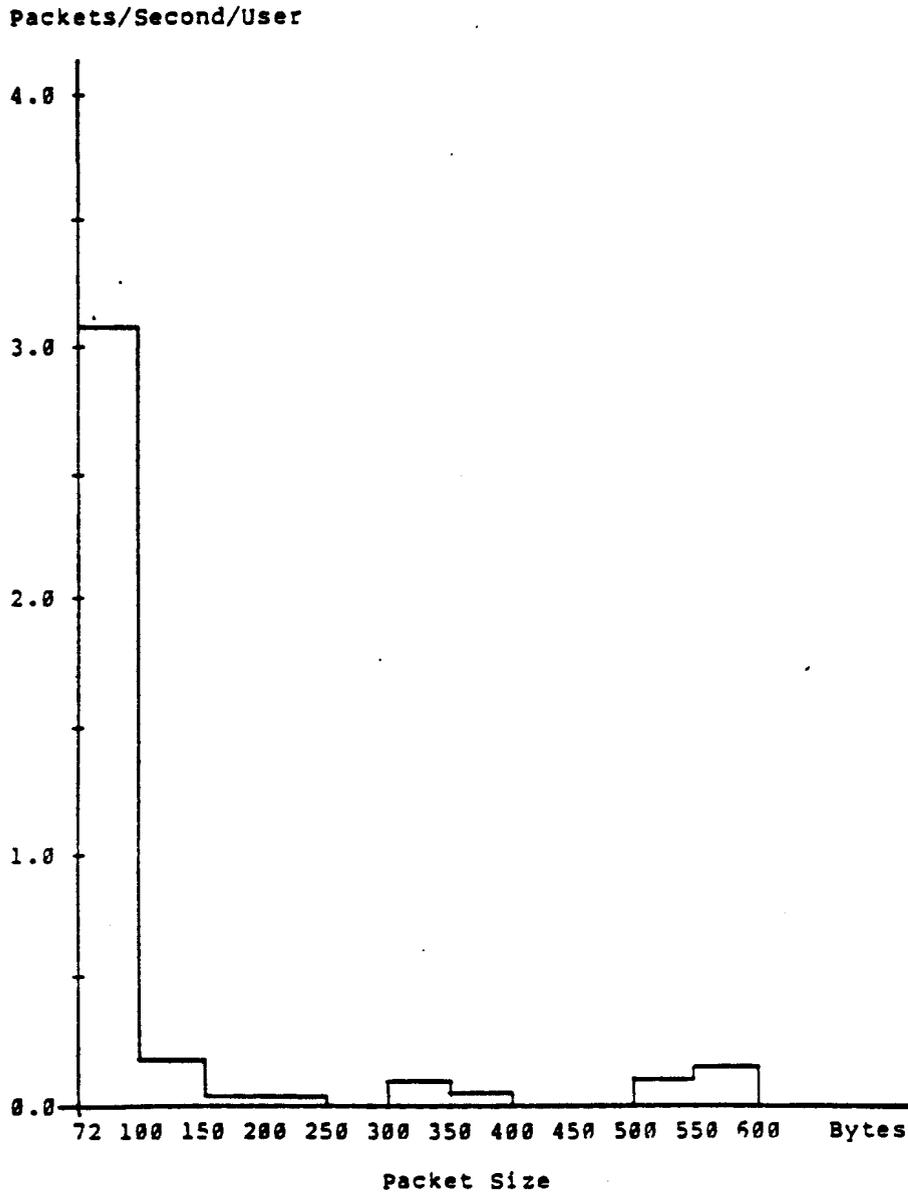


Figure 2. "Per-User" Ethernet Packet Size Frequencies  
(Low Remote Disk Traffic)

PREDICTED CAPACITY OF ETHERNET IN A UNIVERSITY ENVIRONMENT

Fig 3. % Offered load

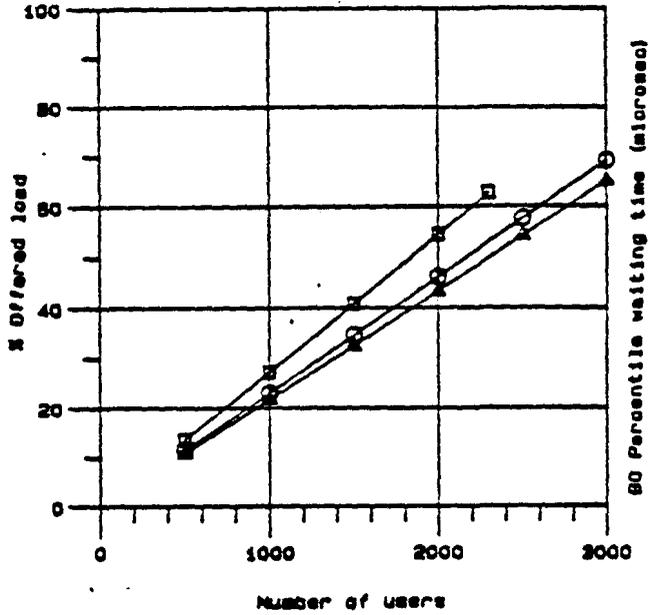


Fig 5. 90 percentile of waiting time

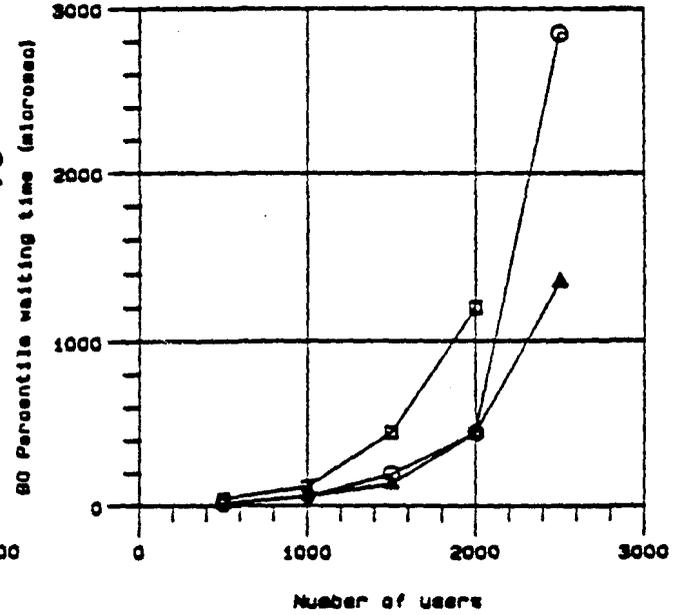


Fig 4. Mean Waiting time

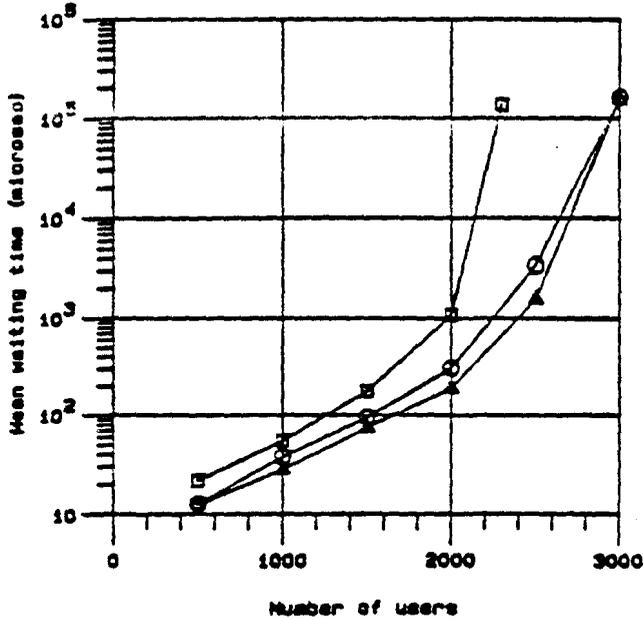
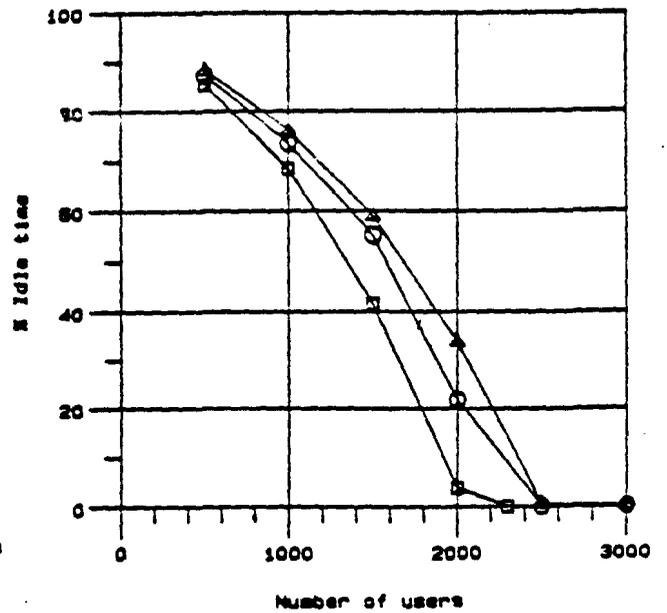


Fig 6. % Idle time on cable



- High network disk transfers
- Medium network disk transfers
- △— Low network disk transfers

PREDICTED CAPACITY OF ETHERNET IN A UNIVERSITY ENVIRONMENT

Fig 7. Mean number of attempts

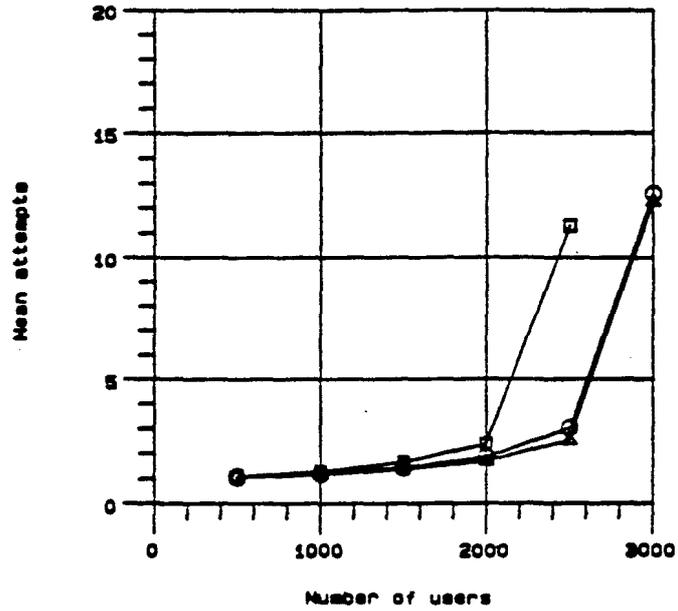
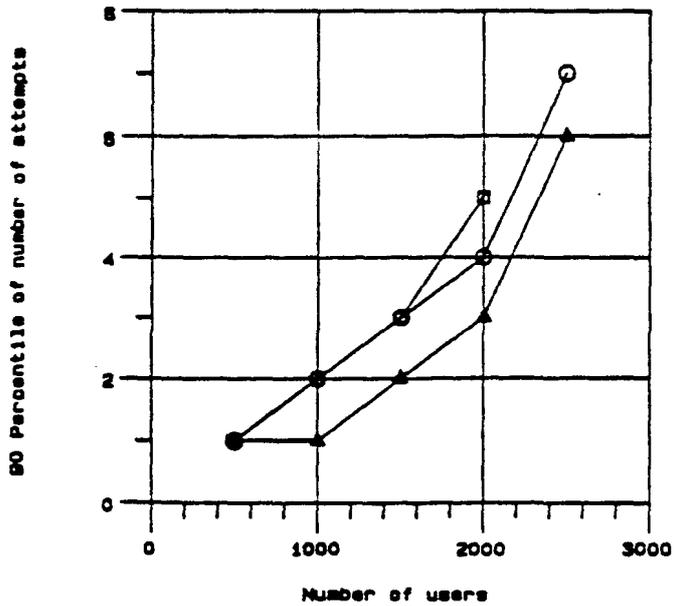


Fig 8. 90 Percentile of number of attempts



—□— High network disk transfers  
—○— Medium network disk transfers  
—△— Low network disk transfers

PREDICTED CAPACITY OF ETHERNET IN A UNIVERSITY ENVIRONMENT

REFERENCES

- [COTT80] TECHNOLOGIES FOR LOCAL AREA COMPUTER NETWORKS, I. Cotton, Computer Networks, Vol.4, No.5, Oct/Nov 1980, pgs. 197-208.
- [METC76] ETHERNET: DISTRIBUTED PACKET SWITCHING FOR LOCAL COMPUTER NETWORKS, R. Metcalfe, D. Boggs, Comm. ACM, Vol.19, No.7, July 1976, pgs. 395-404.
- [DAP80] DECNET DATA ACCESS PROTOCOL FUNCTIONAL SPECIFICATION, Version 5.6.0, Digital Equipment Corporation, October 1980.
- [MQUI80] LOCAL NETWORK TECHNOLOGY AND THE LESSONS OF HISTORY, J. McQuillan, Computer Networks, Vol.4, No.5, Oct/Nov, 1980, pgs. 235-238.
- [DECN80] DECNET TRANSPORT FUNCTIONAL SPECIFICATION, Version 1.3, Digital Equipment Corporation, March 1980.
- [NSP80] DECNET NETWORK SERVICES PROTOCOL FUNCTIONAL SPECIFICATION, Version 3.2.0, Digital Equipment Corporation, October 1980.
- [DIGI80] THE ETHERNET - A LOCAL NETWORK, Version 1.0, Digital, Intel, Xerox, September 1980.
- [SESS80] DECNET SESSION CONTROL FUNCTIONAL SPECIFICATION, Version 1.0.0, Digital Equipment Corporation, November 1980.
- [FREE80] UPDATED BIBLIOGRAPHY ON LOCAL COMPUTER NETWORKS, H. Freeman, K. Thurber, ACM Comp. Comm. Review, Vol.10, No.3, July 1980, pgs. 10-18.
- [SHOC80] MEASURED PERFORMANCE OF AN ETHERNET LOCAL NETWORK, J. Shoch and J. Hupp, Comm. ACM, Vol.23, No.12, Dec. 1980, pgs. 711-721.
- [MAR80] DESGIN ANALYSIS OF A LOCAL AREA NETWORK, M. Marathe, Comp. Network Symp., Wash. D.C., Dec. 1980, pgs. 67-81.
- [WECK80] DNA: THE DIGITAL NETWORK ARCHITECTURE, S. Wecker, IEEE Trans. Comm., COM-28, No. 4, April, 1980, pgs. 510-526.

## PROGRAPHICAL INFORMATION

### MADHAV MARATHE

Madav Marathe is a Consulting Engineer in the Systems Performance Analysis Group at Digital Equipment Corporation. He is involved in performance analysis of local area networks and data base systems. His present research interests are in distributed data base architectures, data base machines and file server architectures for local networks. He is a member of ACM and the Local Area Networks Performance Working Group of the IEEE Project 802 Local Area Networks Standards Committee.

He received a Ph.D. from Carnegie-Mellon University for research in the performance of the hardware architecture and operating systems kernel levels of a computer system. While there (and at Digital), he also did research in memory and data contention in multiprocessor systems.

## **DAVID E. LIDDLE**

Dr. Liddle received his undergraduate degree in engineering from the University of Michigan in 1966. He later received a M.S. and Ph.D. from the University of Toledo, the latter being awarded in 1972. From 1970 to 1972, Liddle was Project Manager at Owens-Illinois, Inc. in Toledo, Ohio.

In 1972, Dr. Liddle joined Xerox Corporation as a research scientist in the Palo Alto Research Center, where he worked on the "POLOS" office system project, early design issues for the Alto, and various file servers.

In 1975, Liddle became Manager, System Architecture in the Systems Development Department where he was responsible for the development, definition, and specification of the overall OIS architecture for Xerox. In January of 1976, he wrote and published the "OIS Architectural Principles". Until 1979, Liddle was Manager of the Systems Development Department. This organization developed the Mesa programming environment and the Ethernet communications network. It also produced the Pilot Operating System, and was responsible for a new family of software products for office applications.

In 1979, Liddle was appointed Vice President of the Office Products Division of Xerox. In this position, he had continued responsibility for the development of Ethernet, Integrated Network Services, and Advanced Information Processing Systems for OPD. He became a member of the Technology Review Group, a corporate committee which review strategic technical issues on a corporate-wide basis.

Currently, Dr. Liddle is Vice President and General Manager, Office Systems Business Unit for the Office Products Division. He has general management responsibility for office automation systems and products, including network services and professional workstations.



## ETHERNET EDITORIAL SEMINAR

### PRODUCTIVITY IN THE OFFICE ENVIRONMENT

In discussing 'productivity in the office', it's important to first clarify just what is meant by those terms. In this context, the concept of the extended office is being used: any structured association of people working with information. Essentially, improving productivity equates to producing more work, of higher quality, at less cost.

Business Week reports office costs are rising at the rate of 12-15% a year and will probably double over the next six years. Those costs are rising faster than any other cost factor - even faster than the costs associated with generating business revenue. Direct costs of office operations in 1980 were over \$920 billion and are likely to rise to \$1.5 trillion by the end of the decade. What's important to look at is that this productivity factor in the extended office is really people productivity.

Over 50% of American workers now work with information on a full-time basis. Nevertheless, nationwide, over ten times as much is invested in technology for factory workers as for office workers. When companies first attempt to "automate", their attention traditionally has been on the secretary. However, secretarial functions account for only 23% of office costs and only around 12% of salaries. At the opposite end of the spectrum, managers and executives have also benefitted somewhat from technology. The mainframe computer has typically provided data processing reports of various forms for use by the manager.

The person in the middle of the office hierarchy, the professional, has not had the benefit of technology even though professionals make up 80% of payroll costs - and their numbers are expected to grow 30% during this decade. Can their tasks benefit from technology and thereby, make them more productive? A number of studies say "yes".

To increase professional productivity, the tools they use must be improved, and the barriers to productivity must be eliminated or minimized in the four basic areas of information processing: creating knowledge, reproducing it, getting it in and out of files, and distributing it to others. Putting these two thoughts together says that the "better tools" must be able to work together; Ethernet provides the interoperability for that solution.

Every network user has the option of selecting the piece of equipment that best meets his or her own individual needs, whether that be the need for a recording typewriter for short letters and memos, or a personal computer to run accounts payable or inventory, or access to a mainframe computer. The user must not be limited to equipment from just one vendor.

It's important that the network and the products on it can grow in an evolutionary manner. A company should not be penalized by starting small. The evolution into automated office systems integrated on a network should not require a massive, all-encompassing galactic plan.

Ethernet has over 7000 person-years of testing and user experience. All of this experience supports the important premise of ease of growth and interoperability. The specifications for Ethernet were published jointly in September, 1980,

by Intel, Digital Equipment Corporation, and Xerox. Since that time, over 275 requests for license applications have been made; over 70 applicants have paid their license fee; 22 have publicly announced their intentions to build Ethernet compatible products. This speaks for itself; no other network technology has attracted such a broad allegiance. The fact is, it works. Over 50 installations of Ethernet networks within the past four months prove it.

Ethernet has provided a truly integrated approach to automating office tasks. Systems connected to the Ethernet operate simultaneously, and can be both standalone office machines and part of the network system, sharing resources or files or printing devices. The open architecture of Ethernet allows multi-vendor connectibility. The specifications have been published to allow other vendors that ability. The Ethernet customer is not forced to purchase all their equipment and services from one vendor. The higher level protocols that Xerox recently published take this connectibility a step further and allow any vendor to be truly compatible with other products on the Ethernet.

Ethernet's interconnectibility and interoperability is transparent to the user. The barriers to productivity can all be hurdled. Input, output, filing, retrieval, distribution - all can be accomplished from any system on the Ethernet. A secretary can print on the laser printer from an electronic typewriter. A manager can call up records files from a mainframe and manipulate them on a personal computer. The professional can access massive stored reports and extract information to prepare a summary report, complete with graphics, on a professional workstation. And everyone can distribute information to every other workstation on the system without the delays of mails and unanswered telephone calls. Network capabilities are driven by user needs, and Ethernet provides these integrated services critical for office productivity.

David E. Liddle  
Vice President & General Manager  
Office Products Division  
Xerox Corporation

**ETHERNET EDITORIAL SEMINAR**

**PRESENTATION HANDOUTS**

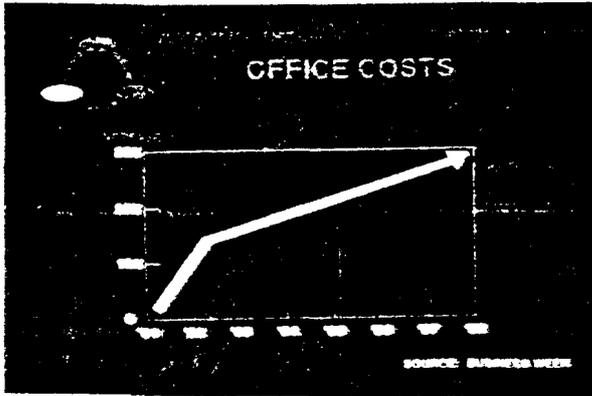
**PRODUCTIVITY IN THE OFFICE ENVIRONMENT**

**DAVID E. LIDDLE**

**XEROX CORPORATION  
OFFICE PRODUCTS DIVISION**

**February 10, 1982**





Office costs are rising 12-15%, doubling over the next six years. These costs are rising faster than any other cost factor, even faster than the costs associated with generating business revenue.

---

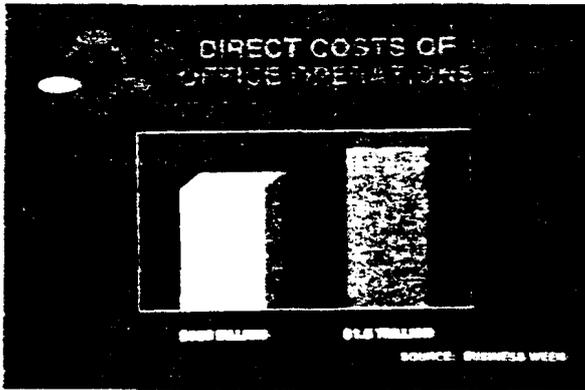
---

---

---

---

---



Direct costs of office operation in 1980 were over \$920 billion. Overhead expenses are expected to rise to \$1.5 trillion by 1990.

---

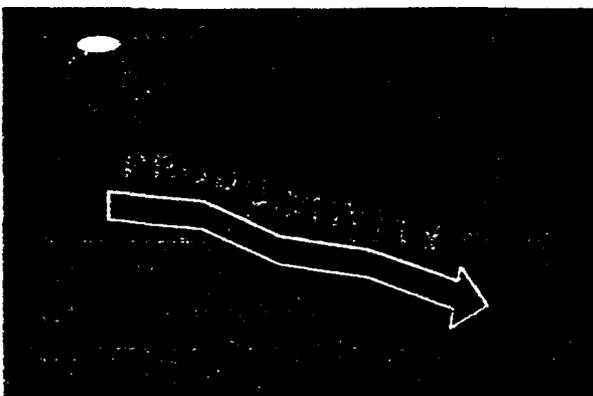
---

---

---

---

---



While office costs are rising, however, office productivity is declining.

---

---

---

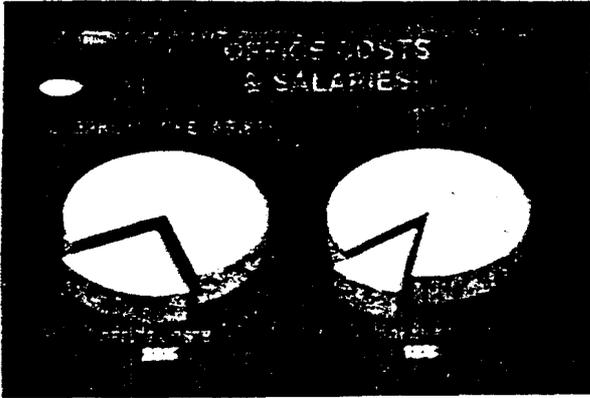
---

---

---







When companies first attempted to automate their offices, their attention traditionally has been on the secretary. However, secretarial functions account for only 23% of office costs and only around 12% of salaries.

---



---



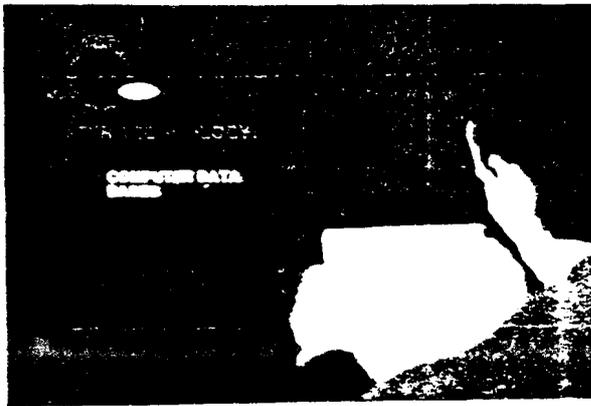
---



---



---



At the opposite end of the spectrum, managers and executives have also benefitted from technology. The mainframe computer, with its elaborate processing power, has typically provided data processing reports of various forms for use by the manager/executive.

---



---



---



---



The person in the middle of the office hierarchy, the **professional**, has not had the benefit of technology even though they make up 80% of payroll costs, and their numbers are expected to grow 30% during this decade.

---



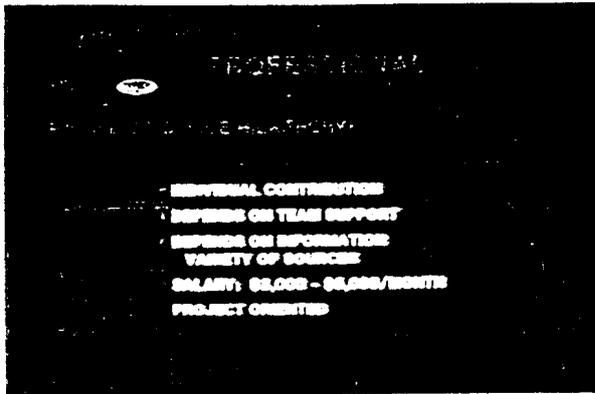
---



---



---



---

---

---

---

---

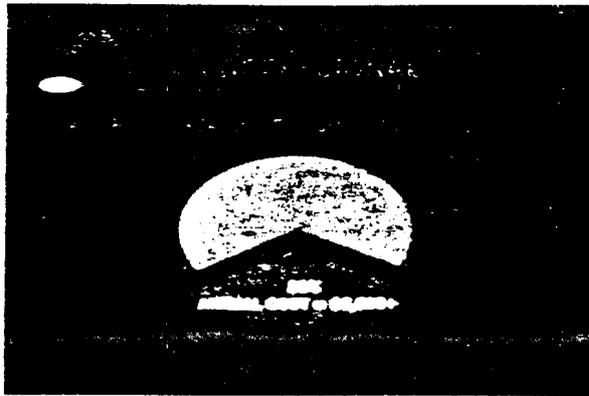
---

---

---

---

---



Can the professionals' tasks benefit from technology and thereby, make them more productive? A number of studies say "yes". For example, approximately one third of a professional's time is spent in creating documents at an average cost of \$6000/professional.

---

---

---

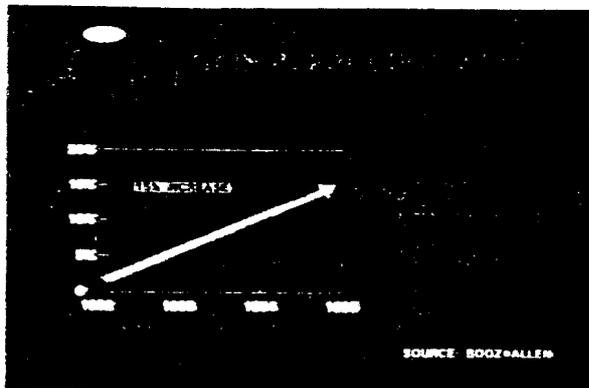
---

---

---

---

---



A recent Booz Allen study indicated that by utilizing office automation technology, a 15% gain in professional productivity could be realized by 1985. That's an average annual savings of \$5,500/information worker.

---

---

---

---

---

---

---

---

ELECTRONIC MAIL	2.5 DAYS
FILE TRANSFER	2.5 DAYS
REDUCED PHONE USAGE	47 MIN/DAY
EXAMINE DOCUMENTS SIMULTANEOUSLY	100%
REVERSE SHOOT DOCUMENTS	50%
REVERSE CYCLES	20%
COMPLEX DOCUMENT TURNAROUND	12.5 DAYS

In internal Xerox probe locations utilizing professional workstations and network services, these productivity gains were realized.

---

---

---

---

---

---

---

---

<b>PRODUCTIVITY</b>	
• COMMUNICATION OF DATA & IDEAS	
• CREATION OF DOCUMENTS	
• FILING & RETRIEVAL OF DOCUMENTS	
DISTRIBUTION OF DOCUMENTS	

These probes and studies indicate that to increase productivity, the barriers to productivity must be eliminated or minimized in the four basic areas of information processing: creating and communicating ideas and data, creation of documents, filing and retrieval of documents, and the distribution of documents.

---

---

---

---

---

<b>PRODUCTIVITY</b>	
INTEGRATION/COMPATIBILITY	
- FROM ELECTRONIC TYPEWRITERS TO MAINFRAMES	
- FROM SPECIALIZED TO MULTIFUNCTIONAL SYSTEMS	
FROM ONE TO MANY VERSIONS	

The office tools must be improved, and these tools must be able to work together if all four areas of information handling are to be impacted.

---

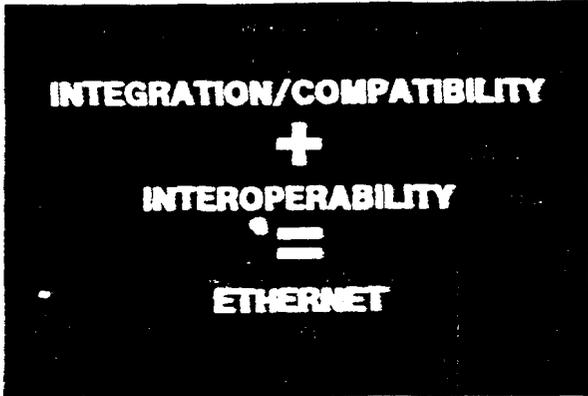
---

---

---

---

---



The integration of compatible products and the interoperability of products are available today on Ethernet.

---

---

---

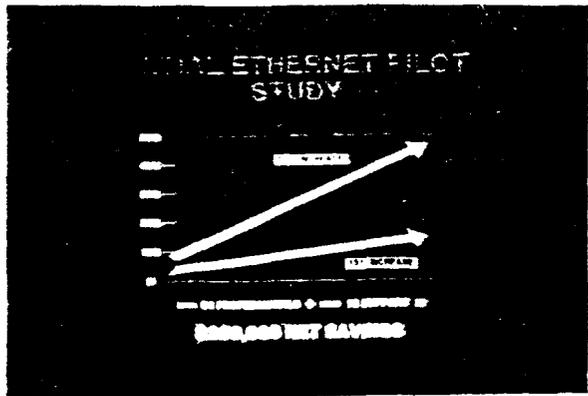
---

---

---

---

---



An Ethernet installation in a Fortune 100 manufacturing company showed a \$250,000 net savings during the initial year. Because of that immediate realization of productivity increases, we asked Booz Allen to analyze this installation in light of their original 1985 projections.

---

---

---

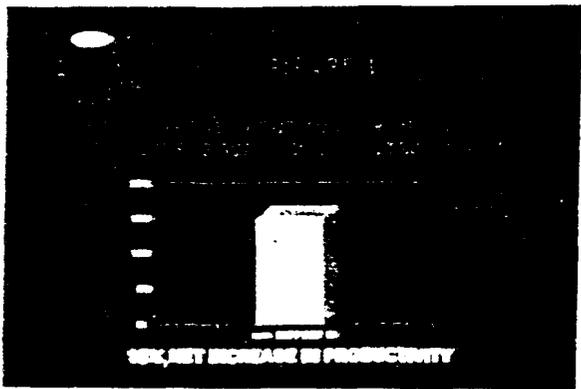
---

---

---

---

---



In Phase I of their extrapolation, a 15% increase in productivity could be expected with the support group workstations and file server connected to the network.

---

---

---

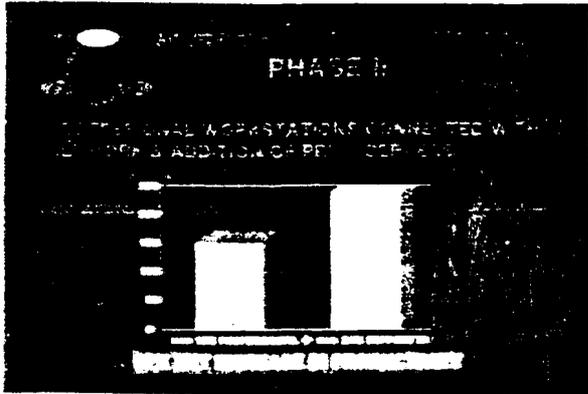
---

---

---

---

---



Phase II projects a 17% net increase in productivity with all professional workstations on the net and the addition of print services.

---

---

---

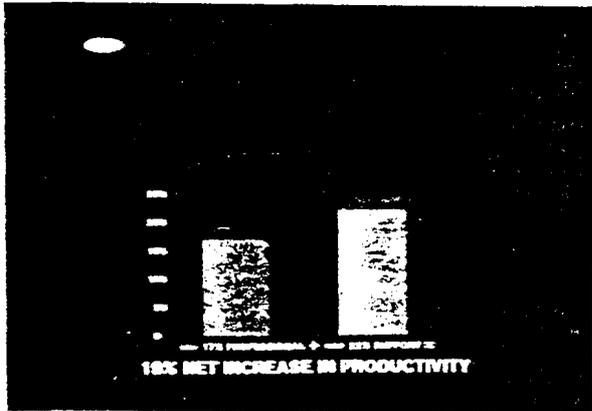
---

---

---

---

---



With the further addition to the net of electronic typewriters, personal computers for managers/executives, and communication services, an additional 19% net increase in productivity was realized.

---

---

---

---

---

---

---

---

**AFTER HAVING COMPLETED OUR ANALYSIS OF THE ETHERNET SYSTEM, WE CONCLUDED ITS POTENTIAL PROFESSIONAL TIME SAVINGS ARE VERY CLOSE TO OUR PREVIOUS 1985 ESTIMATE.**

HARVEY POPPEL  
SENIOR VICE-PRESIDENT  
BOOZ+ALLEN

**THE SOLUTION IS AVAILABLE TODAY**

This analysis clearly indicates the technology to increase office productivity exists now and is being utilized by Ethernet customers today.

---

---

---

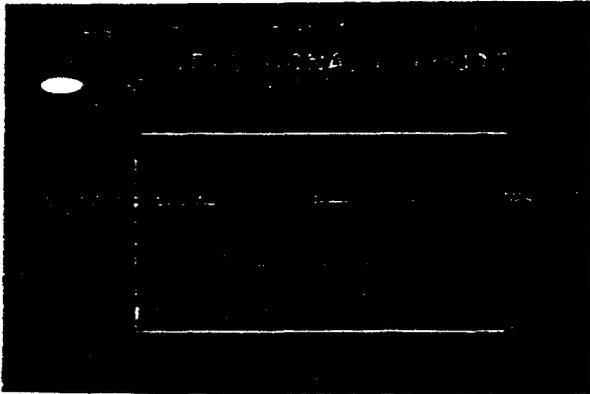
---

---

---

---

---



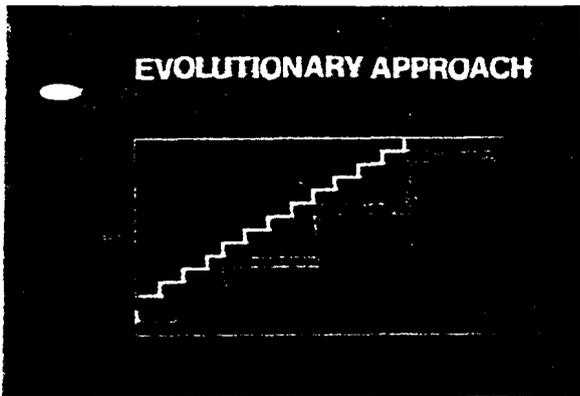
The traditional approach to "buying" technology was that the user had to buy giant pieces at a time; when the user began optimizing all of the capabilities, another big piece of equipment was purchased. The classic example of this is the mainframe computer.

---

---

---

---



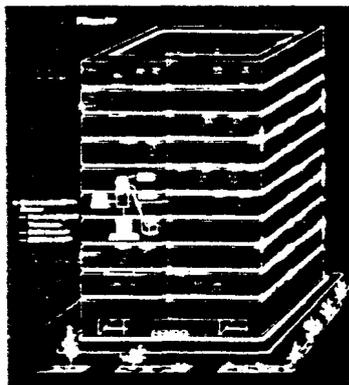
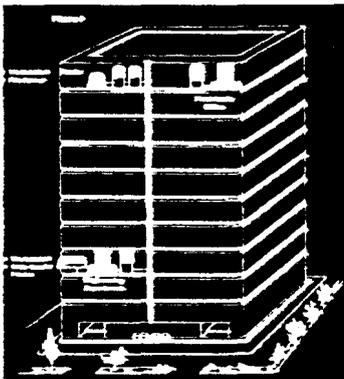
Ethernet, however, allows the user to start small, one work group or a department at a time. It is not necessary to have a comprehensive, long-term automation plan to begin automating an office.

---

---

---

---



A sample phased Ethernet installation.

---

---

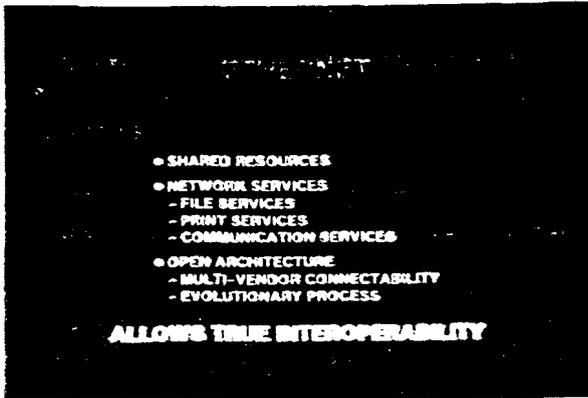
---

---

---

---

---



Ethernet provides generic local area network capabilities and with the published availability of higher protocol specifications, interoperability is a reality.

---

---

---

---

---

---

---



By providing integrated services and capabilities and interoperability with "foreign" products, Ethernet allows productivity improvements for all user groups.

---

---

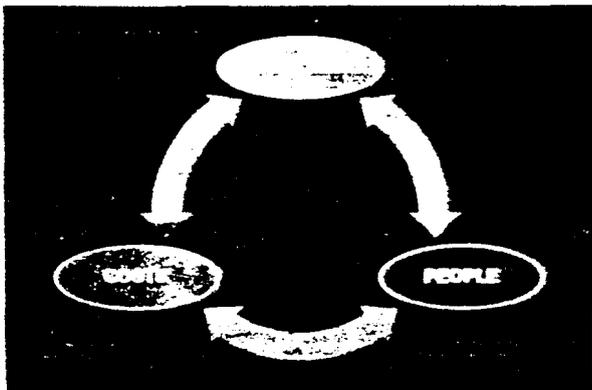
---

---

---

---

---



Ethernet has a proven history of providing the integrated services and capabilities, the true interoperability necessary to increase office productivity for all user groups.

---

---

---

---

---

---

---

DISPELLING ETHERNET MYTHS  
WILLIAM C. LYNCH  
MANAGER, TECHNICAL PLANNING  
OFFICE SYSTEMS BUSINESS UNIT  
OFFICE PRODUCTS DIVISION  
XEROX CORPORATION

Often when some new idea, some new concept is introduced you often hear a list of stories that are told about this new concept. The stories are repeated often, embellished and after a time, they become believed by everyone. What I am going to discuss with you today are some of these types of stories and embellishments, the myths we have heard that have collected over the last two years that we have been working on Ethernet. Frankly, some of them are rather remarkable. You have heard or will hear from the Intel speakers where we are in the program, that we are close to having a chip. As you well know, Xerox has been delivering Ethernet products. We are on our way to seeing Ethernet become a recognized standard. So I'd like to clear up as much as I can about the misconceptions that exist about Ethernet. What I want to do is go through several of these items, tell you what the myths are, and then what the truth is.

Before I do this, I would like to remind you of what the Ethernet specifications are. Ethernet exists in the two lower protocol levels of the ISO model, the physical and data link layers. They meet the ISO architecture, and it is an open architecture. We published these specifications jointly with DEC and Intel in September of 1980. The specifications of Ethernet define the electrical and

mechanical rules so that when you connect machines together that meet the specifications, they work. We specify in Ethernet a protocol called CSMA/CD. Its purpose is to multiplex data between machines.

This brings me to the first misconception: Ethernet costs too much. It costs too much to attach a terminal or a device to the network. You heard the presentation by Intel which covers that item.

Another myth about Ethernet is that it performs poorly under heavy load. DEC has presented a paper regarding performance, and perhaps some of you have also seen reports of our own Xerox experiments that were reported a year or two ago on the 3Mb prototype. The traffic patterns are about the same between the 3Mb prototype and today's 10Mb Ethernet. You have heard that in all performance evaluations, Ethernet performed suberbly.

The next misconception I'd like to dispel is the idea that Ethernet has limited bandwidth. After many years of study and experience with Ethernet, we found that there is more than enough band width to handle applications that we perceive for the next ten years.

Some opponents to Ethernet say it is statistical rather than deterministic. The first thing to understand is the term deterministic. How is it presented to you? It is presented to you in the following way: I can guarantee that when I have a message to deliver to you that I can give you an upper bound on when that message is going to be received by you at your terminal or work station. DEC has gone through a very detailed presentation on the Ethernet performance. The really

important issue is waiting time. Does the message get there fast enough for the application that we want? If it gets there fast enough, it doesn't really matter if it is statistical or deterministic. It gets there. It does the job. That is the issue I am bringing out here.

What you are told is that because Ethernet is statistical, it can not do certain things. I can give you an example of token ring and buses, etc., that are also statistical. For example, a token bus. As long as there are, say, ten workstations on my bus, and I'll never get any more than ten, and everybody is sending the same stream of traffic all the time, that is, a terminal user at a constant rate pushing the same button - the return or enter button - what I will get is a stream of data coming out from everybody. And if nothing breaks and nobody else wants to get on the bus, I can guarantee there will be a response time that is fixed. But now you've got to solve the following problem. You come in to your office in the morning and want to check your mail. All the people in your office probably do not come in the same time every morning. You come in and flip the switch and ask, "What is my mail today?" But then, what about messages that are coming back to you? I don't think that a deterministic process will handle this situation, either.

My point is there is nothing deterministic in this business. The reason these channels are shared is to take advantage of the fact that what you are going to do is probabilistic. But you can't predict in advance. Otherwise, you would give a fixed pair of wires for every terminal that is going to use the network. If you really want that - you can do it.

I want to show you this slide to give you an average to look at. These are numbers, response time. Response time is waiting time plus transmission time. You already heard that transmission time for Ethernet is a small value. Waiting time is a probabilistic number: some number of tenths to milli-seconds. If the question really is, "Can I guarantee that my message will get there in a tenth of a second if I use Ethernet?", the answer is: Yes. Data can get from sender to receiver in a tenth or hundredth of a second, almost always. Conservatively, data can get there in one-hundredth of a second, upwards of 90% of the time.

On the other hand, what does micro-second response time mean? It means that you have this packet that is only a few bits wide. What can you get in a few bits - maybe the preamble.

Another myth that we hear is that we have put too much intelligence into the terminal. What is happening now is that VLSI is allowing us to inexpensively locate a lot of intelligence at the work station. Let's take advantage of that. Improve performance and improve capability. More and more of this pattern is showing up. As a matter of fact, the trend is to put even more of it in the work station. You will find smarter, not dumber work stations in the future.

Next Myth: Ethernet protocol has no error control. Truth: Ethernet uses the Autoden 2, 32 bit FCS for error detection. Next truth: RS232 (X21 bis) has the same error control. This is exactly what the standards have been using for years. You have heard at least four of the presentations talking about layering. What has happened is that with computer networks, we are allowed to start layering some of

these functions, with the error detection at the lowest layer and error correction at the middle layers. Xerox provides error correct in the Transport Layer.

We have found that there are certain applications that don't require error control. Time of day, for instance. We send a clock down the network every few milliseconds. I am sure that I don't want to have to retransmit that because I lost it. I know another one is coming later. What we have done is take advantage of the kinds of things that you do on a network. You don't just send files, you also send control information from time to time. Or just plain information. We have a packet we call "Breath of Life" - it sort of floats around to initiate down-line loading. The communication server just sends it out and says here I am, does anybody want me? If it's lost, it's lost; another one is coming.

Myth: The Ethernet protocol (CSMA/CD) does not work with any other media. We have seen since at least 1969 CSMA/CD on every media that you can think of. Xerox implemented CSMA/CD on a fiber optic technology. We have had a Fibernet experiment running since around the 1977. There are a number of other vendors outside of Xerox looking at putting Ethernet on fiber. There are going to be differences in physical architectures, but a CSMA/CD takes advantage of multiple access. Second, broadcast or broadband technology has been using CSMA/CD technology for a long time. Miternet is an example. Of course, Wang's Wangband is CSMA/CD - it's the same protocol. That says something. We must be right. It is a basic way to tie computers on a network and communicate in a multi-access environment. In the Wang or Miternet implementation - in the first two levels of the protocol - what is different is that the physical channels, the physical implementations, are going to be different.

Question: How do you configure this local network called Ethernet? A misconception is that Ethernet has limited topology and topography, that there are only a few ways you can configure Ethernet. That is utter nonsense. We have Ethernets in high rise buildings, and in single floors. The installation we have of the Ethernet, for example in Palo Alto - at my office - is essentially a single cable - a snake between the floors. One Ethernet. We have installations where on each floor, there is a backbone Ethernet going down an elevator shaft and a single Ethernet on every floor. I think I know the source of this particular myth: If you look in the specifications, it says 100 taps per segment and then you look at the next picture - a maximum of 1500 meters cable in a network - linear - difference between two stations. And people start counting because the configuration you see in the specs is 500 meters, 500 meters - where do you get the 1000? It must be limited. Therefore, you can only get 1500 meters between stations. Again, that is nonsense. You can have at least three dimensions of a network topology.

Myth: Ethernet has a limited number of attachments. Truth: Ethernet has 1,024 tap locations. Each of these tap locations can interface several terminals. You can have just about any application to topology that you want. You have seen already from the DEC presentation that 2000 terminals is not a problem. Actually, more than that is not a problem. Referring back to my deterministic slide, remember that response time is in such a small time value range that the applications never see this response time. You sit at the terminal, working out what you are going to do next, which process to serve next. It will always take you longer to do that than it takes Ethernet to send the message on the network. So you are allowed all kinds of applications.

Myth: Application coverage is limited. This has the flavor of, "Ethernet can't do factory applications." That is, of course, not true. Ethernet has been used in factory applications. As a matter of fact, right now there is an Ethernet network in a manufacturing facility in Dallas that is wrapped around a power distribution cable. And every now and then the big switch on the wall goes "grrchunck". And the packets keep going and there is absolutely no problem. This takes us back to response time: a tenth of a second - no problem. Ethernet can be used in some applications of manufacturing. Ethernet is planned to be used and is being used outside of the office, dispelling the myth of Ethernet being suitable only for office applications.

These next issues have to do with the acceptability of Ethernet itself, the Ethernet protocol, how has it been received in the public, how has it been received in industry. The first myth here, that Ethernet has limited acceptance by the business and communities, is nonsense. You have already heard that at least 22 vendors that have publicly declared their intentions to be compatible or make components with Ethernet. There are more that are not yet public. There are numerous companies who intend to and who are investigating supporting Ethernet in the field as a product. This slide shows the types of products or components that are currently offered by non-Xerox vendors. You can see that the entire spectrum of the things that you need to do with an Ethernet are available: transceivers, controllers, controller chips, cable, systems, compatible stations. For instance, there are at least five transceiver vendors world-wide. By the way, there is an overlap in this list. Some of the vendors that are making transceivers are also making controllers. There are at least nine controller vendors. We have four chip vendors, three cable vendors. The system vendors consist of people who have

decided to supply complete Ethernet compatible systems. They range from software, computer-based systems all the way to the entire network. And the compatible station vendors are people making smart terminals to talk on Ethernet, which range from highly talented terminals or work stations to fairly low functions but direct connections.

All this reinforces Dave Liddle's earlier statement: our goal is interoperability. Open up the marketplace and let other vendors get into this. For example, Xerox does not make transceivers. We buy transceivers. Xerox is into the chip business. We buy chips. So it is important to us that we have received this wide acceptance of Ethernet. The point is that the acceptance in the community has been very high, very wide, and very complete where companies have committed money, time, and people to support Ethernet.

Finally, the last myth: Ethernet is just a development project and will never be implemented. We actually had a question asked of us recently -when will the first Ethernet be installed? The answer is, of course, yes, we have had Ethernet commercially installed for over one year. Two major companies that have discussed their experience with Ethernet to the U.S. press are TransAmerica and Arco. When I made this slide, there were 35 other networks that were up and running. There are another 50-60 networks in different stages of installation and operation. This does not count networks we have inside Xerox.

It is true. There is an Ethernet. It works, and it works reliably. What I have tried to do was go down the list of what I consider the really crucial myths, discuss

them, and tell you the facts that dispel these myths. Ethernet itself is certainly no longer a myth - it is a reality.

Thank you.

## **ROBERT S. PRINTIS**

Dr. Printis was born in Washington, D.C. He received a B.S. in Electrical Engineering from Howard University, an M.S. in Electrical Engineering from the University of New Mexico, and a Ph.D. in Electrical Engineering in 1974 from the University of Maryland with major research interest in systems theory. He is a member of Tau Beta Pi and Phi Kappa Phi honor societies.

Printis conducted research in systems theory and applied mathematics, with specific interest in applications of systems theory to large scale systems, at IBM Research Center in Yorktown Heights, New York, between 1973 and 1977. He spent a short time at Bell Laboratories, in Holmdel, New Jersey, in 1977-1979 working in the ACS project on network control and network management protocol development. He has been employed at Xerox Corporation since 1979, working in the area of communications networks. His work assignments included the design and development of the communication protocols for Xerox' long-haul data communications networks.

Dr. Printis is currently Manager, Network Standards in Xerox Corporation's Office Product Division, Systems Development Department in Palo Alto, California and is developing network management standards for local networks.

## **ETHERNET EDITORIAL SEMINAR**

### **DISPELLING ETHERNET MYTHS**

On September 30, 1980, Digital Equipment Corporation, Xerox Corporation, and Intel Corporation published version 1.0 of the Ethernet Specification. Despite the support of the standard by a wide range of institutions, both commercial and academic, there is still some confusion expressed about Ethernet, its design, and its operation. In today's presentation, several of these misconceptions will be discussed, using the experience gained from the installation and operation of the network since 1975 for the experimental network and since 1980 for the commercial 10Mbit version.

This discussion is separated into six areas of concern: design, configurability, application coverage, acceptability, performance, and costs of Ethernet. Any local area network technology must deal with these concerns. We will discuss how Ethernet addresses these issues. The presentations by DEC and Intel will have addressed the issues of performance and costs. The Xerox presentation will address the remaining four areas.

Robert S. Printis  
Manager, Network Standards  
Office Products Division  
Xerox Corporation

# **DISPELLING ETHERNET MYTHS**

**Robert S. Printis  
Xerox Corporation  
Office Products Business Unit  
Palo Alto, California**

## MYTH CATEGORIES

- DESIGN
- CONFIGURABILITY
- APPLICATION COVERAGE
- ACCEPTABILITY
- PERFORMANCE
- COSTS

## DESIGN

- **ETHERNET HAS LIMITED BANDWIDTH (capacity)**

- 10Mbit per second capacity more than adequate for local computer network applications envisioned for next ten years

- **ETHERNET PROTOCOL IS "STATISTICAL" (rather than 'deterministic')**

- Definitions

- **Waiting Time** -- elapsed time from the time that the packet is ready for transmission until the packet successfully begins transmission.

- **Transmission Time** -- propagation time of the packet on the medium.

- **Response Time** = Waiting Time + Transmission Time.

- **Deterministic System** -- Waiting Time known to a fixed upper bound. Therefore, the Response Time is bounded, under normal operation of the channel.

- **Statistical System** -- Waiting Time's upper bound known with probability.

**--Response time requirements met by Ethernet**

<b>.100 sec</b>	<b>Yes, unless system is broken</b>
<b>.010 sec</b>	<b>Almost always, unless system is broken</b>
<b>.001 sec</b>	<b>Misses this requirement if long packet (Maximum packet size = 1518 bytes)</b>
<b>.0001 sec</b>	<b>Possible for small packets</b>
<b>.00001 sec</b>	<b>Forget it</b>

**--For a point of comparison, 9600 baud line**

<b>.100 sec</b>	<b>no, if message exceeds 120 bytes</b>
-----------------	---

● **PARTITION OF ETHERNET FUNCTIONS PLACES TOO MUCH OF THE COMMUNICATIONS RESPONSIBILITY ON THE 'TERMINAL'**

-- Reduction in price-size-performance due to VLSI permits introduction of more communication function in the station.

-- Permits the design of more efficient communications.

-- The direction of the future is to place more, not less, communications in the station.

● **ETHERNET PROTOCOL HAS NO ERROR CONTROL FOR DATA TRANSMITTED ON THE CABLE**

-- The situation is the same as that in RS232C data communications.

-- The Autodin II 32 bit FCS is specified by the Ethernet Specification.

-- It is common in computer communications networks to place error recovery in the transport layers of the communications protocols.

- **ETHERNET PROTOCOL (CSMA/CD) CANNOT WORK ON OTHER MEDIA**

- Fiber Technology -- Fibernet

- Broadband Technology(coaxial cable)

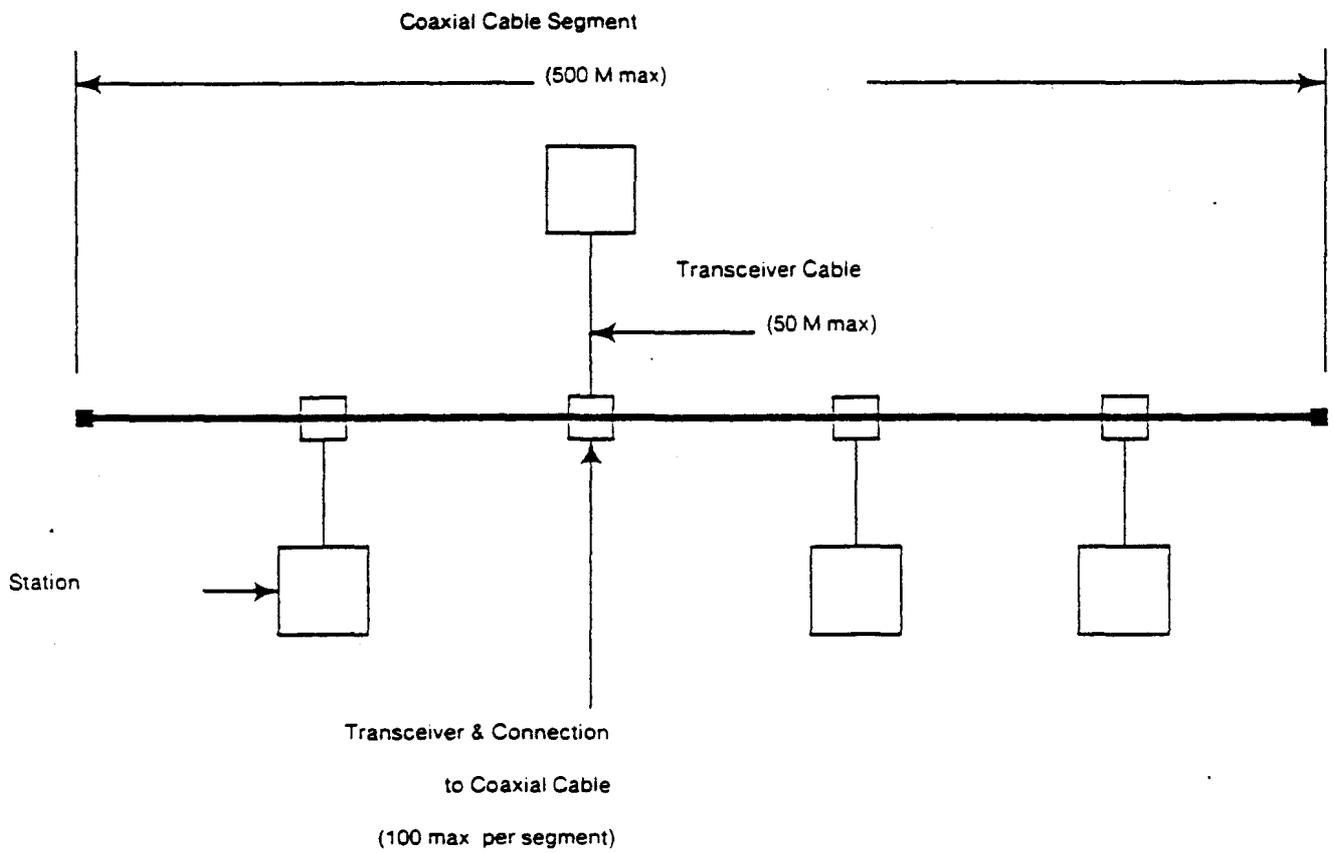
- Mitrenet

- Wangnet's Wangband

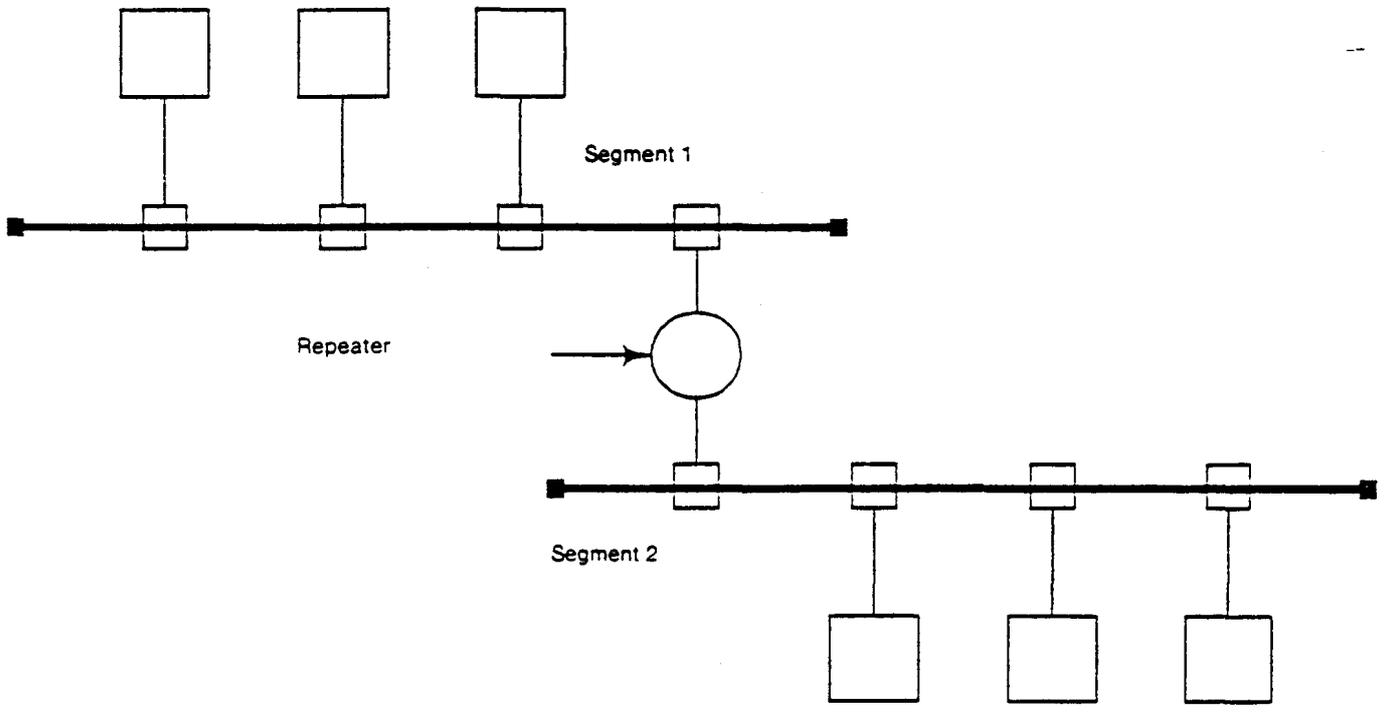
- Layering permits this -- must define physical channel interface, but protocol is media independent.

## CONFIGURABILITY

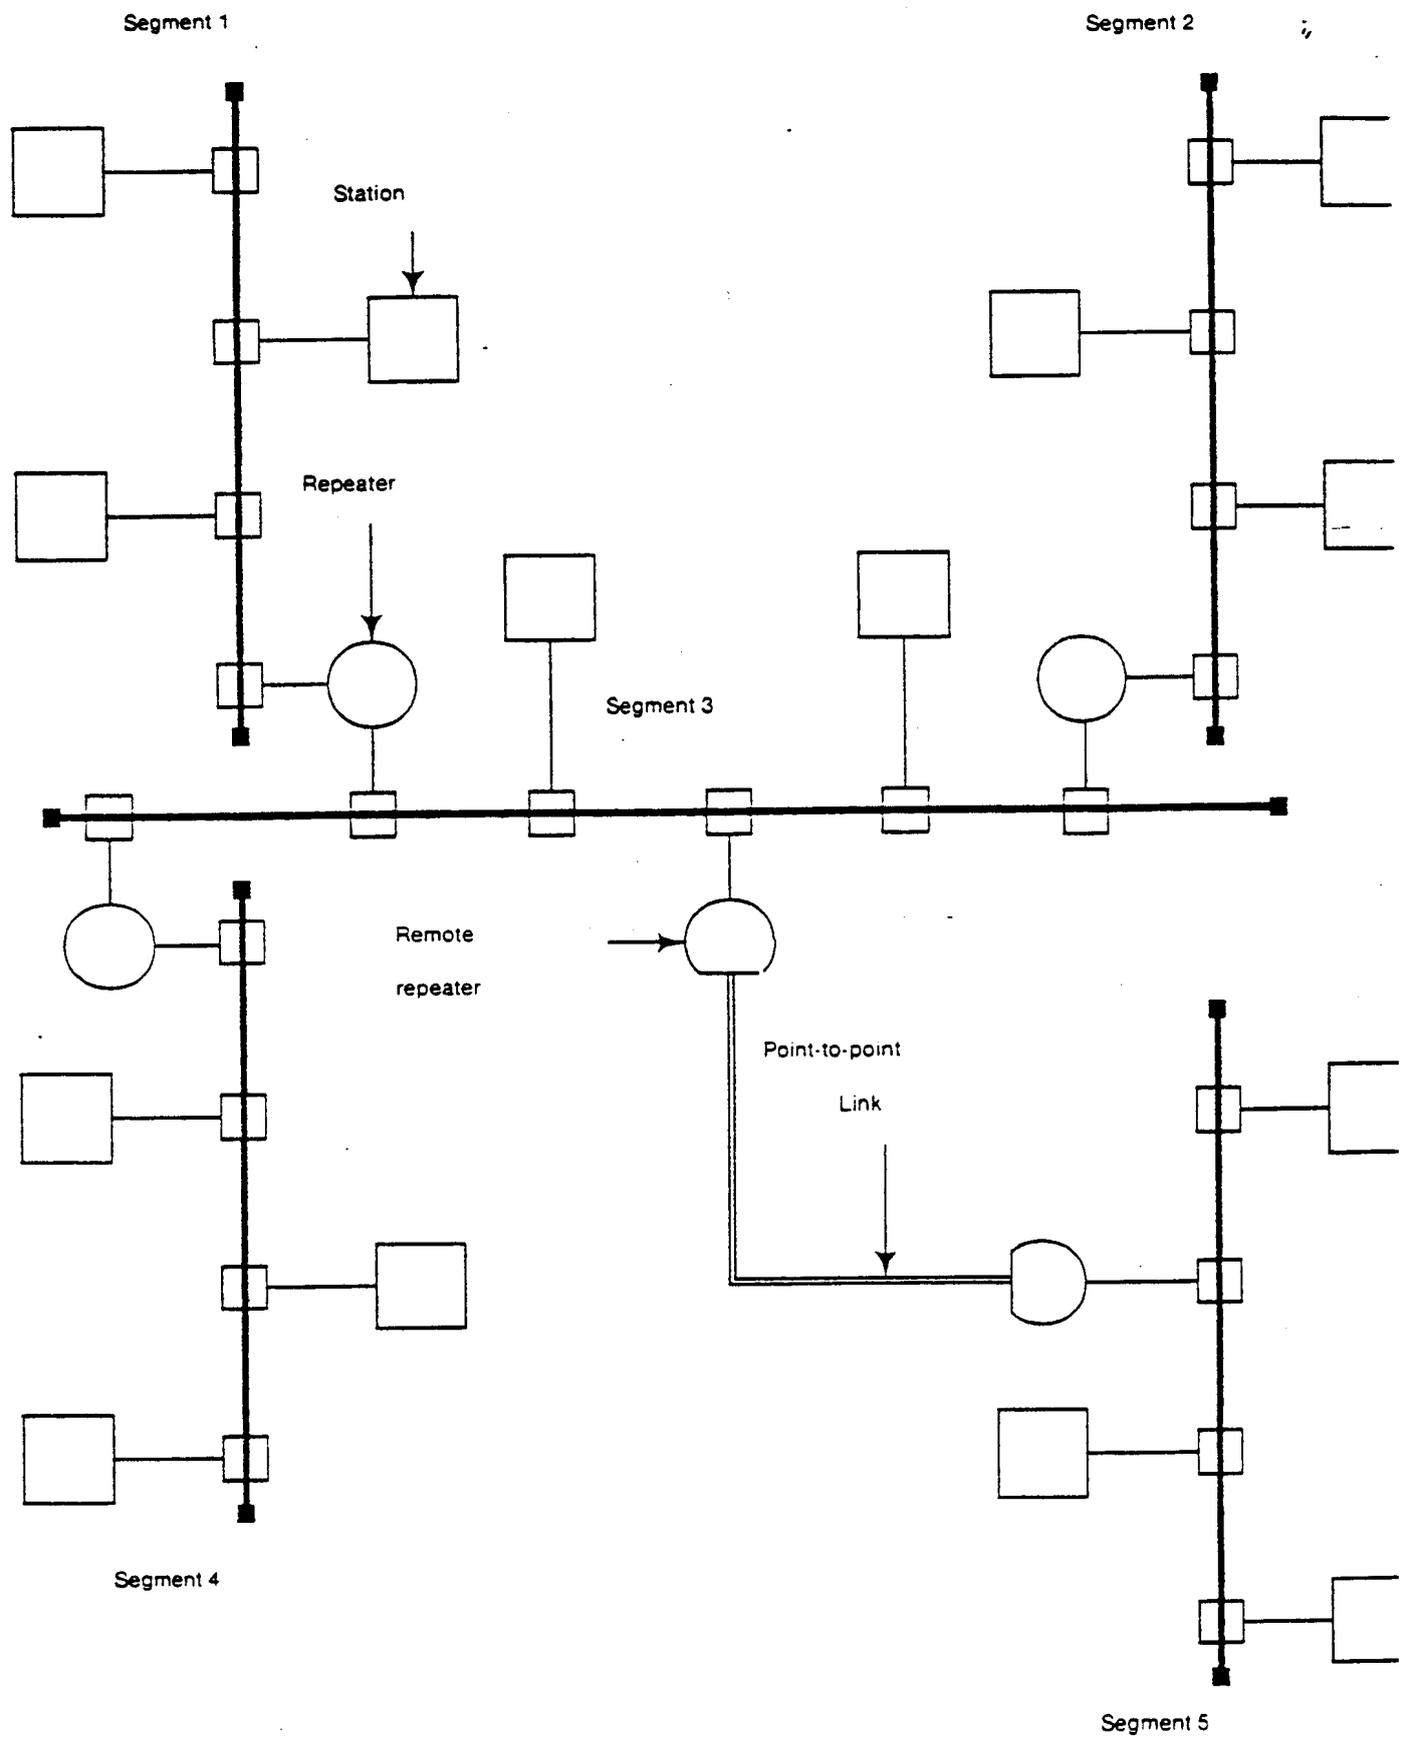
- **ETHERNET HAS LIMITED TOPOLOGY AND TOPOGRAPHY**



Small Ethernet Installation



A Medium-scale Ethernet Installation



Large-scale Ethernet Installation

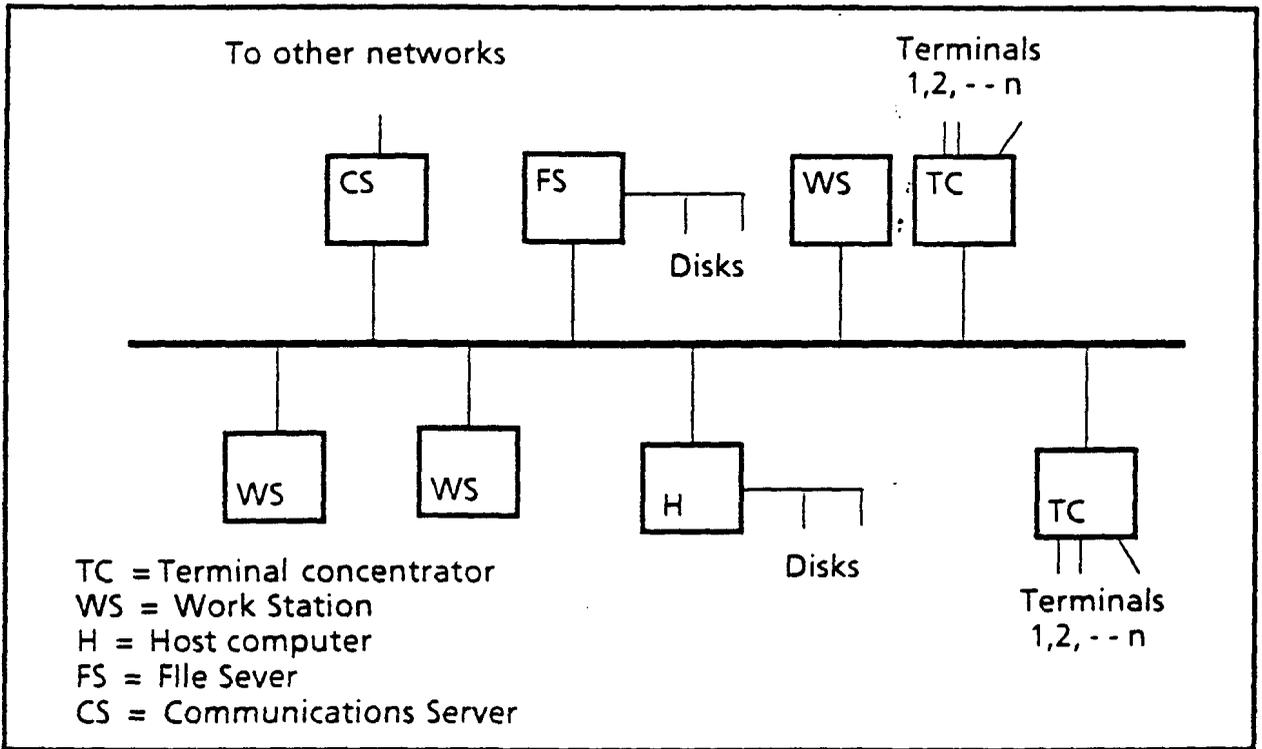
- **ETHERNET HAS LIMITED NUMBER OF POSSIBLE ATTACHMENTS**

- 1024 attachments possible, i.e., tap locations

- Several stations can share a tap.

- For example, with the Xerox 873 Communications Server -- 8192 RS232C Ports

- Numbers of users depends upon applications



## APPLICATION COVERAGE

- **ETHERNET CANNOT BE USED IN REAL-TIME PROCESS CONTROL APPLICATIONS**

--If "real-time" means message delivery by 0.1 sec then, may use Ethernet

-- Ethernet can be used in some applications in manufacturing.

-- DEC and Intel briefings give more examples of Ethernet in other environments.

- **ETHERNET IS SUITABLE ONLY FOR OFFICE APPLICATIONS**

**DEC intends to use Ethernet as local network for applications which include the office and traditional data processing**

**Many of the companies licensed to use Ethernet are not in the office automation business, but are data processing companies**

**No technical reason which restricts use to office applications.**

## ACCEPTABILITY

### ● ETHERNET HAS HAD LIMITED ACCEPTANCE BY THE BUSINESS AND TECHNICAL COMMUNITIES

Twenty two vendors announced their intention to provide compatible systems

Transceivers	Five vendors
Controllers	Nine vendors
Controller Chips	Four vendors
Cable	Three vendors
Systems	Ten vendors
Compatible stations	Ten vendors

● **ETHERNET IS JUST A DEVELOPMENT PROJECT AND WILL NEVER BE IMPLEMENTED OR INSTALLED**

**TransAmerica**

**Arco**

**35 other nets**