01 May 1987

Ethernet Performance: Design and Implementation Study

Michael M. Chaney

Tachen L. Lo

Daniel C. St. Clair
Missouri University of Science and Technology

ETHERNET PERFORMANCE: DESIGN AND IMPLEMENTATION STUDY

Michael M. Chaney, Tachen L. Lo,
Daniel C. St. Clair

CSc-87-8

Department of Computer Science

University of Missouri-Rolla

Rolla, MO 65401 (314) 341-4491

*This report is substantially the M.S. thesis of the first author, completed May, 1987.

ABSTRACT

General concepts concerning local area network designs, functions and topologies will be presented. Ethernet as a multipoint bus topology local area network will be presented in detail. The Carrier Sense Multiple Access/Collision Detect (CSMA/CD) method of fairly regulating access to the shared network bus is studied. The Ethernet Network in relation to the Open Systems Interconnect (OSI) is reviewed, but only the layers pertaining to Ethernet are discussed throughout the majority of the paper. The specifications as described by Xerox, Digital and Intel are presented to help the designer understand the network's physical limitations. Analytical models are used to predict performance and actual measured performance studies will be used to make performance assumptions. The performance is studied under varying load conditions. The data gathered concerns both limits imposed on the number of users by the finite bandwidth of the channel and efficient utilization of that channel. In conclusion, design specifications and performance data will be used together to formulate a design methodology for building the most efficient Ethernet network.

ACKNOWLEDGEMENT

I would like to thank Emerson Electric Company for the use of their Ethernet Network in this paper. Observing an operational Ethernet installation was a key factor in understanding Ethernet's design and performance. I would also like to thank Dr. T. L. Lo. His participation as my Co-Thesis Professor was very helpful. His networking expertise added many insights into this work. Both Dr. Lo and Emerson Electric's cooperation are greatly appreciated. Finally, the use of Patricia Anderson's thesis as a working format model is also appreciated.

TABLE OF CONTENTS

TABLE OF CONTENTS (CONTINUED)

LIST OF ILLUSTRATIONS

LIST OF ILLUSTRATIONS (CONTINUED)

Page

LIST OF TABLES

I. INTRODUCTION

It has been over 10 years since (Metcalfe, 1976) gave us the first formal description of an Ethernet Local Area Network (LAN). Since then, it has become a very popular method of networking workstations. As network management becomes a key issue in the 1980's, questions concerning Ethernet come to mind. What type of performance does it give? How does it delegate access to the shared bus? Because of its ease of installation and simplicity, the specifics of its operation are transparent to the user. This paper will examine some important areas of Ethernet concern: Ethernet and LANs (Figure 1 shows a general Ethernet design), Ethernet and the Open Systems Interconnect (OSI), and how Ethernet transmits and receives data on the shared bus. More importantly, Ethernet performance will be discussed from both a predicted and a measured viewpoint. Analytical models will be presented and verified which will allow the reader to estimate Ethernet performance for their own installation. Finally, a design methodology for using performance to fine tune Ethernet installations is described. A case study (design to implementation) for the Emerson Electric Company is presented as an example of the technologies discussed.

A two-segment Ethernet.

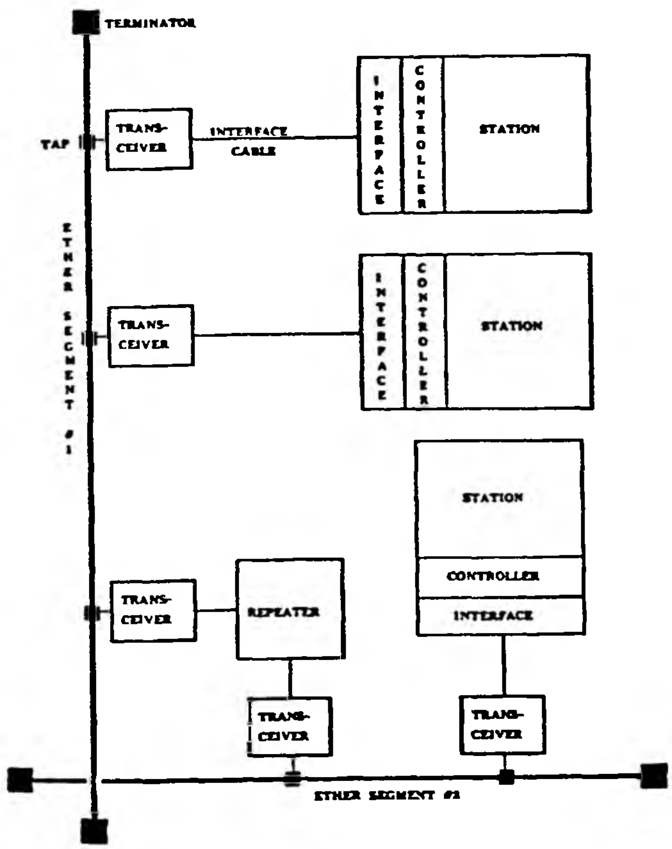

FIGURE 1 - (METCALFE, 1976) GENERAL ETHERNET CONFIGURATION

II. LOCAL AREA NETWORK DESIGN

Until the LAN push came along, a multiple host based computer facility was a cabling nightmare. For every different host at the facility, a point-to-point link was needed from a terminal (3 cables from the office to the CPU room for three different hosts) to its destination. As the terminal users moved within the facility, their cables had to be moved. This procedure was not very cost effective. Now, the LAN allows access to many hosts through a single connection. Other benefits of local area networks are the ability to share devices (Laser printer, disk drive, etc.) between systems, electronic mail, distributed processing and on-line file transfer. Overall, the LAN has improved the efficiency of the office operations from both a cost and time standpoint.

The push for LAN came because of a growing need to share information. Digital Equipment Corporation (DEC) and Micom-Interlan both agree the need for LAN's can be seen in the 80/20 corporate model (Micom, 1987). As we examine the corporate hierarchy (Figure 2), within each group of the organization, 80% of all information flows locally within the group and 20% flows outside the group (Digital, 1982). As data flows upward within the corporate structure, the same philosophy holds true. This localization of data prompted the need for a local communications link. Another driving force behind local networking was the need for distributed processing. Today, as microprocessors become cheaper and more powerful, the need for these standalone machines to communicate with each other becomes greater.

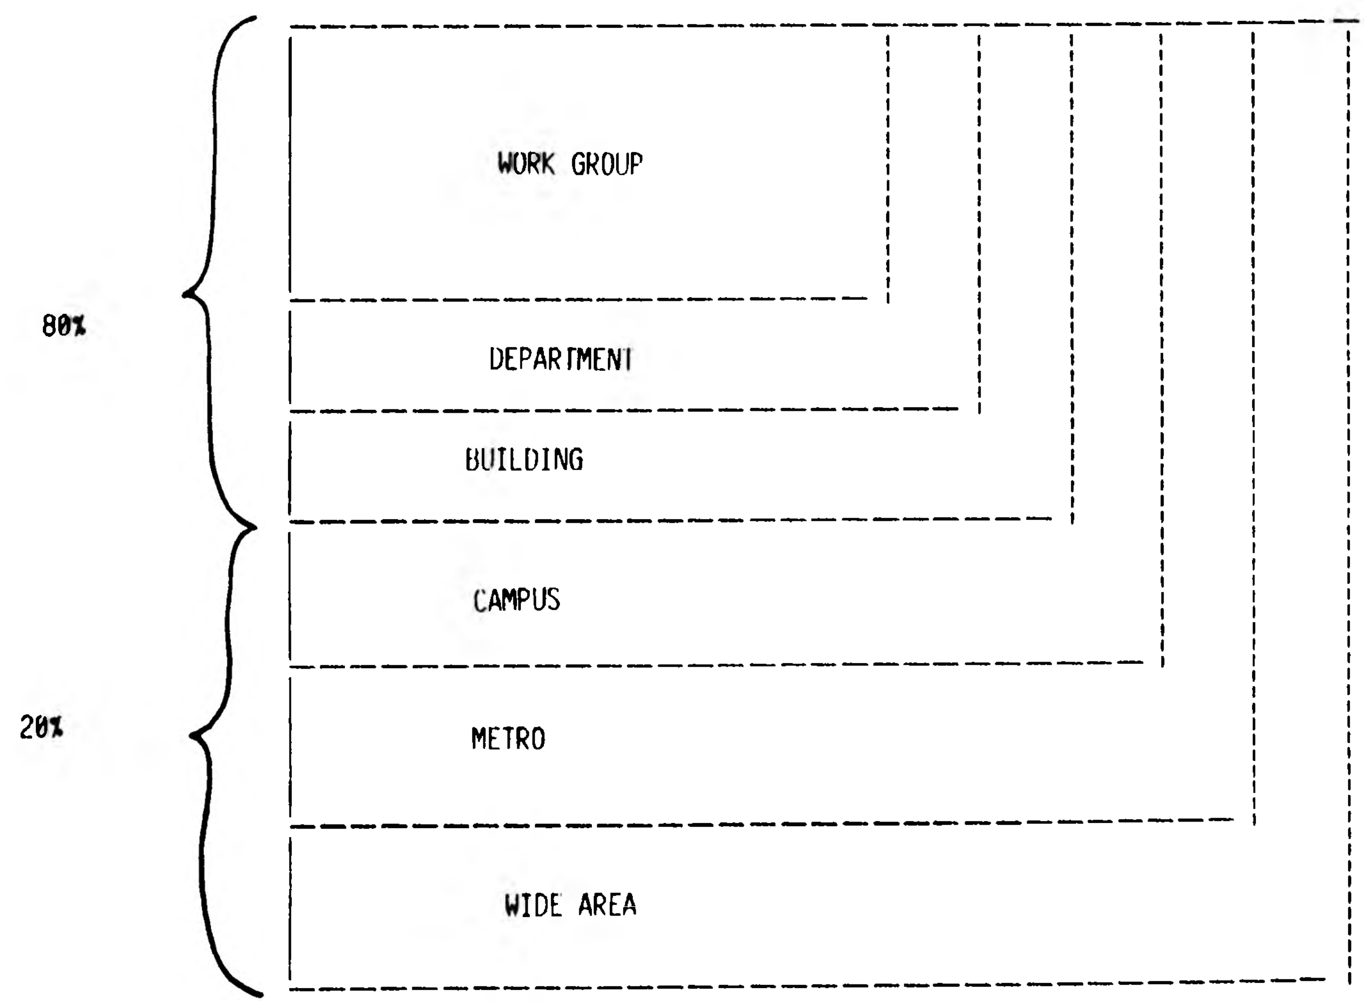

FIGURE 2 - 80/20 CORPORATE HIERARCHY MODEL

Local area networking concerns communication not computer technology. A LAN is a communications tool which allows a great number and variety of machines to exchange large amounts of information at high speeds over limited distances (Digital, 1982). LANs have the following characteristics: privately owned, total distance of a few square kilometers (.1 - 50 km), high data rates (.1 - 100 mbps), low error rates (1/(10**8) - 1/(10**11)), multiple vendor facilities, high reliability, any-to-any communications, low cost, and proven technology (Lo, 1984).

The best way to examine the different LAN designs is to look at their topologies and medium access methodologies. The nature of the LAN is determined by its topology and transmission medium. The general LAN topology can be either Point-to-Point or Multipoint. There are 3 major topologies used in the design of LAN's: point-to-point (star), point-to-point (ring), and multipoint (bus) (Brindley, November 10, 1986).

Point-to-point is the simplest LAN connection. Here, two nodes are connected via a single connection without passing through an intermediate node. Before LAN's, point-to-point links were needed for direct connection between a terminal and a host. Today, the point-to-point topology is used to connect different LAN's together.

The Multipoint link is a single line shared by more than two nodes (Digital, 1982). It is intended to yield a less costly cabling solution by having several nodes share the transmission capabilities of a single trunk cable. Since there may not be a controller node associated with this type of network, each node on the multipoint link may need to be more intelligent than its point-to-point counterpart.

The star network topology allows multiple nodes to connect to each other through a central controller node (Figure 3-A). Different nodes are connected to the central controller via point-to-point links. All communications must go through this central node before being routed to its destination. New nodes are added to the network by connecting them to the controller (a new point-to-point link). The main drawback to this design is that if the controller fails, the entire network crashes.

Ring networks consist of multiple nodes connected together in a ring fashion (Figure 3-B). All communications flow unidirectionally from source to sink around the ring. Each node must be able to recognize its own address so it knows whether to process the incoming data or pass it on to the next node. Line access can be managed/controlled using methods like token ring (data can be transmitted when the node receives a network token) and slotted ring (a node transmits data during its designated time slot). Adding a new node is difficult. The current connection must be cut and the the two free ends are reconnected to the new node. In the ring network, if the ring is broken or the controller node fails, the entire network crashes.

The last LAN type is the multipoint (bus) topology (Figure 3-C). It consists of a single shared medium used by all nodes to communicate on the network. The nodes use a line access technique of either the polling or contention variety to transmit the data over the line. Each node must be able to recognize its own address to receive data. Adding a node to this type of network occurs by attaching a node to the trunk

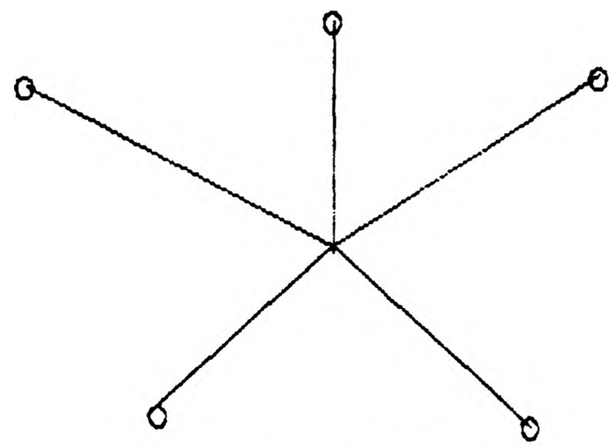

Star Network
FIGURE A

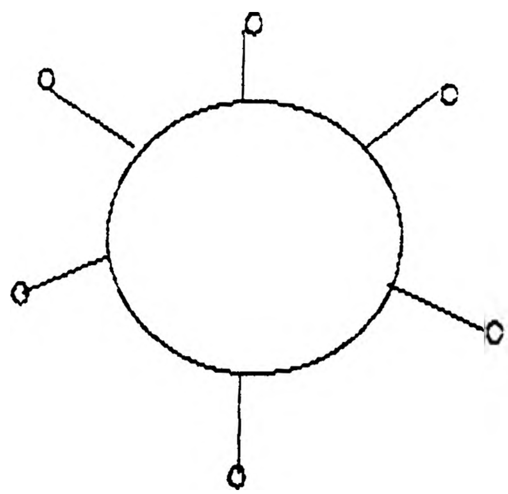

Ring Network
FIGURE B

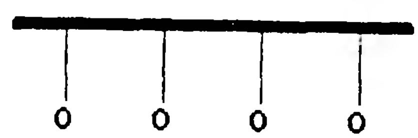

Multipoint (bus) Network
FIGURE C

FIGURE 3 - LAN NETWORK TOPOLOGIES

line (main medium). The advantage of this topology is that if a contention type line access scheme is used, a failure by any node will still keep the network operational.

III. UNDERSTANDING ETHERNET COMPONENTS AND OPERATIONS

Ethernet is a variation of the multipoint topology. Ethernet, a bus LAN, is a type of feedback communications channel where all stations can sense only the channel state and their own state and use this information to arbitrate transmission distributed access methods (CSMA/CD) (Stuck, 1984). The early designers wanted a reliable network. The bus topology was choosen because it was simple (reliability through simplicity). Metcalfe (Metcalfe, 1976) indicates that Ethernet gets its name from the luminiferous ether through which electromagnetic radiations were once alleged to propagate. Because of this, the central bus medium of the Ethernet network is usually called the ether.

Digital, Xerox and Intel created the specifications for a data delivery service with a very high transmission speed (10 mbps). Their specifications (Digital, 1980) describe the medium (trunk line), connections to this trunk line and data delivery techniques. The following two lists show the goals and non-goals of this Ethernet design. (Digital, 1980)

GOALS:

1) Simplicity
2) Low Cost
3) Compatibility
4) Addressing flexibility
5) Fairness
6) Progress
7) High Speed

8) Low delay

9) Stability

10) Maintainability

11) Layered Architecture

NON GOALS:

1) Full Duplex

2) Error control

3) Security

4) Speed Flexibility

5) Priority

6) Hostile user

This Ethernet specification can be seen as a small piece of a much bigger networking picture, the Open Systems Interconnect (OSI).

A. ETHERNET AND THE OSI

The International Standards Organization (ISO), a Class D member of the International Consultative Committee for Telegraph and Telephone (CCITT), devised a modeling schema for designing communications and networks. It was called the Open Systems Interconnect Reference Model (OSI) (# 7498) (Lo, 1984). This same organization has devised standards like RS-232 in the past. The OSI is a seven layered model (Figure 4) describing how communications architectures should be implemented to allow connections between heterogeneous computing environments. The seven layers are as follows (Hancock, 1985) :

1. Physical Layer - the touch-and-feel layer - provides for the transparent transmission of bit streams from one physical

entity to another (or many as in the case of datagram oriented services such as Ethernet).

2. Data Link Layer - handles the transfer of data between the ends of a physical link.
3. Network Layer - handles the routing and switching of information to establish a connection for the transparent delivery of data.
4. Transport Layer - provides for error-free delivery of data and also acts as the control area for quality of service requirements for the selected data service.
5. Session Layer - provides the coordination between communicating processes between nodes ("virtual connectivity").
6. Presentation Layer - provides for any format, translation, or code conversion necessary to put the data into an intelligible format.
7. Applications Layer - allows the end application to communicate with the communications architecture by providing the appropriate communications service(s) to the application.

The Ethernet specification pertains to layers one & two only (Physical Link and Data Link Layers). In an actual network configuration, Ethernet performs the functionality of the bottom two layers of the OSI, while a network protocal such as IBM's SNA, Department of Defense's (DOD) TCP/IP, Digital's DNA or Xerox's XNS will drive the upper layer network functions

```
7                     Application                   7
6                     Presentation                  6
5                        Session                    5
4                      Transport                    4
3                        Network                    3
2                         Link                      2
1                        Physical                   1

^                                                   ^
|                                                   |
               Communications channel
---------------------------------------------------
Party 1                                       Party 2
```

FIGURE 4 - OSI REFERENCE MODEL

(Digital, 1982). Ethernet is only responsible for the physical delivery of data gathered and driven in the upper layers of that networking protocal (Figure 5). At the Ethernet level, packets are delivered with high probability only. If a residual error rate greater than that provided by the bare Ethernet transport mechanism is required, then mutually agreed upon packet protocals must be followed. The entire network protocal is a very important part of an Ethernet network. This paper limits discussions to the OSI layers which concern Ethernet only. The hardware needs of an Ethernet network will be considered first followed by the software tasks which drive this hardware.

B. ETHERNET CONFIGURATION

Of all the hardware needed to install an Ethernet network, the least obvious component is probably the most important. The trunk line that supports all multipoint (bus) networks becomes almost forgotten once installed; but without it, there would be no network. The Ethernet's shared access facility, the ether, is a passive broadcast medium with no central control. Messages over the ether are flooded to all areas and nodes on the Ethernet (no routing is needed within an Ethernet). There are four standard types of media used as trunk lines in Ethernet installations: twisted-pair, thin-wire Ethernet (DEC calls it "Cheapernet"), baseband coaxial cable and fiber optical cable (Brindley, November 24, 1986). Twisted-pair is the oldest transmission medium. It is used for low speed (1-2 mbps) transmissions. It is susceptible to interference and noise including cross talk from adjacent

ISO SEVEN LAYERS	DNA LAYERS
APPLICATION	USER
	NETWORK MANAGEMENT
PRESENTATION	NETWORK APPLICATION
SESSION	SESSION CONTROL
TRANSPORT	END COMMUNICATIONS
NETWORK	ROUTING
DATA LINK	DATA LINK (Ethernet)
PHYSICAL	PHYSICAL LINK (Ethernet)

FIGURE 5 - ETHERNET, THE OSI AND DNA

wires (this can be minimized with proper shielding) (Lo, 1984). Thin-wire Ethernet is used in places where flexibility and cost are an issue. The thin-wire cable is more flexible than the normal trunk lines and is often used to network personal computers and workstations. Though fiber-optical cable will probably be the Ethernet medium of the future, baseband coax is the most prevalent today. The specification for the trunk line calls for up to 10 mbps data rates with a maximum length for any segment of 500 meters. The only failures that should be expected from the ether are breaks in the medium and decomposition due to bad or unplugged taps.

Since a distance limitation exists on the trunk line, what can be done to increase the allowable length of the cable? Repeaters can be used to chain a maximum of three trunk segments (500 meters each) together. The repeater takes the weakened signal and regenerates it to full strength. This hardware may be the most vulnerable of the Ethernet network (other than the cable itself). If a repeater crashes, then the two segments of cable on either side of the repeater will still operate standalone, but will not be able to communicate with one another.

Similar to the repeater are the packet filter and gateway. The packet filter (bridge) will pass traffic from one Ethernet to another (as well as regenerating the signal). The gateway allows the Ethernet to connect to networks of other varieties.

It should be noted that Ethernet communicates on the network by flooding it with transmissions (i.e. all areas of the network get the same message). It waits for the return of the message to see if there was any transmission problem. For this reason, each open end of the trunk line must be terminated. This allows the Ethernet transmissions to bounce off the cable end and return to sender.

There are two pieces of hardware required to attach nodes to the network: the transceiver and transceiver cable. The transceiver attaches directly to the trunk line itself. It transmits and receives data from the ether. The transceiver is designed with redundant logic, and has little or no chance of failure. When the transceiver powers off, it should disconnect from the Ethernet. This allows the network to stay up even when an individual node crashes. There may be a weak link in the network if a bad transceiver tap is made into the trunk line. The transceiver is tapped into the trunk line via a small hole. A transceiver cable is connected off the transceiver. This cable is made up of four twisted pair wires. Each wire carries one of four signals: transmit, receive, collision presence and power. The maximum length of a transceiver cable is 50 meters. The specification says, "All Ethernets must be compatible at the coaxial cable." (Digital, 1980) So they must all work the same at this transceiver connection.

The final piece of hardware on the Ethernet is the controller. It resides at the end of the transceiver cable. According to the specifications (Digital, 1980), the controller is the implementation unit which connects a station to the Ethernet. It typically comprises part of the physical layer, much or all of the data link layer and appropriate electronics for interfacing to the station. Basically, it is the low level firmware or software for getting packets onto or off of the network. Also, as will be seen later, when a source station detected collision occurs, it will be the source station controller's responsibility to generate a new random retransmission interval based upon the updated collision count. The controller may reside in its own housing (a terminal server) or be installed straight into the station (e.g. a DEUNA board in a VAX). This controller represents the connection from the network to the host's I/O bus.

Now, the next step is how to configure the hardware. The following is the Ethernet channel configuration model which tells how to configure the different pieces of the network. The maximum configuration is as follows (Digital, 1980) :

1. A coaxial cable, terminated with its characteristic impedance at each end, constitutes a cable **segment**. A segment may contain a maximum of 500 meters of coaxial cable.

2. A maximum of 2 repeaters may exist in the path between any two stations. Repeaters do not have to be located at the ends of segments, nor is the user limited to one repeater per segment. In fact, repeaters can be used not only to extend the length of the channel, but to extend the topology for one to three dimensional. Repeaters occupy transceiver positions on each cable segment and count towards the maximum number of transceivers on a segment just as do the logically distinguishable stations.

3. The maximum total coaxial cable length along the longest path between any two transceivers may be 1500 meters. The propagation velocity of the coaxial cable is assumed to 0.77 c worst case. (c is the velocity of light **in vacuo**; **300,000** kilometers per second). The total worst-case round-trip delay for all the coaxial cable in the system is therefore 13 μs.

4. A maximum of 50 meters of transceiver cable between any station and its associated transceiver is allowed. Note that in the worst case the signal must pass through six 50 meter transceiver cables, open at the transmitting station, one at the receiving station, and two at each repeater. The propagation velocity of the transceiver cable is assumed to be .65 c worst-case. The total worst-case round-trip delay for all the transceiver cables is therefore 3.08 µs.

5. A maximum of 1000 meters of point-to-point link anywhere in the system is allowed. This will typically be used as a way of linking cable segments in different buildings. Note, that a repeater with this internal point-to-point link can be used to repeat signals between segments many hundreds of meters apart. The worst-case propagation velocity of the link cable is assumed to be 0.65 c; the round trip propagation delay for 1000 meters is 10.26 µs.

The maximum distance of 2.8 kilometers between stations is calculated using the configuration of three Ethernet trunk sections (1.5 kilometers), plus six transceivers (used at the source, sink and at the two repeaters) at 50 meters plus 1000 meters of point to point link.

As indicated earlier, the Ethernet specification is for the lowest two levels of the OSI (physical and data link layers). Products such as DNA, XNS, TCP/IP and SNA would be network packages that drive the total networking operation of the system. Ethernet controls the functions needed by the data link and physical link layers to transmit/receive data over the trunk.

C. ACCESS METHOD

It is important to know the method used by Ethernet to access the shared medium in order to understand how it delivers data. Since it is a shared medium, there must be some ways of allocating this single medium to multiple users. To perform this function, there are two options: polling and contention. Polling, uses a central controlling node which constantly interrogates each node if it has any data to transmit. With contention, each node acts as a controller for its on network access on the shared bus. Contention is the access method used in the Ethernet environment. The actual type of contention used by Ethernet is Carrier Sense Multiple Access with Collision Detection (CSMA/CD). There are three main protocals associated with CSMA/CD systems: 1-persistent, non-persistent and p-persistent. They are used to minimize the chance of collision on the shared medium (The 1-persistent method is used in Mitrenet and the IEEE 802 standard). This access method was created from the Aloha Network's initial packet switching ideas. In Hawaii, a way to connect computing resources on different islands was needed. A single radio frequency was used to transmit and receive network data. This single channel was accessed by sending data whenever data was ready to be sent. If no acknowledgement was received in a predetermined time interval, the data was retransmitted. Since, this method worked well, the general principle was altered to include an initial carrier sense before transmission yielding the 'CS' of the current access method. Under the revised technique, when any node is ready to transmit data, it must first listen to the network to see if anyone else is using it. The 'listen' function is actually a power sense. The transceiver knows how much power is generated on the trunk during a normal network transmission. If the network is busy, it waits until the line is free. When the line is free, it transmits its data. It then listens again. If it senses a collision, it transmits a jam signal to let everyone know there was a collision, and then halts its transmission. It then waits to retransmit its data. The amount of time it waits is determined by a binary exponential backoff technique. The more collisions that occur on the same data packet, the longer it waits for the next retransmission. This process continues until a successful transmission occurs or until an overload condition is determined.

Even thought the CSMA/CD technique is the key behind Ethernets movement of data through the network, much more takes place in the data link and physical link layers: data encapsulation/decapsulation, link management, collision handling, encode/decode, transmit and receive, collision detect and carrier sense. There is an interface between the data link layer and the network layer which will not be discussed. This interface allows the different network products to pass data to and from the Ethernet.

1. Layer 2 - Data Link Layer. This layer performs two major functions; data encapsulation/decapsulation and link management. In the encapsulation/decapsulation section, it creates the Ethernet packets needed for network transmission using data supplied by the upper layers. For received data, it checks for transmission errors and removes the Ethernet packet and sends the appropriate data to the upper layers. The link management section checks the physical link layer to see if the link is busy before sending a packet down for transport.

Figure 6 shows the Ethernet packet format. Using this predescribed format, Ethernet takes data sent from the upper layers, and creates a packet to be sent to the network. The preamble is an 8 byte (64 bit) section sent for synchronization by receiving nodes. It is actually generated at the physical link layer. It ends with a minimum spacing period of 9.6 ms to allow recovery time between packets and for other data link controllers and channels to stabilize. There are two types of destination addresses used in the packet: a physical address and a multicast address (determined by the first bit set to '1' or '0'). The physical address is the unique address of a node on any Ethernet. The multicast address specifies either a subset of all nodes or a broadcast packet to go to all nodes (for a broadcast packet all bits are set to '1'). The source address is incorporated from above. The type node specifies to the receiving node which high-level protocal & network architecture are being used with the data (Decnet, SNA, etc.). The data field can vary in size from 46 to 1500 bytes in length. The data sent from the network layer is placed here. Finally, there is a four byte cyclical-redundancy-check (CRC) frame check sequence appended to the end of the packet. The CRC is used to detect errors in the ether, but not to check against errors in the parallel portions of the interface hardware or station. This must be performed by high-level software checksums (when a high degree of reliability is needed) (Metcalfe, 1976). Its value is calculated using the contents of the other 4 fields. In Schoch and Hupp's (Schoch, 1980) experimental Ethernet, the number of packets damaged in transmission was 1 in 2,000,000. So the

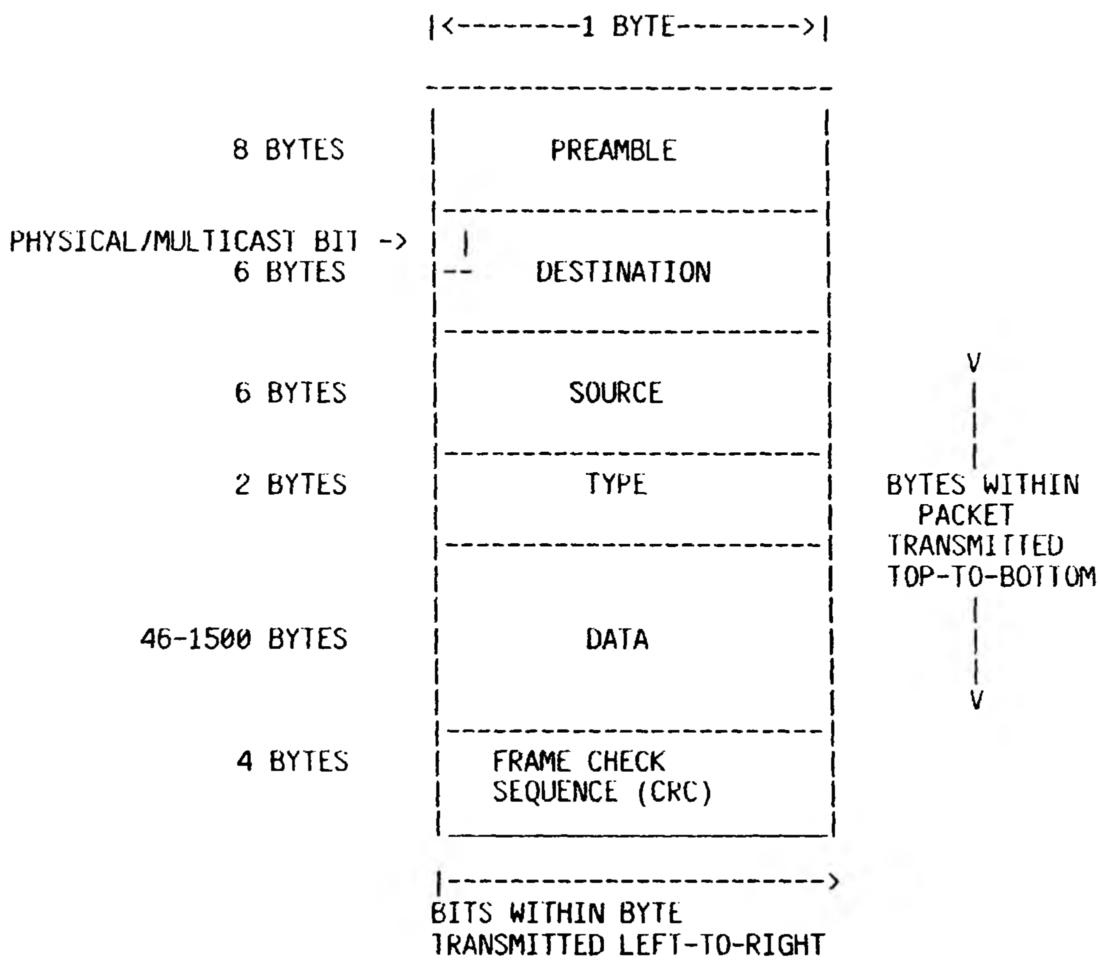

FIGURE 6 - ETHERNET PACKET FORMAT

CRC is not needed very often. The entire Ethernet packet can vary in length from 64 to 1518 bytes (preamble not included).

2. <u>Layer 1 - Physical Link Layer.</u> Two major functions are performed by the physical link layer. The primary function is to provide a 10 mbps channel over the coaxial cable. The physical link layer also performs the encode/decode of data and the channel access. The channel access is done using the CSMA/CD technique discussed earlier. First, a preamble is generated which is sent just prior to the transmission of the actual data frame. The encode/decode function takes the binary data ('1' and '0') and changes it into the Manchester Phase Encoding used by Ethernet on the coaxial cable. The Manchester Phase Encoding changes the bits in the data stream into the electronic pulses transmitted on the channel. It does a reverse of this process on the receiving end (decoding and removal of the preamble). By looking at the following examples of transmission/reception with and without collision, one can see how these ethernet functions are actually performed.

3. <u>Transmission Without Collision.</u> When the higher layer protocals need to transmit data, they will send information down to the data link layer. The information sent is source and destination addresses, protocal type used and data. The data link layer will build the Ethernet packet using this data. It will compute and append the CRC on the end. Then, it will check with the physical link layer to see if the medium is free. If not, it waits. When the medium is free, it will send the packet down to the physical link layer. The physical link sends a carrier down the link to alert everyone that it is sending data.

It then generates and sends a preamble to allow all nodes to synchronize their clocks. Next, it encodes the binary data and puts the manchester phase encoded data onto the channel. If the physical link detects no collision while transmitting, a success message is sent up to the higher levels and it waits for the next transmission. It is important to note that in the experimental Ethernet (120 concurrent stations) at Palo Alto (Schoch, 1980), approximately 98% of the transmissions/receptions occurred without collisions.

4. Reception Without Collision. On the receiving end, when the carrier signal is detected at each node, the physical link layer realizes data is in transit and alerts its data link layer that data is on the way. In a sense, each node becomes a receiving station. When the data arrives, the physical layer used the preamble to synchronize and then discards it. Then, the packet is decoded as it arrives and the data is passed up to the data link layer. The destination address portion of the packet is read. If the packet is not intended for the node, the physical link layer stops receiving. If the packet is intended for the node, it is inspected for damage and alignment. The data is then passed to the higher levels.

5. Collisions. There is only a very short time period in which two nodes can both send a packet and cause a collision. Suppose nodes A and B are on opposite ends of the ether. When node A sends a packet down the ether, it will take the one-way propagation delay time for the packet to arrive at node B. If node B would send a packet just before node A's packet arrived, a collision could occur. Node B would sense

the collision almost instantly. But, it will take another one-way propagation delay time before node A will realize two packets have collided. Thus, the maximum time required for any node to realize a collision would be 2T (T is the maximum one-way propagation delay between the two farthest nodes on the network). This time period is called the collision window or slot time and is the only time a collision can occur. The slot time for a particular network is thus a function of the physical transmission medium and its bus length. For a 1000 meter trunk line at 3 mbps the slot time is 28.86 µs while for the 2.8 kilometer trunk (the maximum distance configuration allowed between nodes), using a 10 mbps channel, the slot time is 26 µs. The maximum slot time (2T) is just less than the time needed to send 64 bytes (minimum packet size) one-way down the Ethernet. If a collision occurs during transmission, a packet of less than 64 bytes (a runt packet) will be transmitted. Seeing this, the receiving station will know that a collision has occurred. This detection is done by the data link layer.

The transmitting station notices a collision by monitoring the power level of the coaxial cable. It knows that if the level is greater than that of a normal transmission, then two (or more) stations are transmitting at once. When it determines a collision has occurred, it continues transmitting for a fraction of time (a jam) to let all stations know a collision has occurred, then it stops transmission. Next, it begins backoff procedures. The binary exponential backoff algorithum, used by Ethernet for rescheduling transmissions, waits a random amount of time. The mean waiting time is doubled every time the transmitting station experiences another collision on the same packet. This algorithum has the advantage of being fair to all nodes on the net since it is executed by all. It continues the backoff procedures until the packet has been successfully sent or an overload has been determined.

IV. ETHERNET PERFORMANCE

The concept of "network management" is becoming very popular in this decade. Network security, reliability and performance all fall into the scope of "network management". Thus far, it can be seen from both the software and hardware side, how Ethernet gets data to its destination. Whether the network is installed or is still in the design stages, it is important to be able to calculate the Ethernet network's performance. In the next few sections, issues of Ethernet network performance will be examined by using analytical modeling techniques and by examining actual measured performance. This will lead to design and management methodologies for Ethernet networks

A. WHAT TO MEASURE?

One finds that depending on the role in network management: user, network manager or designer; one may have a different idea about what is needed from the network. In general, these three groups will have some network needs that overlap. But, they will also have needs that contradict one another. What statistics should be measured in order to satisfy all three? One thing that should be mentioned is that no matter what performance criteria an individual starts with, it eventually takes the shape of a "cost-performance tradeoff" (Ilyas, 1985). The Micom-Interlan Corporate Hierarchical model (Figure 2) can again be used. Micom-Interlan says that the work group (user), department (network manager) and the building (designer) all have different views of a network. The user is interested in his applications and how they are serviced. The network manager is concerned with total network serviceability. The designer is concerned with the actual wiring of the network.

The user, not being a network expert, is going to be more interested in things he can see in the network. The first thing he will notice is how friendly it is. If it is not easy to connect to the network and perform his desired functions, he may shy away from it. Also, he will be interested in its visible speed (how quickly information gets from one place to another or delay per message). He wants his desired application to run fast. When the network user is affected by the network's operational cost, he may sacrifice some service to save some dollars (cost-performance tradeoff). For the most part, the needs of the user pertain to the higher layers of the OSI model.

The network manager looks at the network a little differently than the user. He still has a "cost-performance tradeoff" in mind; wanting the maximum utilization of resources. His best test of maximum but fair utilization of the network is throughput. Ilyas and Mouftah (Ilyas, 1985) suggest that this may not be the best performance calculation. With high throughput, the delay per message may also be high, making users unhappy. A calculation of POWER is suggested instead. POWER is the ratio of throughput to its delay. By increasing throughput or decreasing delay, a network's POWER value will be increased. If it is optimized, both user and manager should be happy.

Finally, consider the designer's needs. Being the most technical of the three, he wants to make sure the actual performance of the network is close to the designed one. He also wants to make sure the design goals are being met. He is interested in different performance criteria such as buffer efficiency, protocal efficiency and flow control effectiveness. He uses the various performance statistics to fine tune the network to fit the desired results.

The following standard criteria will be used for measuring the performance of an Ethernet network: offered load, utilization, efficiency and average packet transmission delay. The offered load is the amount of traffic put onto the network by the stations connected to it. One would like to know the types and amounts of traffic volumes that reside on the network. Predicting offered load will help calculate network traffic capacity in order to support varying numbers of users. Offered load assumes transmission without collision. A network's effective load will include the extra retransmitted collision packets. The utilization or throughput (the percentage of time good packets are being carried on the channel) should also be measured. This usage should be fairly distributed to all nodes on the network. The efficiency is the performance that can be expected under heavy load conditions. It will show whether a network configuration can remain stable under excessive load volumes. The average packet transmission delay calculation will show what effect load and slot time have on response time (delay). Analytical modeling techniques help predict these network performance statistics.

B. ANALYTICAL MODELS

In the design phase, one would like the ability to estimate some measures of Ethernet network performance. These results can help designers to refine their network configuration designs. They can also be used to help check and fine tune an existing Ethernet. Analytical models exist that can calculate such network values as offered load, throughput and efficiency. By supplying analytical models with values from actual Ethernet measurements, one should be able to predict how an Ethernet installation will perform under varying load conditions. These predicted values can then be compared against values from measured Ethernet performance tests. Finally, these results should allow one to make some final conclusions concerning Ethernet network performance.

The first area to examine will be the calculation of offered load. Offered load is the amount of traffic put onto the network by all N stations queued on the network. With offered load, assume all traffic is put onto the network without collisions (a best case scenario). The offered load may be calculated as follows:

$$\text{Offered Load} = \Sigma\ (C*Ts*E(Pk_i)*Lk_i) \text{ for } i=1 \text{ to } N$$

where;

N = number of stations on the network,

C = transmission rate on the ether (i.e. 10 mbps),

Ts = slot time,

C*Ts = bit transmission time in bits per packet (i.e. 64 bytes per slot time),

E(Pk) = mean packets per message, and

Lk = mean message arrival per second.

Using approximated values from the time sharing environment study (Digital, 1982), one can predict the offered load for a time sharing installation like the Emerson Electric Ethernet. The results are represented by graphs like the one shown in Figure 7. The predicted

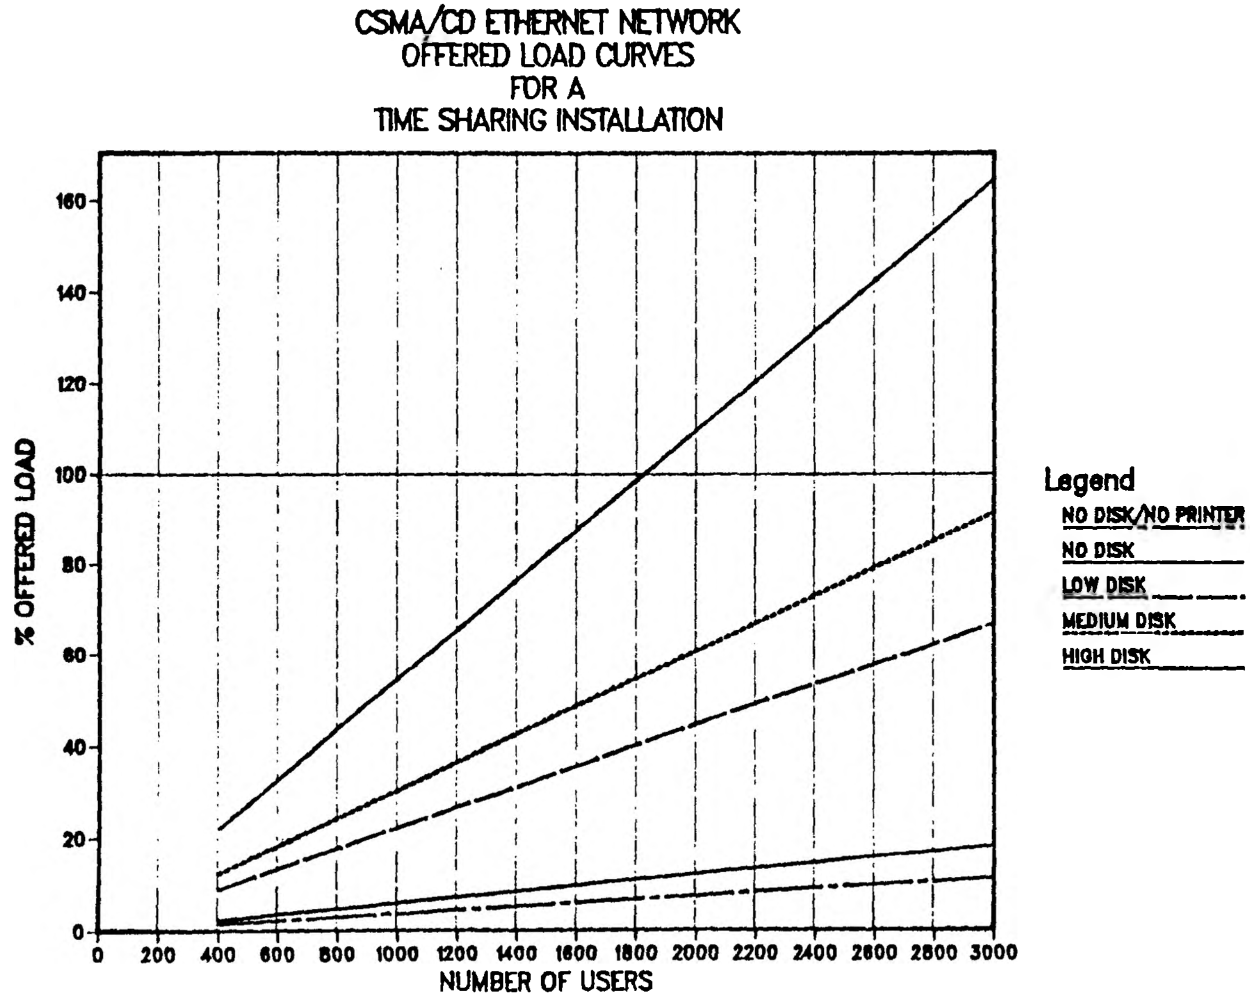

FIGURE 7 - PREDICTED OFFERED LOAD CURVES

data will be compared with measured offered load data to see how closely the analytical modeling predicts actual results. The results will be discussed in detail in IV.C.

To examine the concept of predicting utilization or throughput. First look at the Aloha Network Utilization Curve such as the one in Figure 8. The low channel utilization value and unstable utilization both indicate why the original Aloha line access method had to be made more efficient. The analytical model for measuring the Aloha utilization as presented by Reiser (Reiser, 1982) is:

$$L = L1*e^{-2*L1}$$ where L1 is the effective load.

The effective load is the higher channel load value which takes into account the added packet retransmissions due to collisions. The Aloha Network Utilization Curve (Figure 8) shows that after the Aloha utilization reaches its maximum (18.4%), it drops off. This shows that as the effective load on the network becomes greater than 100%, the network becomes unstable. Overall, the Aloha network yields both low and unstable utilization results.

Section III.C stated that the CSMA/CD access method was spawned from the original Aloha concept. Its better line utilization stems from its carrier sense function. With the Aloha method, each station transmits onto the shared medium whenever it has a packet to transmit whether or not there is a packet already in transmission. Thus, a collision can happen at any time. The carrier sense, in the Ethernet

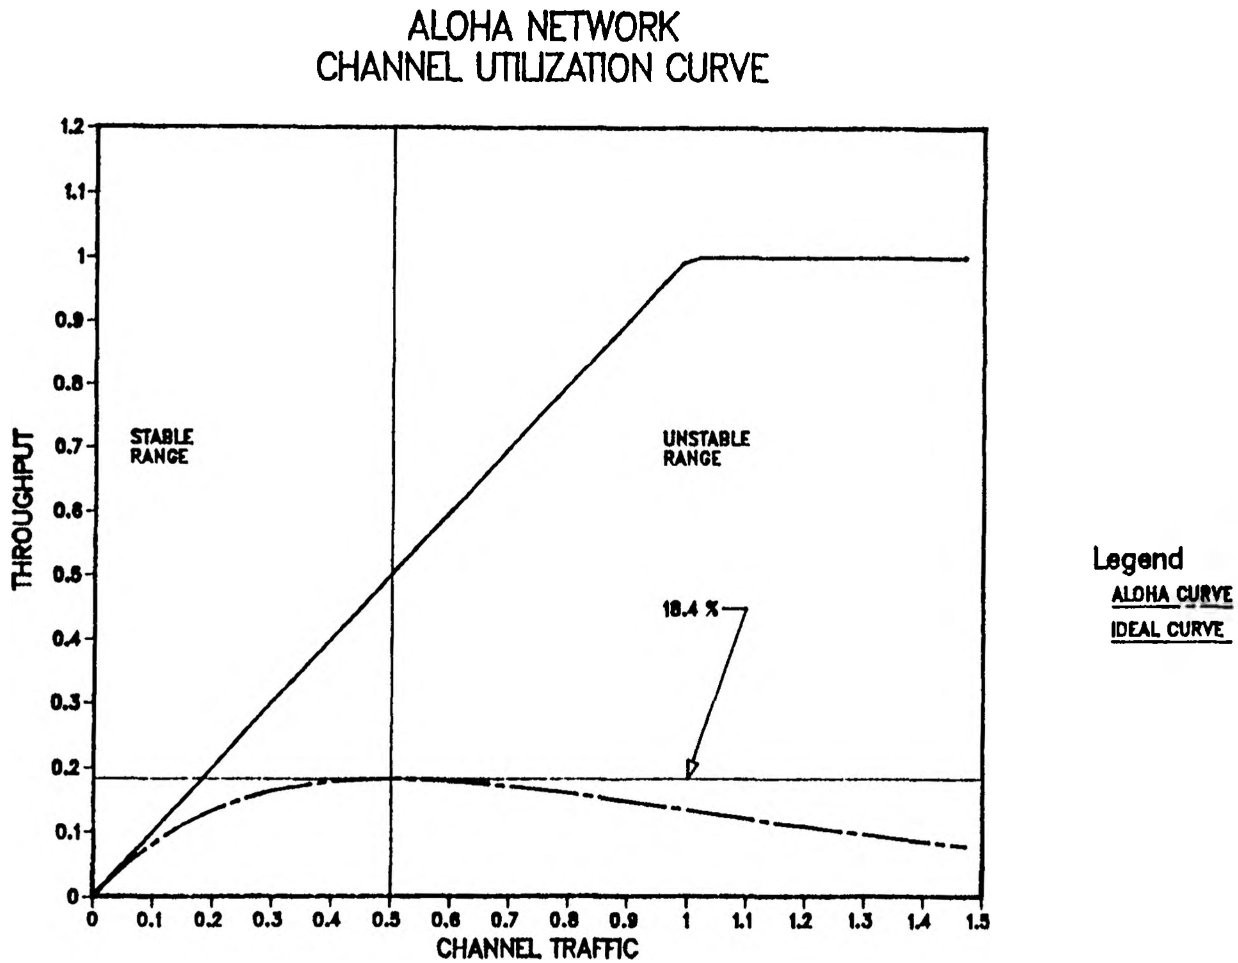

FIGURE 8 - ALOHA NETWORK PERFORMANCE CURVE

CSMA/CD network, only allows a station to transmit when the line is free. This makes the collision window very small (equal to the round-trip propagation delay of the minimum packet size 64 bytes). The smaller collision window increases the channel utilization.

Five different protocals are commonly used with the CSMA/CD shared medium access method: non-persistent, slotted non-persistent, 1-persistent, slotted 1-persistent, and slotted p-persistent. They are based on the premise that collisions can not be avoided, but their probability of occurring can be minimized. With p-persistent slotted CSMA, the user sends data with probability p whenever an empty slot is found, and postpones it to the next slot with probability q=1-p. In the next slot, the same algorithm is repeated again. With non-persistent CSMA, if the user finds the channel busy, he reschedules his next retransmission after a random duration of time. In the following CSMA throughput (percentage of time good packets are being sent over the channel) equations (Ilyas, 1985), a is the normalized propagation delay and G is the effective load.

Non-persistent CSMA:

$$S = (Ge^{-aG})/(G(1+2a)+e^{-aG});$$

Slotted non-persistent CSMA:

$$S = (aGe^{-aG})/(1+a-e^{-aG});$$

1-persistent CSMA:

S = Numerator/Denominator where

$$\text{Numerator}= G[1+G+aG(1+g+ag/2)]e^{-G(1+2a)}$$

$$\text{Denominator}= G(1 +2a)-(1-e^{-aG})+(1+aG)e^{-G(1+a)};$$

Slotted 1-persistent CSMA:

S = Numerator/Denominator where

$$\text{Numerator} = (Ge^{-G(1+a)})[1+a-e^{-aG}]$$

$$\text{Denominator} = (1+a)(1-e^{-aG})+ae^{-G(1+a)};$$

Slotted P-persistent CSMA (zero propagation delay):

$$S = (Ge^{-G})[1+pGX]/(G+e^{-G})$$

$$X = \Sigma\ [(qG)^k/(1-q^{k+1})k!]\ \text{for } k=0 \text{ to infinity.}$$

These calculations will be used with the data from the Xerox Ethernet (Schoch, 1980) to predict expected utilization of the different CSMA/CD variations. The values will be compared for accuracy against measured utilization values from the Xerox test site.

Finally, a method for predicting Ethernet performance under heavy loads (a worst case scenario) is needed. This calculation is called efficiency. The network model assumed here is one that contains a number of constantly queued stations, each transmitting at 100% capacity (packets are transmitted one at a time continuously). The network offered load for this model will be 100*Q where Q is the number of stations queued on the network. This calculation will show if an Ethernet's channel utilization remains stable (does not drop) as the channel bandwidth is saturated.

Three equations will be needed to predict the efficiency of an Ethernet network (Metcalfe, 1976). The first equation is the acquisition probability A, or the probability that exactly one station attempts a transmission in a slot and therefore acquires the ether:

$$A = (1-1/Q)^{Q-1}$$

where Q is the number of constantly queued stations on the network.

Next, the A calculation is used to calculate W, the mean # of slots waiting in the contention interval before a successful acquisition of the ether by a station's transmission:

W=(1-A)/A   Note: A is the probability of waiting no time.

Finally, the efficiency of the network can be predicted using the equation.

$$E = (P/C)/((P/C) + W*T)$$

where,

T = slot time in seconds,
C = the transmission speed of the shared medium, and
P = avg packet size on the network.

Using the above calculation and the specific constants from each installation, the efficiency of both for the Emerson Electric and Xerox test Ethernet can be computed. One can see what effect the number of stations, channel transmission speed and average packet size all have on the efficient channel bandwidth utilization of an Ethernet network. The predicted results will also be compared to the measured results from Schoch and Hupp's (Schoch, 1980) Xerox test.

Mitchell (Mitchell, 1981) discusses the delay components associated with end-to-end response time in a LAN. The five components that contribute to delaying the user's response are applications processing, nodal communications software, node to Bus Interface Unit (BIU), BIU communications software/hardware/firmware and transmission plant. The delay he associates with the transmission plant considers the capacity of the transmission media, the access method overhead and distance between nodes. A methodology is given which includes delay components for the CSMA/CD access method. The average packet transmission delay, D, is given (Mitchell, 1981) by:

$$D= (A_1+A_2+A_3)T,$$

where A_1 is the normalized waste time due to collisions, A_2 is the dead time due to retransmissions and rescheduling and A_3 is the propagation and transmission time. A_1, A_2 and A_3 are computed using the following equations:

$$A_1= N_2(W+a);$$

$$A_2= R_2(2^{N_2+1}-1)+(N_1-N_2)R_1;$$

$$A_3= a+1;$$

where

$$W= (1-e^{-aG})/(G) - ae^{-aG};$$

$$N_1= G/S-1;$$

$$N_2= (1+aG)e^{aG}-1;$$

with N_1 representing the average number of times a packet encounters a collision or busy state and N_2 being the number of times a packet encounters a collision. R_1 and R_2 are the normalized mean retrial intervals after detecting a busy condition and a collision respectively. Also, a different equation for normalized throughput is used by Mitchell. Normalized Throughput, S, is:

$$S = \text{Numerator}/\text{Denominator} \quad \text{where}$$

$$\text{Numerator}= Ge^{-aG}$$

Denominator= $(!+a)Ge^{aG}+(1+aG)(1-e^{-aG})^2+1$,

a is the normalized propagation delay (slot time) and G is the offered load. The A_1, A_2 and A_3 equations will be used later to see how different propagation delays affect the performance of the transmission plant.

C. ETHERNET PERFORMANCE: MODELED VERSUS MEASURED

The analytical modeling techniques presented previously will be used to predict values for offered load, throughput and efficiency. The necessary inputs (slot time, channel capacity, number of stations, etc.) will come from the configurations of the Ethernet networks at Emerson Electric Engineering and Space Division (see appendix A) and at the Xerox Palo Alto research facility. The following static inputs from the two Ethernets will be used in the computations:

	Emerson	Xerox
Channel Capacity:	10 mbps	2.94 mbps
Total Stations:	379	120
Slot time:	10 μs	32.53 μs

Two Ethernet performance studies will be used as comparisons for our predicted values. A brief description of the motivations and specifics of each study follows.

The first study (Schoch, 1980), was done at the Xerox Palo Alto facility. At the time of this study (1980), Ethernet though new, was already one of the largest used LAN designs. The study was intended to further the understanding of the actual behavior of an Ethernet installation under normal and heavy usage.

The study measured traffic volumes, packet sizes, error rates, and efficiency. Schoch and Hupp pointed out the importance of high level protocals on network performance but restricted their study to the behavior of the Ethernet itself. Many papers have referenced the results of this study to validate analytic and simulation models of Ethernet systems.

The second, "Predicted Capacity of Ethernet in a University Environment" (Digital, 1982), was motivated by the desire to find a limit to the number of users imposed by the finite bandwidth of the ethernet channel. The Digital, et al. study was performed for users in a time-sharing environment. The study also noted that the higher layers often dictate the performance of the network and agreed they should be carefully studied. The high level protocals will produce extra control packets that may contend with user packets on the channel. The higher levels may also contend with network resources (CPU cycles and memory) at the transmitter and receiver. The study also only addressed the issues that relate to the shared channel. Since this study was done for a time-sharing environment like Emerson Electric, traffic data (average input/output rate and size, etc.) along with Emerson data (channel capacity and slot time) will be used to help predict offered load. The results of this study will give statistical insight into average packet size, offered load, mean waiting time per message, mean number of transmission attempts and total number of users that can be effectively using an Ethernet network.

The values used to predict offered load in a time sharing environment come from the Emerson Electric network and the Digital, et al. study. The load per station for different volumes of network traffic is given below. These values are used to predict the network offered load for varying numbers of stations on the network. The network load is broken into five traffic volume types represented by 1 curve each in Figure 7. The data was taken from the Digital, et al. study

OFFERED LOAD TRAFFIC TABLE

	TERMINAL LOAD	PRINTER LOAD	DISK LOAD	TOTAL LOAD PER STATION
CURVE 1	37.2 bps	-	-	37.2 bps
CURVE 2	37.2 bps	23.28 bps	-	60.48 bps
CURVE 3	37.2 bps	23.28 bps	162.57 bps	223.05 bps
CURVE 4	37.2 bps	23.28 bps	243.7 bps	364.67 bps
CURVE 5	37.2 bps	23.28 bps	487.24 bps	547.9 bps

When looking at Figure 7, note that without any disk traffic, the load offered by even as many as 3000 stations is small (less than 20%). The curves show that it takes 1800 stations at high loads to saturate the network.

By comparison, the two studies, (Digital, 1982) and (Schoch, 1980), give the following information on offered load. The Digital, et al. study (Figure 9) shows results similar to those predicted in Figure 7. The highest offered loads shown are from 60-70% with a 10-15% variation between the load for high disk and low disk traffic. Schoch and Hupp's (Schoch, 1980) network had a daily traffic load of 2.2 million packets (300 million total bytes transmitted). The mean inter-packet arrival time was 39.5 ms and the mean packet length was 122 bytes. The offered load per station was 3128 bytes/second (25024 bits/second). This value is a little higher than those seen before. However, the traffic on Schoch and Hupp's network was very disk intensive. Overall, the average load recorded on this system was 0.8%. The maximum loads seen were 3.6% for an hour, 17% for a minute and 37% for a second (Figure 10). Coinciding with offered load, both studies show that the average packet

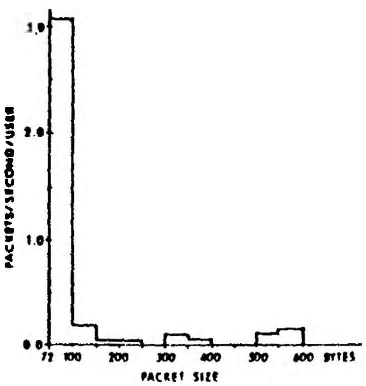

"Per-User" Ethernet Packet Size Frequencies
(Low Remote Disk Traffic)

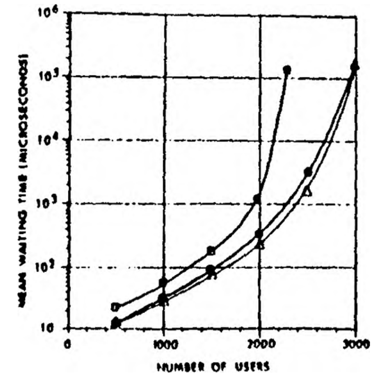

Mean Waiting Time

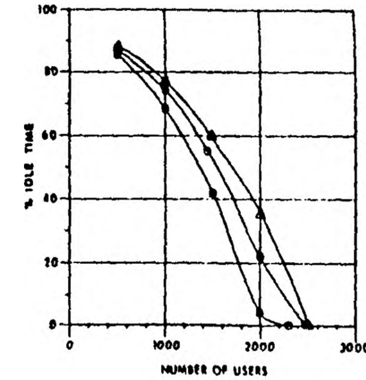

% Idle Time On Cable

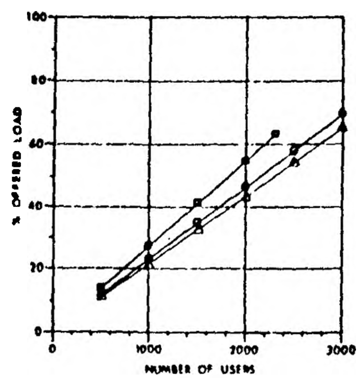

% Offered Load

HIGH NETWORK DISK TRANSFERS
MEDIUM NETWORK DISK TRANSFERS
LOW NETWORK DISK TRANSFERS

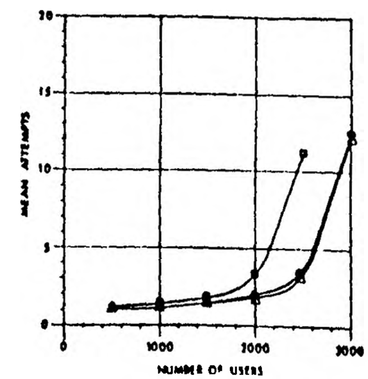

Mean Number of Attempts

HIGH NETWORK DISK TRANSFERS
MEDIUM NETWORK DISK TRANSFERS
LOW NETWORK DISK TRANSFERS

FIGURE 9 - (DIGITAL, 1982) PERFORMANCE CURVES

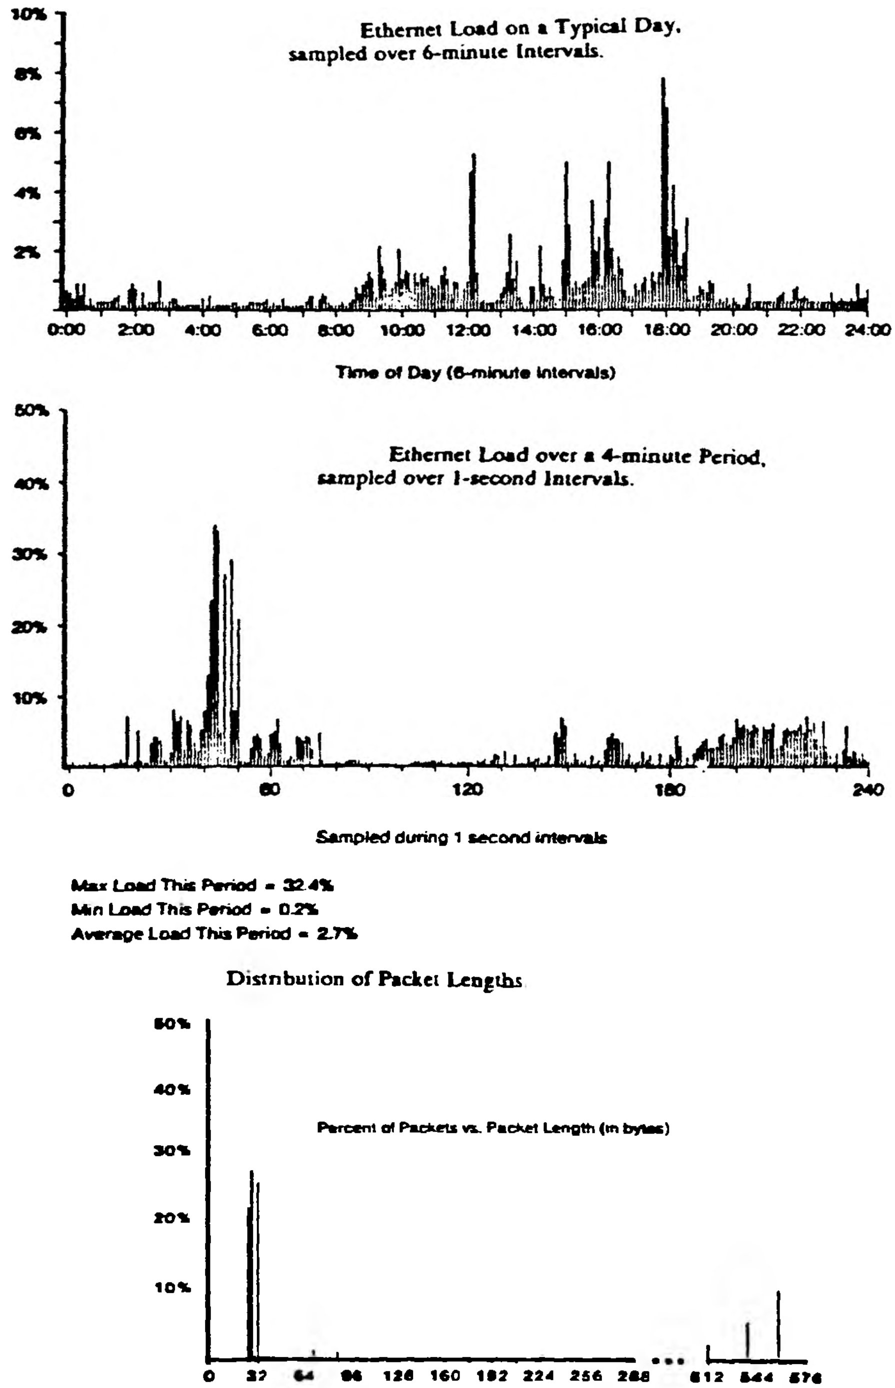

FIGURE 10 - (SCHOCH, 1980) NORMAL LOAD PERFORMANCE CURVES

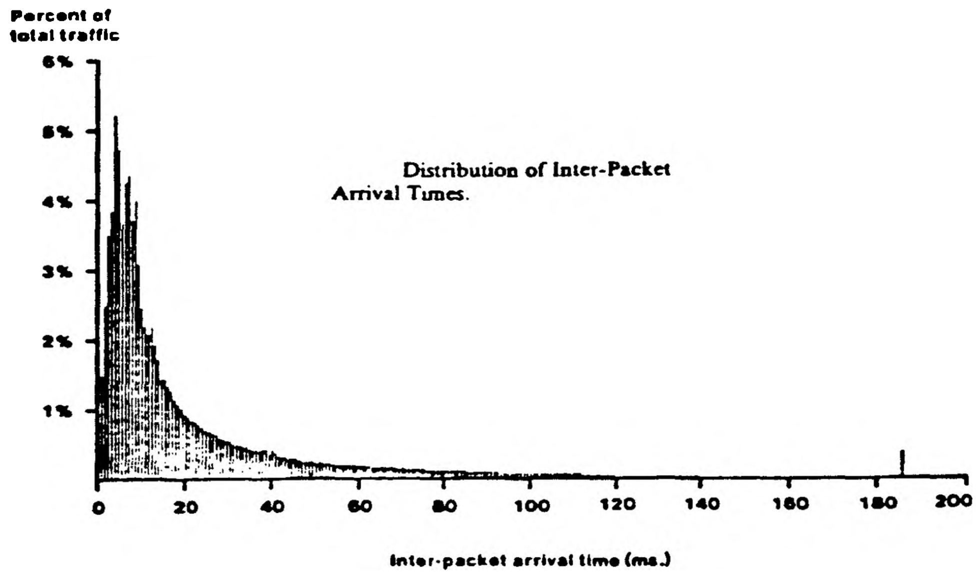

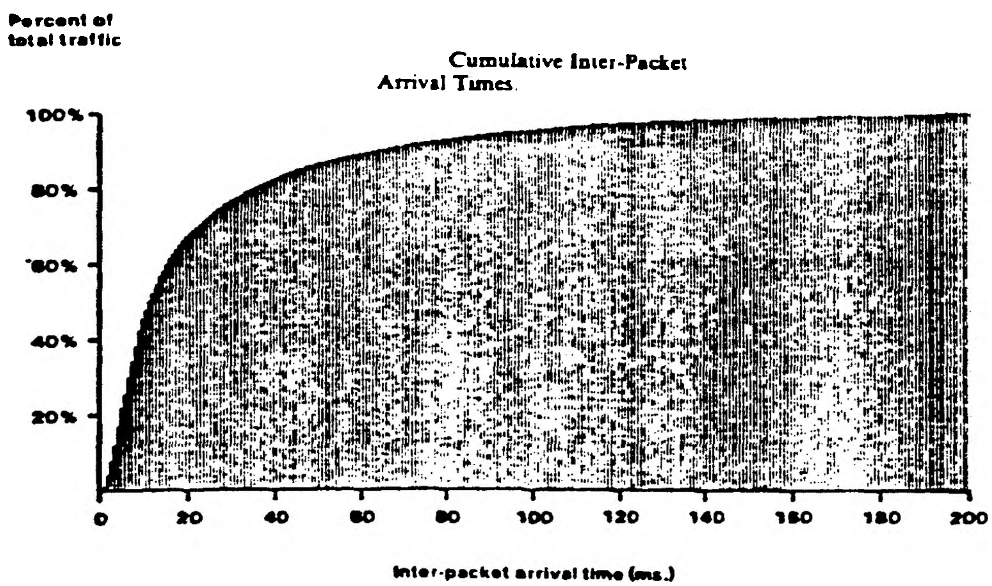

FIGURE 10 - CONTINUED

size hovers in the 64 to 130 byte per message average. The low byte per packet value can be explained by the terminal traffic. Since users can not type fast enough, most packets are sent with data fields of the minimum size (yielding minimum size packets). The Xerox study noted that even though most packets were of a small average byte size, the majority of the volume transmitted on the network was sent via large sized packets (around 512 bytes).

The throughput predicted values can be seen in the following four graphs (Figures 11-14). The predicted values are low and some instability can be seen for both slotted and non-slotted 1-persistent CSMA/CD. Also, it can be noted that whether a slotted or non-slotted CSMA technique is used,
it makes almost no difference in the predicted utilization values. Figure 15 shows the measured results from the Schoch and Hupp study including the measured throughput for high loads. Not that the measured throughput looks different than the predicted values. As the offered load increased, the throughput increased. It leveled off after about 90% offered load and remained stable. Depending on the average packet size, the throughput approached between 83% and 96% channel utilization. The network remains stable at offered loads as high as 150%. For the purposes at hand, it will be assumed the results from the measured study are accurate. For network configurations different from the Xerox test site, similar throughput results to those of Schoch and Hupp will be assumed.

Finally, the network efficient utilization for an increasing number of continuously queued devices was predicted. In this calculation, each station described is generating an offered load equal to 100% of the

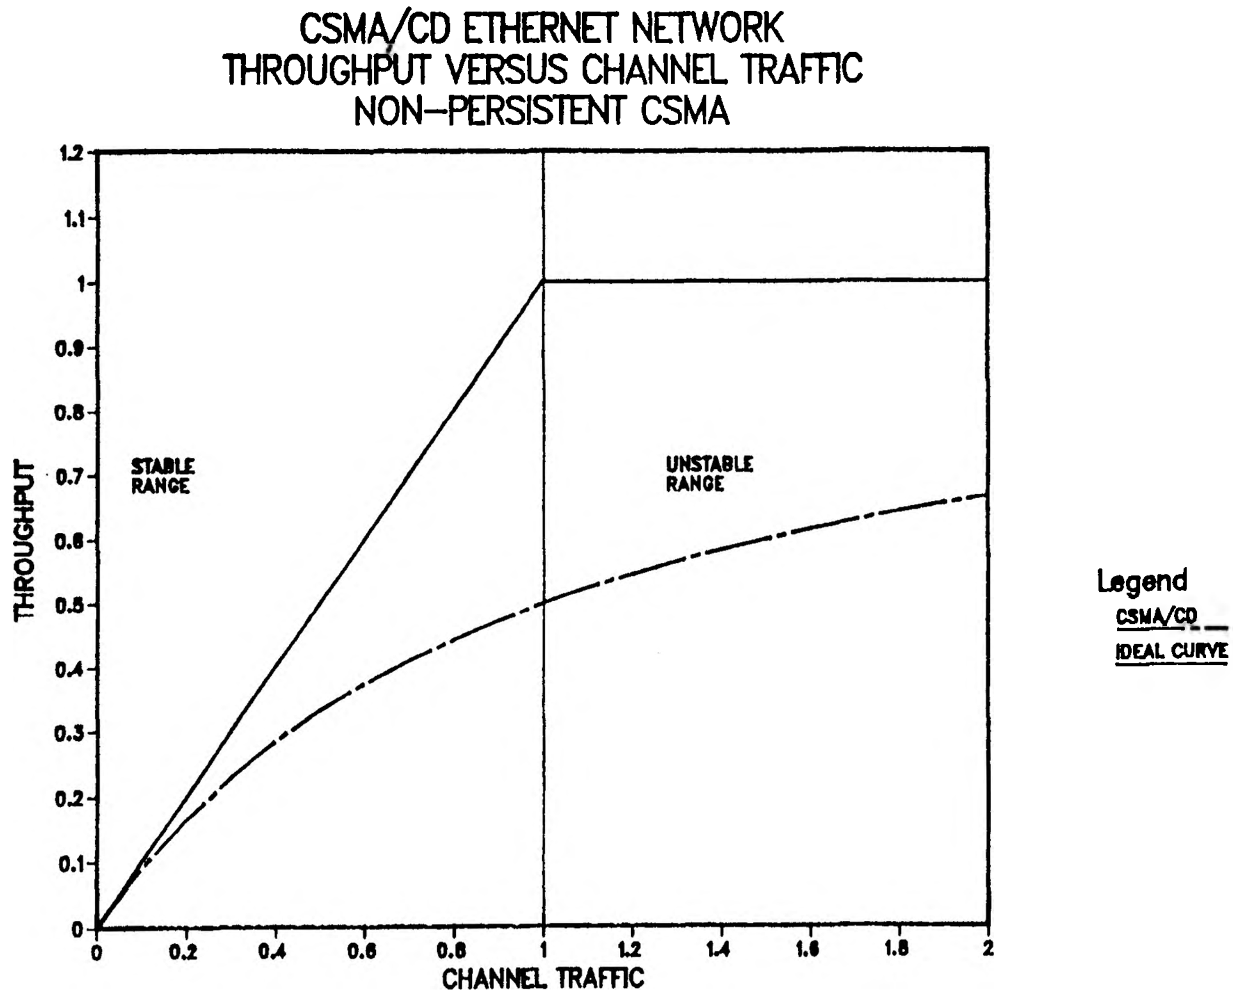

FIGURE 11 - NON-PERSISTENT CSMA/CD THROUGHPUT CURVE

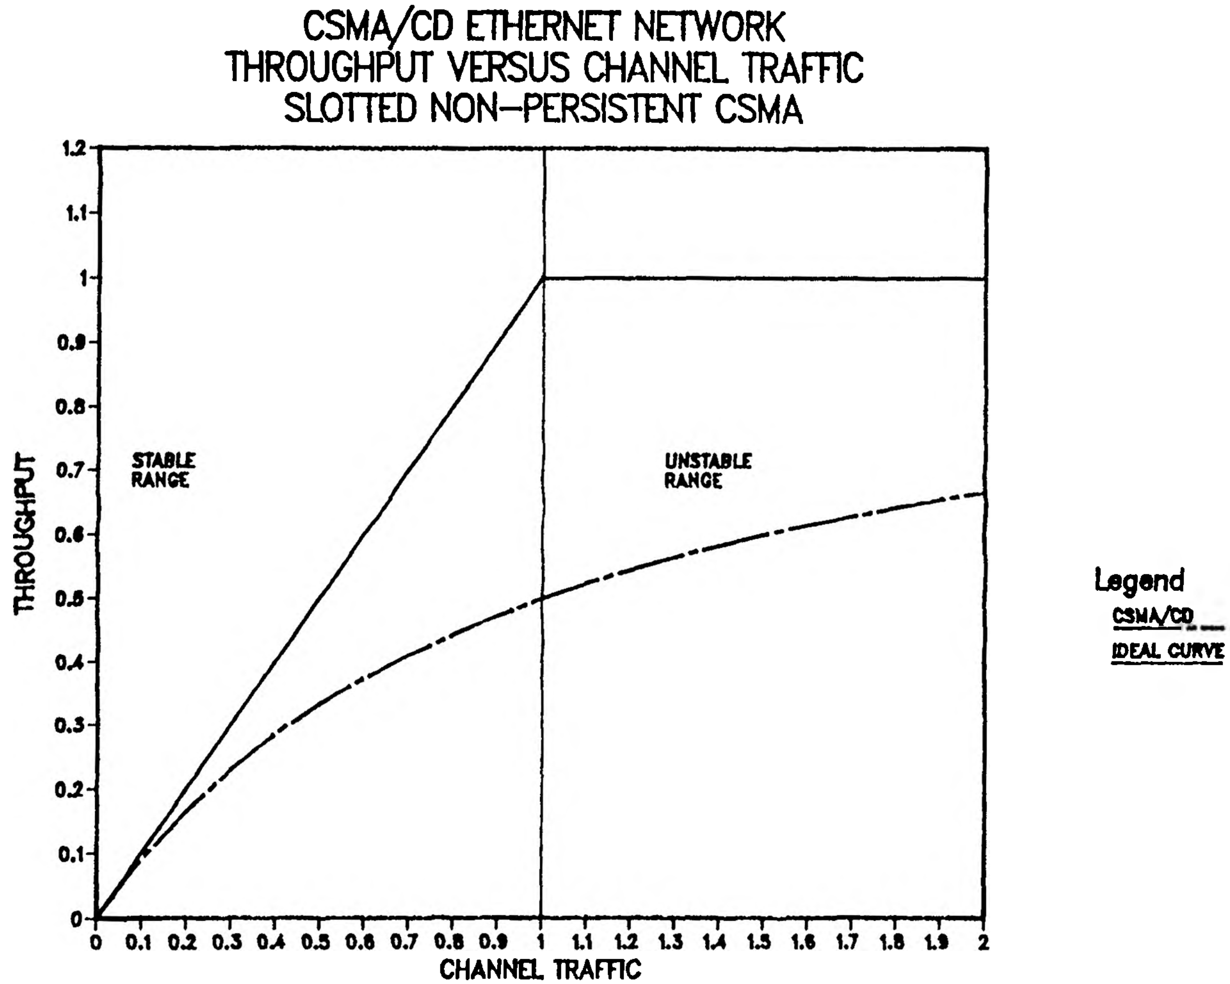

FIGURE 12 - SLOTTED NON-PERSISTENT CSMA/CD THROUGHPUT CURVE

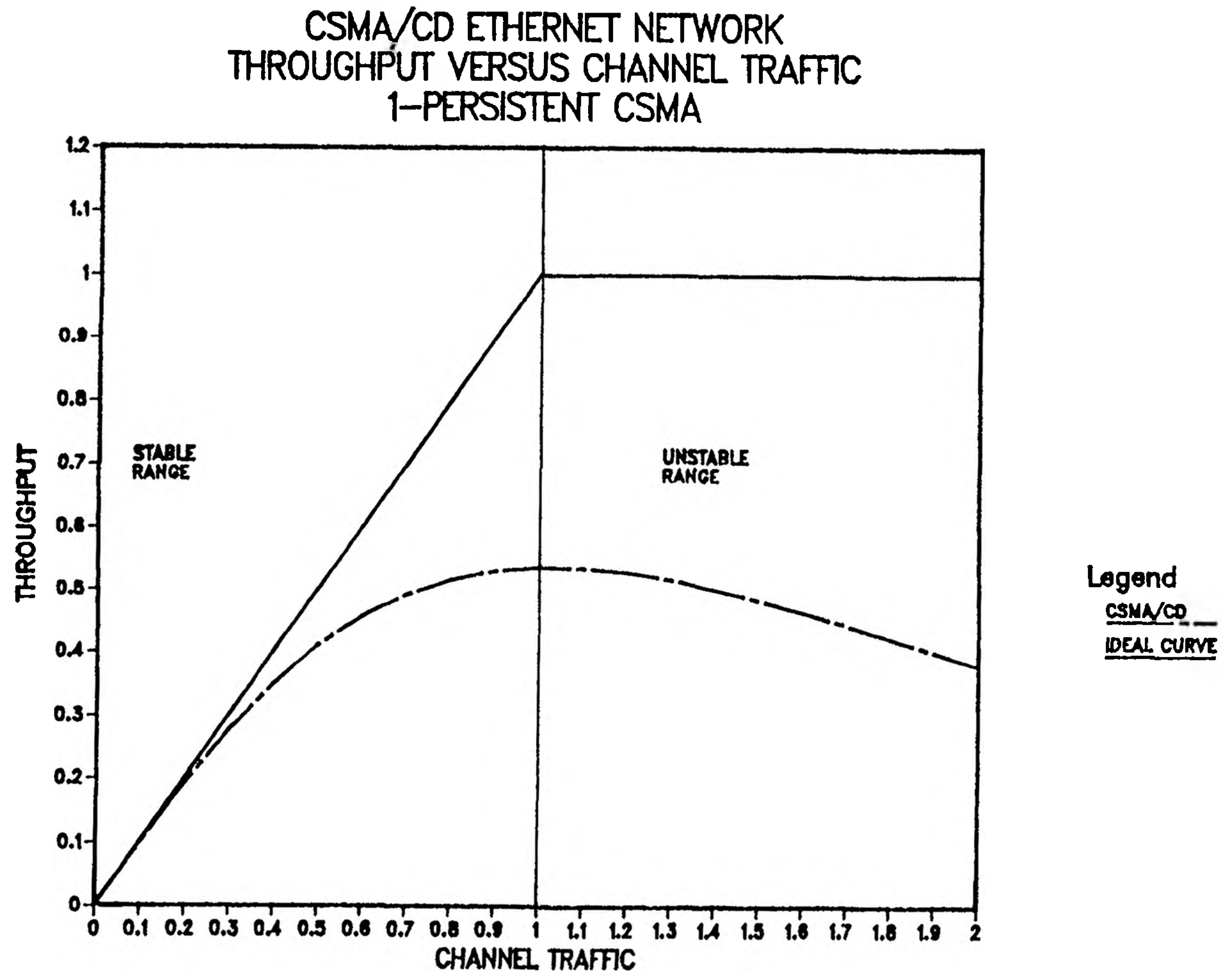

FIGURE 13 - 1-PERSISTENT CSMA/CD PERFORMANCE CURVE

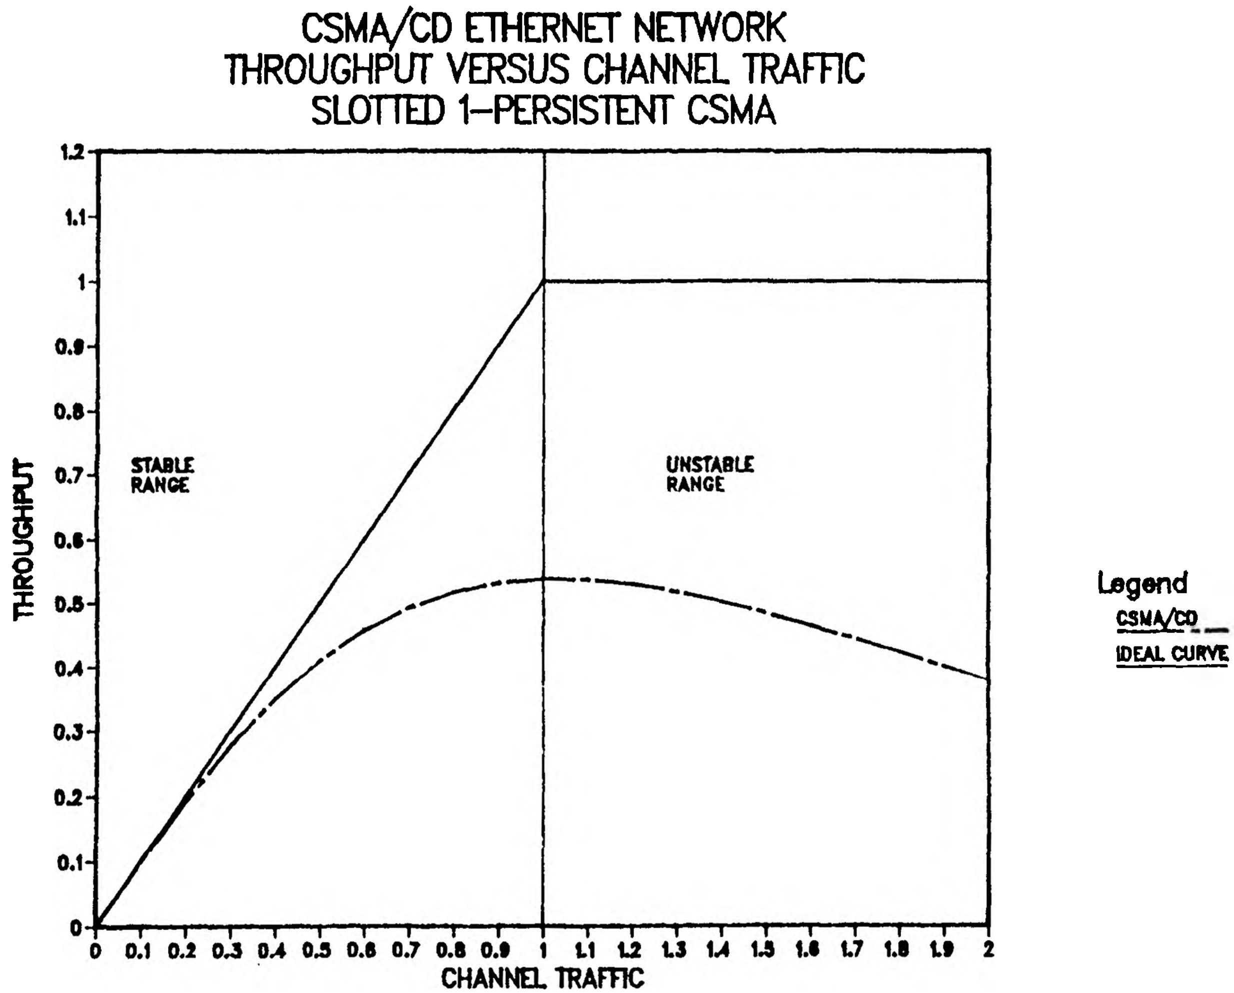

FIGURE 14 - SLOTTED 1-PERSISTENT CSMA/CD THROUGHPUT CURVE

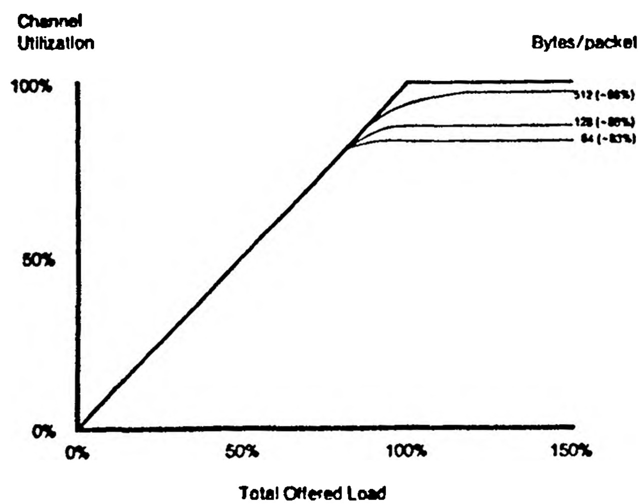

Measured Utilization of the Ethernet Network under High Load.

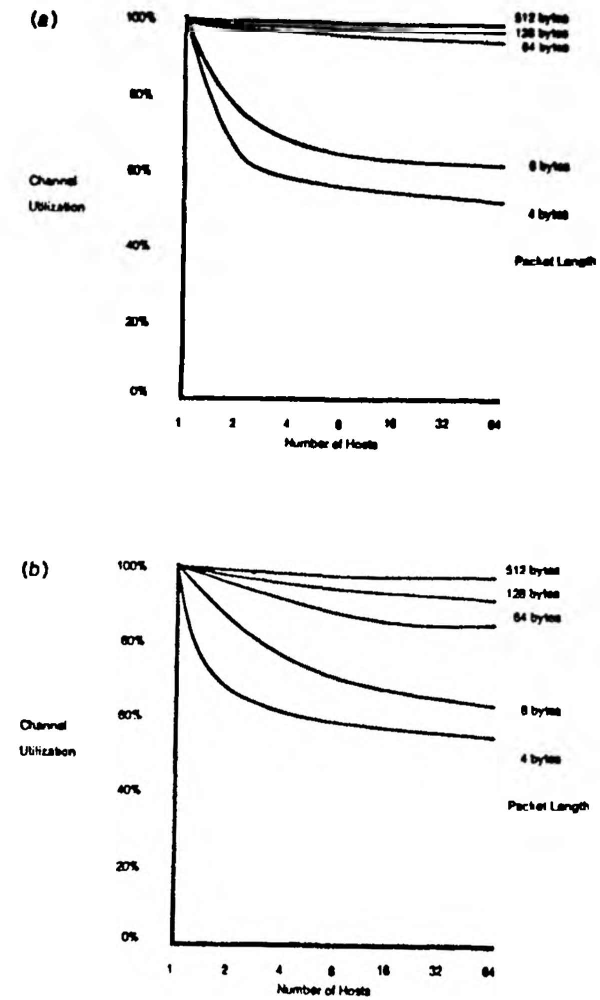

FIGURE 15 - (SCHOCH, 1980) HIGH LOAD PERFORMANCE CURVES

ether's channel bandwidth. The predicted efficiency values for the Emerson and Xerox network can be seen in Tables 3 and 4 and Figures 16 and 17. Contrast this with Schoch and Hupp's predicted and measured utilization calculations shown in Figure 15. The predicted values for the Emerson and Xerox networks closely match Schoch and Hupp's predicted and measured values. If the slot time for the Xerox network is increased without changing the channel capacity, the predicted efficiency values (Table 5 and Figure 18) decrease. This shows that by decreasing a network's slot time, the performance should increase. Overall, since the predicted efficiency values accurately match Schoch and Hupp's measured results, it will be assumed the given efficiency calculation is an accurate measure for predicting the efficient bandwidth utilization of the channel during high loads. It shows that the utilization of an Ethernet network should not decrease (remains stable) as the load generated on the network becomes very high. The efficiency calculation will be used later to help optimize Ethernet network design.

The last area of Ethernet performance to cover is the average packet transmission delay, D. Analytical models have previously been presented which may be used to model this calculation. Here, the three delay components associated with the CSMA/CD access schema, (A_1, A_2 and A_3), will be calculated and examined. The following inputs are used in this performance prediction:

R_1 and R_2 = 5 ms and
a (propagation delay) = 1000 μs and 1 μs.

The two values for normalized propagation delay are used to examine

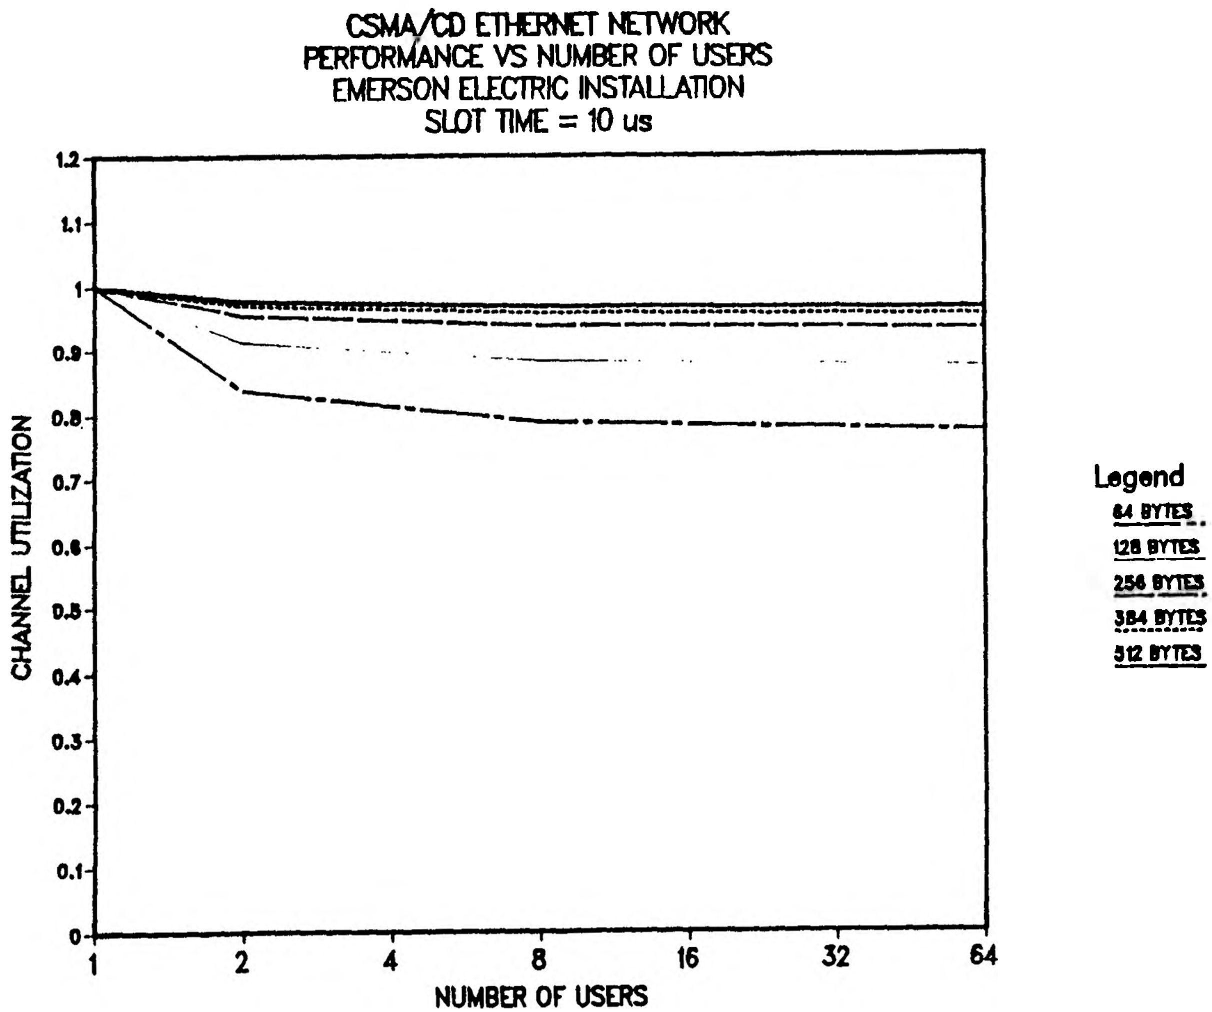

FIGURE 16 - EMERSON EFFICIENCY CURVE

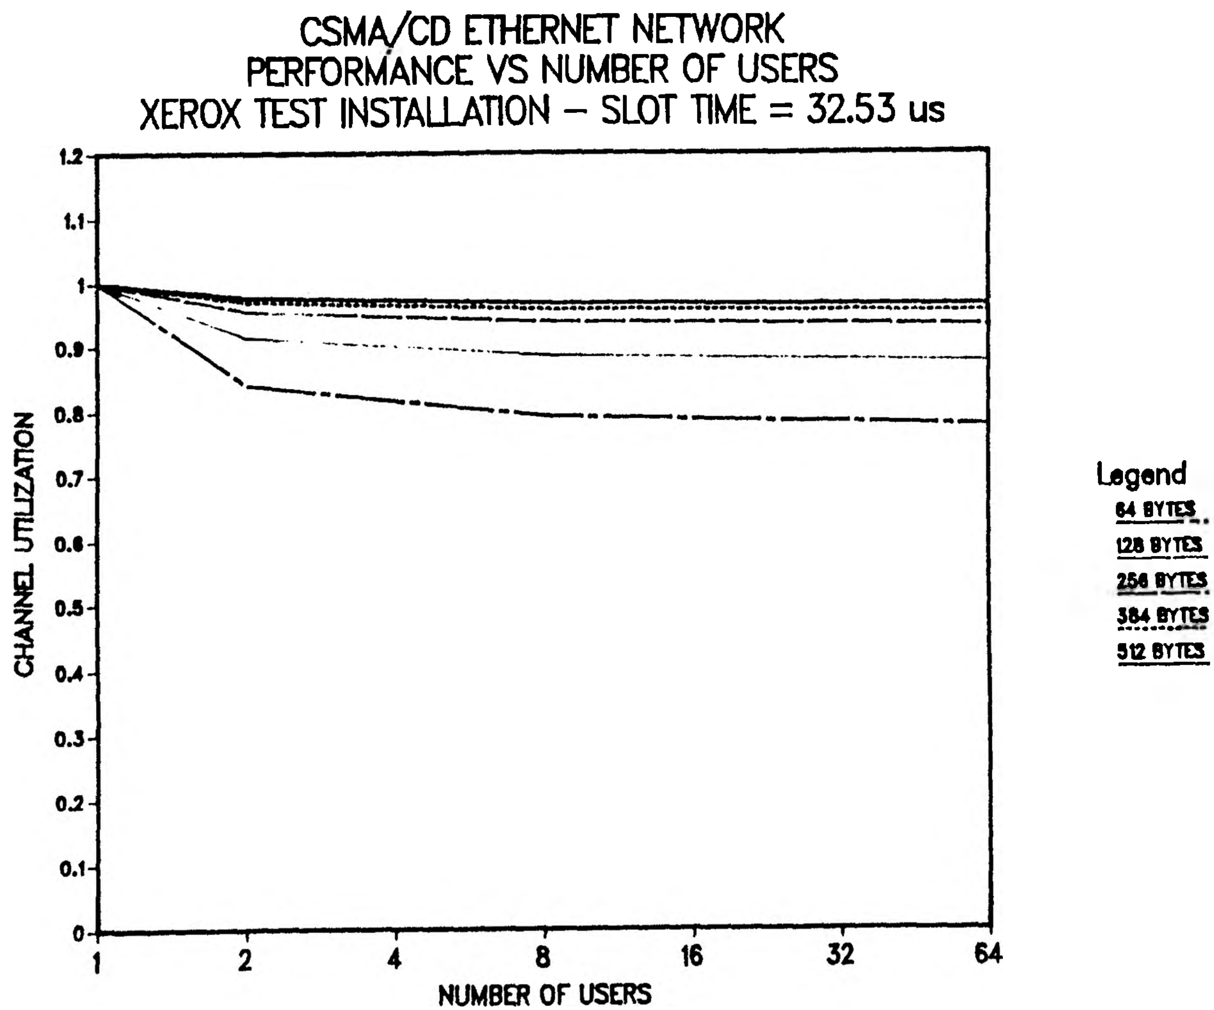

FIGURE 17 - XEROX TEST SITE EFFICIENCY CURVE

CSMA/CD ETHERNET NETWORK
PERFORMANCE VS NUMBER OF USERS
XEROX TEST INSTALLATION – SLOT TIME = 65.06 us

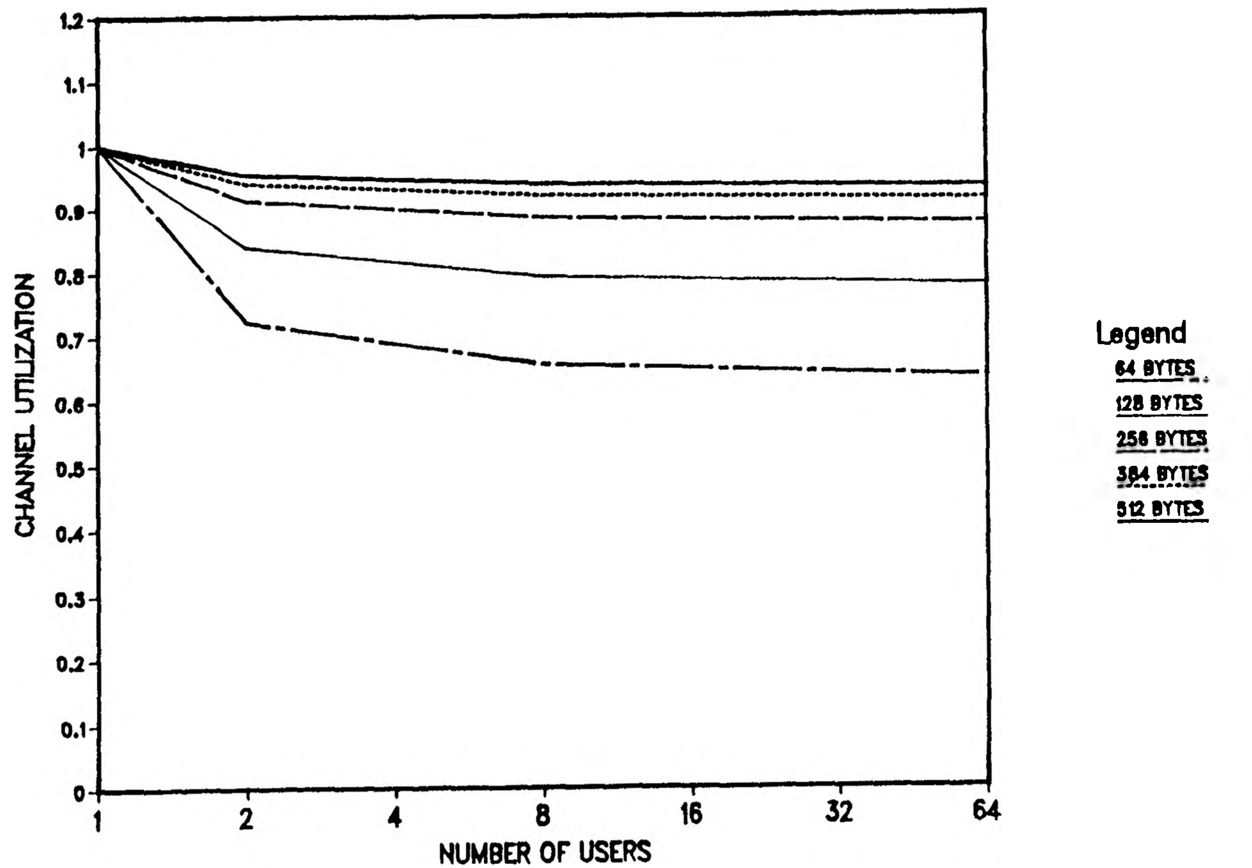

FIGURE 18 - UPDATED XEROX EFFICIENCY CURVE

possible delay changes due to the propagation delay. The predicted results can be seen in Table 1. It should be noted that the difference in D/T between the two propagation delays (1 and 1000 μs) is small. Even though one propagation value (1000 μs) is 1000 times greater than the other, the summed delay components (A_1,A_2 and A_3) for this value are only 1.0013316 times greater than for the other. When the components of D/T (A_1,A_2 and A_3) are taken apart, the greatest delay is generated by the A_3 (propagation and transmission delay) segment. Overall, A_3 comprises 97.27% of the delay, A_2 comprises 2.73% and A_1 represents only a small component of the total delay. The low value for the A_1 (normalized waste time due to collisions) segment, agrees with the previous data concerning collisions. Both studies, (Schoch, 1980) and (Digital, 1982), have show collisions to be less important than expected. The biggest components of delay at the cable level are caused by the propagation delays (A_3) and transmission time.

D. PERFORMANCE CONCLUSIONS

The Ethernet's performance has been tested and predicted under normal and heavy loads. The Ethernet channel's user capacity has also been examined along with average packet transmission delay. After viewing these results, what conclusions can be made about Ethernet performance and how can the performance for a designed or installed Ethernet be predicted?

Because of Ethernet's shared medium topology, the limiting factor

TABLE 1
LOAD VERSUS DELAY
FOR DIFFERENT PROPAGATION DELAYS

	1000 μs	1 μs
LOAD	D/T	D/T
.50	1.008510440853	1.007501009432
1.00	1.011025905494	1.010001023865
1.50	1.013546401463	1.012501043298
2.00	1.016071936313	1.015001067730
2.50	1.018602517609	1.017501097163
3.00	1.021138152922	1.020001131595
3.50	1.023678849836	1.022501171028
4.00	1.026224615943	1.025001215461
4.50	1.028775458847	1.027501264894
5.00	1.031331386162	1.030001319327
5.50	1.033892405511	1.032501378760
6.00	1.036458524526	1.035001443193
6.50	1.039029750854	1.037501512626
7.00	1.041606092148	1.040001587059
7.50	1.044187556071	1.042501666493
8.00	1.046774150298	1.045001750927
8.50	1.049365882516	1.047501840360
9.00	1.051962760418	1.050001934794

in performance would be collisions on the shared medium. If the number of collisions on the medium were high, there would be idle time on the channel and the bandwidth utilization would be lower than the offered load. Test data shows that collisions are not a serious concern until high loads exist on the network. In the Digital, et al. study (Digital, 1982), the mean number of packet transmission attempts remain close to 1 until over 1500 users shared the channel. Not many packets were being retransmitted due to collisions. In the Schoch and Hupp study (Schoch, 1980), it was determined that 98.18% of packets were sent with no latency, 0.79% were deferred and 0.03% were involved in collisions. The delay calculations showed the collision component (A_1) as the least important (less than 1% of total delay) in the delay total. With minimal amounts of collisions, as the offered load on the Ethernet channel increases, the bandwidth utilization should also increase at nearly the same rate. The Schoch and Hupp study showed a throughput utilization curve with the characteristics described above. The utilization grew with the offered load until offered loads were high. As the load crossed into a higher range (around 90%), the utilization curves flattened out. In general, under normal load conditions, the Ethernet's performance is predictable.

Under excessive loads and large numbers of users, the Ethernet performed well. As the Ethernet's load increased well over 100% capacity (stations attempting to generate total offered loads greater than 100% capacity), the bandwidth utilization remained stable (did not decrease). The efficiency curves showed that after the bandwidth reached its plateau, it stayed level for 5000 continuously queued stations. The Digital, et al. study also showed that for users number up to 2000; mean waiting time, percent idle time and mean number of attempts all remained reasonable. It should be noted that high load situations are usually only temporary. The predicting of Ethernet performance under high load situations is needed only to understand how Ethernet will perform under these temporary high load situations.

A methodology can be created to help predict the performance of a designed or existing Ethernet network. The first step is to analyze the network's traffic types and volumes. For example, if the network is purely terminal based, then the offered load and the utilization of the channel's bandwidth will both be low. The determination of expected offered load can be made using the predicted offered load graphs in this paper. After the offered load on the network has been determined, the utilization under normal loads can be approximated from the Schoch and Hupp utilization curve. If the offered load is under 80% of the channel's bandwidth, there should be close to a one-to-one correlation between offered load and utilization. For the case of using the Metcalfe and Boggs efficiency equation (Metcalfe, 1976) and the Ethernet site specific inputs, the bandwidth utilization ceiling can be predicted for the Ethernet configuration. This result should reveal the maximum bandwidth utilization that can be expected from the configuration.

Finally, it has been shown that by increasing a network's slot time, the efficient bandwidth utilization ceiling (seen in the efficiency calculation) can be decreased (and more importantly a shorter slot time will increase the efficiency). By decreasing a network's collision window during high load periods, the probability of a collision can be reduced and the maximum bandwidth utilization can be increased. This concept can be used to the advantage of a network designer. The next section shows how slot time can be used to optimize network design.

V. DESIGNING AN ETHERNET NETWORK

A. METHODOLOGY

Hancock (Hancock, 1987) lists 18 benefits a properly designed network can provide for a company (whichever network topology is being used):

1. Proper analysis of existing equipment for network installation,
2. List of requirements for network installation,
3. Proper configuration of network components for optimum cost savings,
4. A network topology that is flexible and adaptable,
5. Correct selection of network hardware and software for the network function,
6. Documentation of the network for future enhancements and modifications,
7. Migration path into future network technologies without redesign,
8. A long network life-cycle (reducing the costs of potential replacement),
9. Interconnect paths and methods for multiple network architectures,
10. User analysis and configuration of network resources for optimal use,
11. Network management plan and methodology to reduce downtime and allow for maximum use of available resources,
12. Expectations for performance, reliability and usability,

13. Optimal programming environment for network(ed) applications,
14. Training needs for programmers, users and network managers,
15. Recurring expense forecasting and budgeting methods,
16. Network support needs (programming, management, user support),
17. Use of mathematical modeling tools to help insure the success of the network design and topology, and
18. Optimal design to prevent network congestion, queueing delay, and proper placement of routing and management resources on the network.

These benefits can be turned into design goals. They will be used as steps that must be done during network design. If the decision has been made to network via Ethernet, than step 4 has already been taken care of.

The network designer can be in one of two positions. He could be designing for a brand new company that is creating its computer facility and network from scratch. All decisions about the network and computer placement are being made together. The other design option is to build the network around existing equipment and facility. Since, for most companies, the need for computers came before the local area networking push, the approach presented will assume the second case. In order to build an Ethernet design methodology, categorize the above network design list into four main stages of network design:

1) configuration worksheet,
2) initial design,
3) growth/upgrade considerations, and
4) optimize the network design.

The first design stage concerns the configuration worksheet. On this document, the current and near-term (within 6 months) equipment and facilities configurations and needs are laid out. Then, the network activities needed from the above documented hardware and facilities are presented.

In the initial design stage, the physical design for the Ethernet network is first created. This design must fall within the maximum physical configuration constraints seen in Digital, et al. specifications (Digital, 1980). Within this portion of the network design, the designer selects the hardware and software required to perform all the functionality for connecting the network components.

The third stage is the growth/upgrade considerations phase. Here, the long-term facilities and equipment anticipated changes are examined. Expected network technological advancements are also considered. Will the initial network design support these new changes or technologies? The cost issue is also very important. If the network is designed with no growth flexibility, it will be very expensive to expand the existing network to handle the new needs. Also, since LAN technology is growing at fast pace, the cost and ease of upgrading the Ethernet must be considered. Overall, forethought should be used when choosing software and especially hardware for the proposed network so that an installation is prepared for the future.

The final design stage is optimization. Methods for predicting Ethernet network performance at the cabling level have been discussed. Using the performance statistics calculated from the initial network design, the designer should analyze and redesign, if necessary any portions of the network that will yield a better cost/performance advantage at this low cabling level. The following design study will show how these four stages of network design can help to design an optimized Ethernet network.

B. DESIGN STUDY USING PERFORMANCE

A hypothetical company New Fix Incorporated will be used for the Ethernet design study. New Fix currently has two separate computing facilities (housed in one location), each oriented toward supporting a single project. The two computer groups perform limited communications between one another (manual magtape data transfers). New Fix Incorporated is currently under a major reorganization. They have purchased another company and are moving its people and projects in-house. This newly purchased company has a functioning computer facility. The new work being brought into New Fix will highly depend on data coming from the two existing projects. A corporate decision has been made to centralize the three current computing facilities into one facility. They have also decided to network the entire company using Ethernet. The design methodology discussed in the previous section will be used to design the New Fix Ethernet network. Even though high level aspects of Ethernet Networking (such as the protocols used) play an important roll in its design, only the cabling design and optimization issues will be addressed.

The first step in the Ethernet network design is to create the configuration worksheet. The company's three currently existing computer facilities and the building itself have the following configurations.

Current configuration (Group 1):

1) 1 Vax 8600

2) 1 HP plotter

3) 1 LP26 line printer

4) 40 directly connected terminal users

Current configuration (Group 2):

1) 1 Dec 2060

2) 1 LP26 line printer

3) 1 HP plotter

4) 60 directly connected terminal lines

Current configuration (New group):

1) 1 Vax 8600

2) 1 Laser printer

3) 1 LP26 line printer

4) 40 directly connected terminal lines

Facilities configuration:

Current;

1) New Fix Inc. is housed in a two story building, 150,000 square feet per floor (each floor 500 by 300 feet).

Future;

2) For environmental reasons, the new computer facility will be built in the northeast corner of the top floor. All three computers, four printers and two plotters will be housed in this area.

3) There will be two terminal rooms, one on each floor, housed at the south end of the building.

The company sees the following as its current networking goals.

Networking configuration goals:

1) All three computers will have file transfer capabilities between them.
2) The Laser printer will be shared by all three computers.
3) The two plotters will be shared by all three computers.
4) Another 100 terminals will be added to the existing 100 (240 total). Every terminal shall have the ability to get to any machine. Terminals will placed in the terminal rooms (20 per room) and on people's desks throughout the facility.

Now that the configuration and the needed network functionality data has been gathered, the next step is to create the initial design. This design configuration must fall within the maximum configurations as specified by Digital , et al. (Digital, 1980). Since the maximum length of any transceiver cable is 50 meters, the network must be designed so that the trunk line is within 50 meters of any possible terminal location. Also, the total design configuration must not exceed the 500 meter limit for any one section of trunk line nor exceed the total distance of 1500 meters of trunk.

The following design (Figure 19) shows a trunk line of approximately 1000 meters. It winds around the New Fix facility as shown. It is designed using three major trunk segments showing similar configurations on both floors. This design places all possible terminal locations within a 50 meter (usually less) range of the trunk. The three trunk sections are of the following lengths; top floor (353 meters), first floor (353 meters) and connecting the middle of both floors (276 meters). The three sections are connected via two repeaters

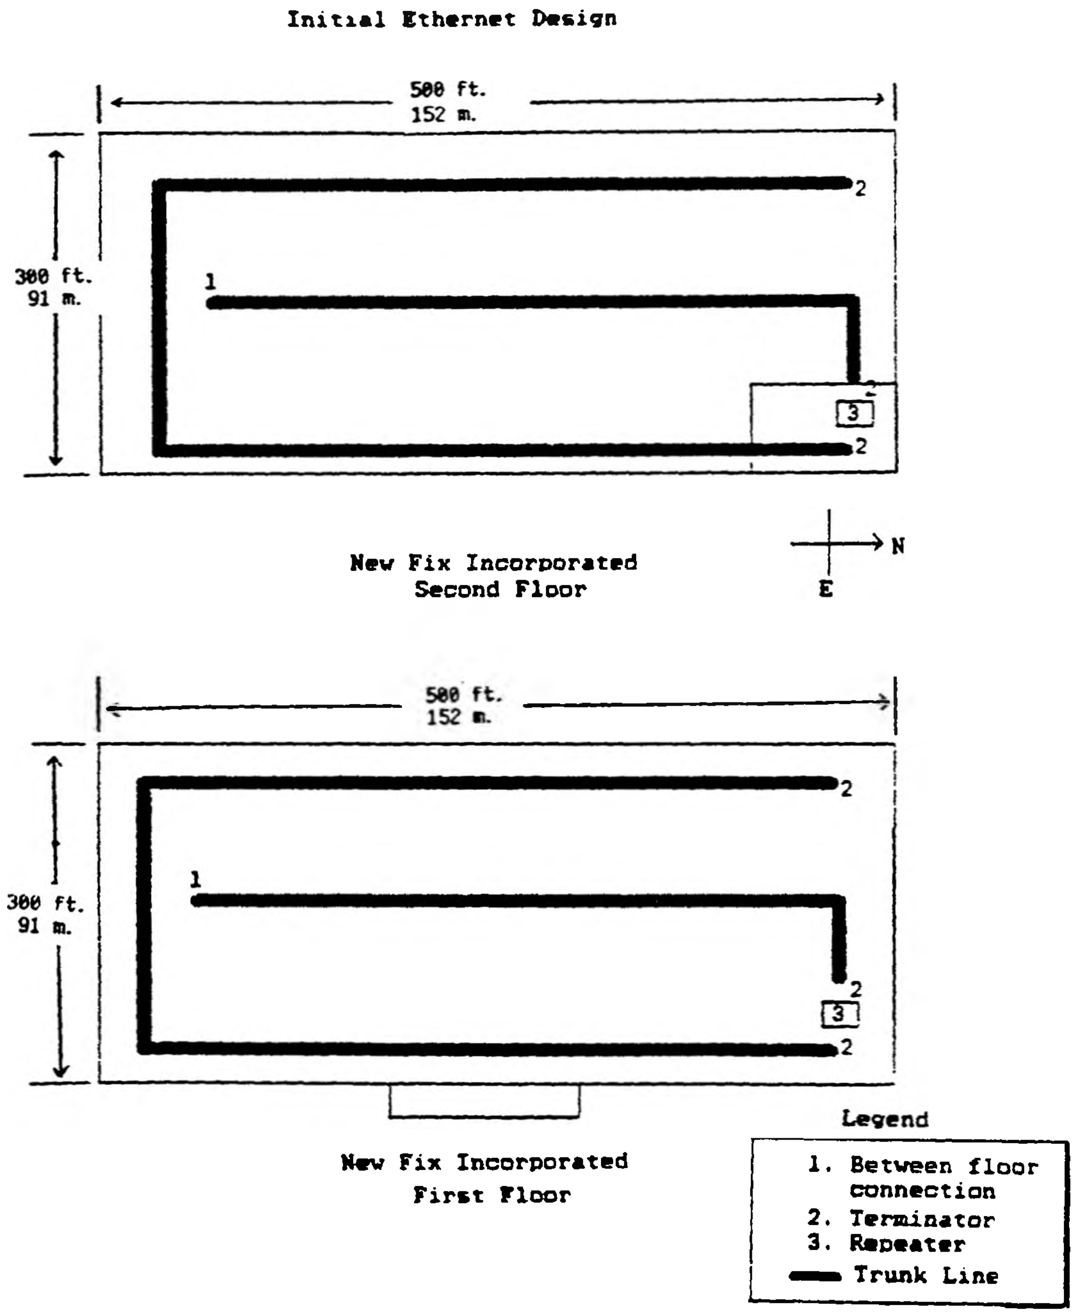

FIGURE 19 - DESIGN STUDY (NEW FIX INCORPORATED) FLOOR PLAN

placed at the north end of the two floors. The bandwidth of the trunk line will be 10 mbps and maximum slot time (982 meters of trunk line at 8.6 μs and six 50 meter transceiver cables at 3.08 μs) will be 11.73 μs.

The third stage of network design is make sure the network design will support future growth and upgrade. Only hardware growth will be discussed since hardware specifics are not part of this case study. The cable design was created to handle future expansion. This design should also allow for computer center expansion and possible relocation. It appears that this design should meet the company's in-house growth needs.

The Metcalfe and Boggs (Metcalfe, 1976) efficiency calculation is used to determine the utilization that can be expected from the initial network design. The calculations for the network with 10 mbps channel capacity and slot time of 11.73 μs are represented in Table 6 and Figure 20. It was previously concluded that Metcalfe and Boggs efficiency calculation yielded accurate results. The current design then yields good results. The final step in the design methodology is to optimize the initial design (if possible). The next step is to examine the possibility of optimizing the network.

To optimize the design, the expected traffic on the network must be analyzed. The expected traffic on the network will come from three sources; terminal traffic (bursty with low average packet size), shared device (Laser printer and HP plotters) traffic and file transfer (disk access) traffic. As seen before, the highest offered load can be expected to come from file transfer traffic. This file transfer traffic should also have the highest average byte size per packet. In the

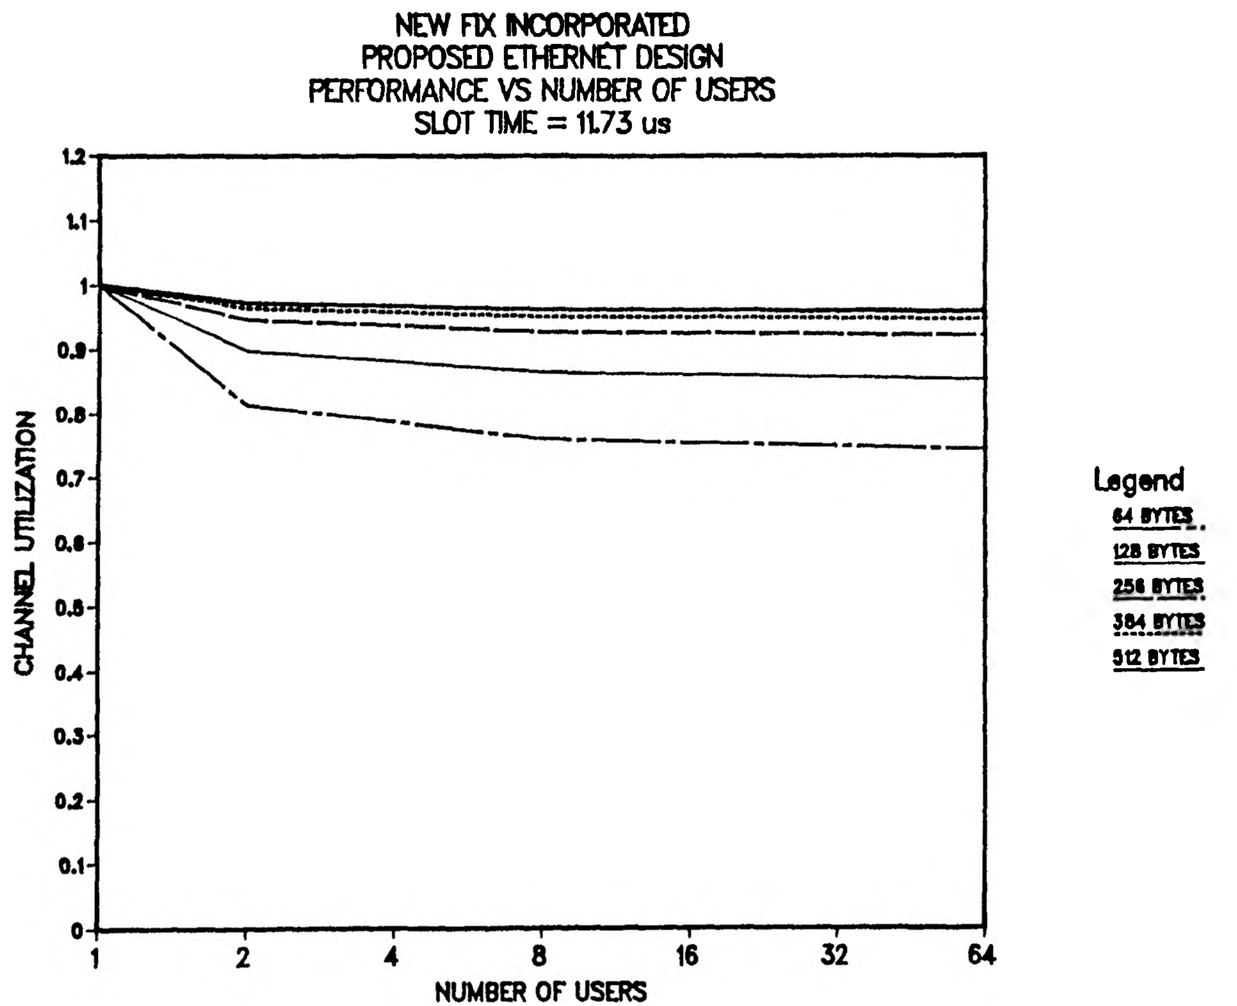

FIGURE 20 - DESIGN STUDY UTILIZATION CURVES

offered load calculations before, a distinct difference in the loads offered by terminals and printers and those offered by disk (file transfer) traffic was observed. It was also noted that on a heavily load network, as the slot time decreased, the expected line efficiency increased.

The traffic volume segregation between terminal/printer and file transfer (disk) traffic fits into another pattern. All the higher volume traffic is generated in one centralized location, the computer facility. If another smaller Ethernet were to be localized in the computer facility, the slot time on this computer room network would be decreased. Would this new design increase any performance on the network?

Table 7 and Figure 21 show the efficiency calculations for a second Ethernet network (length 70 meters) dedicated to file transfers and shared device traffic within the computer facility. These predicted values represent an increase in efficient channel utilization over the original design. For an average packet size of 64 bytes, there is a 37.5% increase. At the high end (512 byte traffic), there is a 5.3 % increase in utilization. With this new design, the performance potential for a possibly high volume area of the network has been increased. The design of the facility's main Ethernet can also be optimized.

If the middle trunk segment used to connect both floors were flipped around, the repeaters would then be connected at the building's southern ends. The connection would be made directly into the middle of both the top and bottom floor trunk segments. In turn, this would

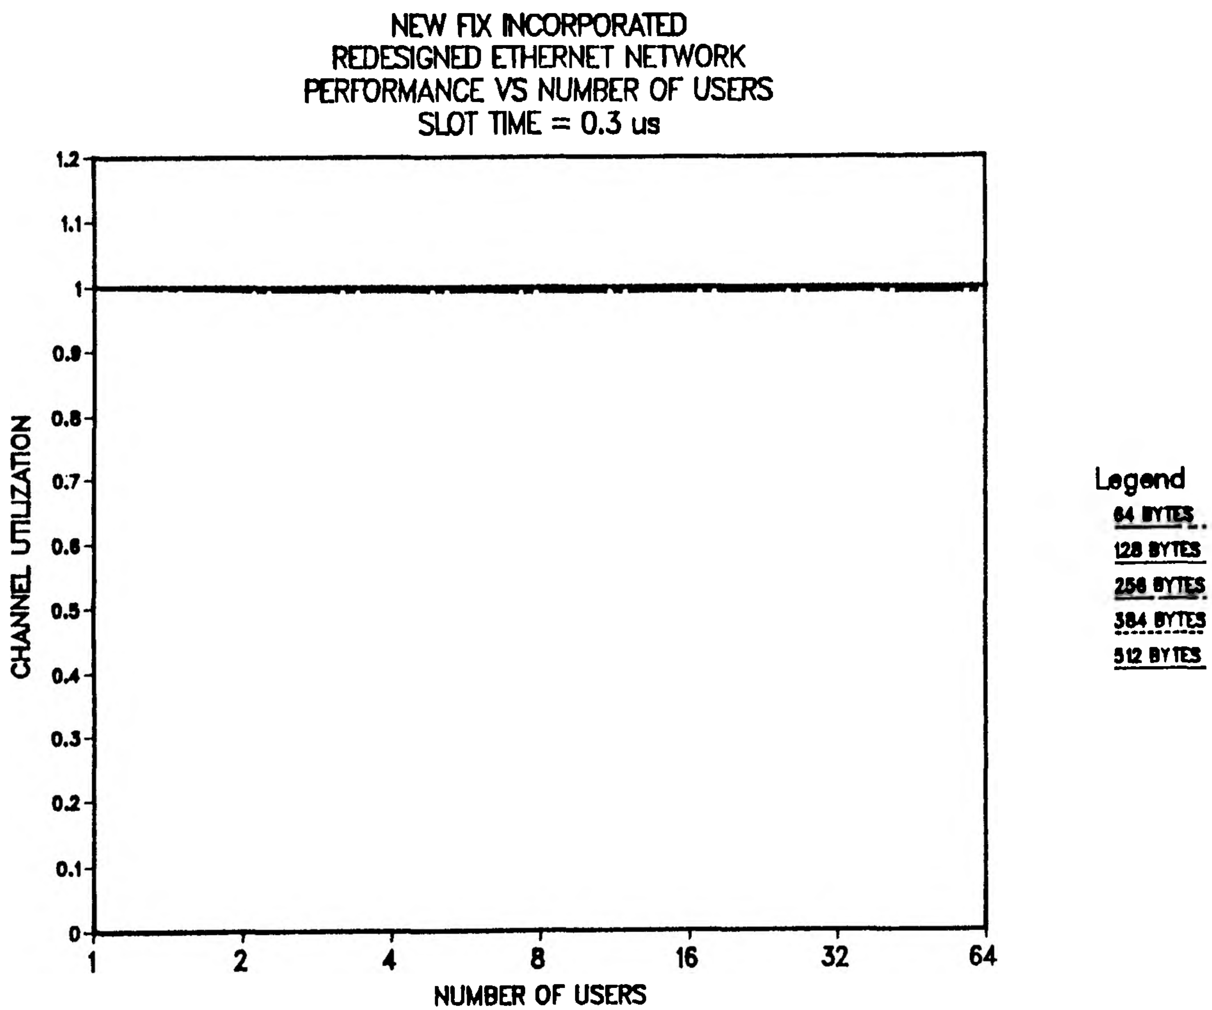

FIGURE 21 - UPDATED DESIGN STUDY UTILIZATION CURVES

decrease the maximum distance between any node from the current 1000 meters to around 800. This would decrease the slot time to 10 μs. It has been shown that by decreasing the slot time, the network's utilization ceiling will be increased. As a general rule of thumb, the trunk segments of a network should be connected to each other at a point which minimizes the network's maximum slot time.

Two ways to possibly optimize the cabling design of an Ethernet network have been discussed. Decreasing the networks slot time is a desirable goal for the second change above. But, it is right to optimize the network by creating a separate computer facility Ethernet? To make this decision, many factors have to be considered. Most importantly, the "cost/performance" tradeoff needs to be considered. The cost of adding another trunk line to the design must be considered because Ethernet cable is not cheap. The cost of the cable must be weighed against the added utilization this new design gives the high traffic section of the network. These things must be discussed at final design time. When the total picture (hardware costs and performance needs) is visible, a decision needs to be made. The ability to optimize Ethernet design is available though if you choose to use it. More insights into Ethernet design and implementation can be seen in Appendix A.

VI. CONCLUSIONS

This work has covered the Ethernet local area network from an informational, performance and design standpoint. The Ethernet LAN topology was first presented from an informational point of view to give the reader the background needed to understand the performance and design discussions and conclusions. Methodologies were presented which should help the user predict performance and/or aid in the design of a new Ethernet network. For a novice or expert, the insights shown may help in the understanding, design and management of an Ethernet network.

In general, Ethernet is an easy to use and understand LAN alternative. All aspects of the network; performance, operation and design, are of a simple nature and remain easy to understand. The simplicity of the Ethernet concept characterizes the Ethernet creators' initial intentions for reliability through simplicity. Being of a simple nature, the ethernet network is easier to manage, operate and understand. Overall, the choice of using an Ethernet network to locally network a facility is a good one.

BIBLIOGRAPHY

Brindley, Bill. "LAN Architectures: Making Connections," Digital Review, (November 10, 1986), 86.

Brindley, Bill. "How Ethernet LAN's keep data moving," Digital Review, (November 24, 1986), 69.

Digital, Intel, Xerox. The Ethernet - A Local Area Network. September, 1980.

Digital. Intro to Local Area Networks, May 1982.

Hancock, Bill. "Decnet and the OSI - What should we expect?," The DEC Professional, 4, 12 (December 1985), 108-118.

Hancock, Bill. "Pay Now or Pay Later," The DEC Professional, 6, 2 (February 87), 122-127.

Ilyas, M. and H.T. Mouftah. "Performance Evaluation of Computer Communications Networks," IEEE Communications, 23, 4 (April 1985), 18-29.

Lo, T.L. "Local Area Networks: A tutorial." 1984 Fall St. Louis CMG meeting, November 13, 1984.

Metcalfe, R. and D. Boggs. "Ethernet: Distributed Packet Switching for Local computer Networks," Communications of the ACM, 19, 7 (July 1976), 395-404.

Micom/Interlan, Slide show at Vendor Demonstration, February 3, 1987.

Mitchell, Lionel C. "A Methodology for Predicting End-to-end Responsiveness in a Local Area Network," Proceedings of the IEEE, (August 1981), 83-91.

Reiser, Martin. "Performance Evaluation of Data Communications Systems," Proceedings of the IEEE, 70, 2 (February 1982).

Schoch, J. and J. Hupp. "Measured Performance of an Ethernet Local Area Network," Communications of the ACM, 23, 12 (December 1980), 711-721.

Stuck, B.W. "An Introduction to Traffic Handling Characteristics of Bus Local Area Network Distributed Access Methods," IEEE Communications Magazine, 22, 8 (August 1984), 46-56.

VITA

Michael Martin Chaney was born in Atlanta, Georgia. He received his primary and secondary education in Kirkwood, Missouri. He attended the University of Missouri-Columbia in Columbia, Missouri and received a Bachelor of Science in Computer Science in May 1984.

He has been employed by the Engineering and Space division of the Emerson Electric Corporation since June 1984. In January 1985, he enrolled in the Graduate School of the University of Missouri-Rolla, Rolla, Missouri.

APPENDICES

A. CASE STUDY: EMERSON ELECTRIC

I. INITIAL NETWORKING MOTIVATION

II. CHOOSING NETWORK HARDWARE

III. EMERSON CONFIGURATION AND DESIGN

IV. IMPLEMENTATION DETAILS

V. EMERSON NETWORK TODAY

VI. FUTURE GROWTH

VII. GENERAL OVERVIEW

Emerson Electric Engineering and Space Division is a Department of Defense contractor housed in St. Louis, Missouri. It is a division of Emerson Electric Company which is headquartered in St. Louis. The Engineering and Space (E&S) Division is located on three main campuses; two in St. Louis and one in Orlando, Florida. The computing resources of the division are separated into two functional groups; one supporting business and financial applications (Management Information Systems - MIS) and the other supporting engineering (Scientific Computing Center - SCC). The MIS department is an IBM shop and is housed on the Corporation's main campus. The SCC is a Digital shop and is located at the Evans Lane engineering facility. The communication between the two groups is performed via manual magtape transfers. Under the direction of the SCC, an Ethernet Local Area Network was installed in the Evan's Lane facility in the summer of 1984.

I. INITIAL NETWORKING MOTIVATION.

In 1983, the computing resources for the Engineering side of the E&S division were small but growing. The SCC had one Digital 2060 mini computer and one Vax 11/780 (recently installed). The SCC also maintained Computer Vision CAD/CAM machines. Dispersed throughout the rest of the two building Evans Lane campus were two 11/780 minicomputers and two 11/750 minicomputers. These machines were project owned and maintained. There was also of network of Mentor-Apollo workstations which was used for Electrical Design. With the arrival of the SCC's second main host, a Vax 11/780, the need to eliminate the cabling problem became evident. The constant moves of personnel caused continuous cabling configuration problems. The troubles of constant cable configuration changes were the biggest motivation for local networking.

Two other networking motivations lay on the horizon. Another DEC 2060 was being acquired and a large staffing increase was expected in the coming year. The current total number of terminals owned by the SCC was around 200. They were concentrated in two terminal rooms. The others were dispersed on desktops and in project labs. For the facility to be able to handle more users connecting to three hosts, the local network need became obvious. At this initial point in time, the local network drive was motivated by terminal connectivity only. Minimal shared device traffic was also planned (printers/plotters), but no file transfer traffic was foreseen for the near future.

II. CHOOSING NETWORK HARDWARE

From the beginning, it was determined that the network hardware choosen was going to dictate the LAN design used. At the time of the networking decision, 1983, the LAN market was still in its infant stage. For the network, terminal connectivity to two Digital machines (one DEC 2060 and one VAX 11/780) was needed . The most used host was the Dec20, so any network hardware purchased had to address the Dec20 terminal traffic. The search was directed to vendors supplying connectivity to DEC equipment. The leaders in the search were DEC, HYPERbus/HYPERchannel and Interlan. The network topologies compared were Star (HYPERbus) and Ethernet. The final decision was made considering both cost and flexibility concerns. Interlan won out over DEC because DEC had no Ethernet terminal server for the Dec20 in 1983. DEC's Ethernet hardware mainly addressed file transfer and some VAX connectivity. The HYPERbus alternative lost out because its cost per connection was higher than Interlan's. Also, the star network topology was not as well liked by the people making the decisions as the Ethernet. All hardware decisions were made at this time knowing that within two years the network hardware would probably be obsolete. The Interlan alternative offered us the best solution in the quickest timeframe.

III. EMERSON CONFIGURATION AND DESIGN

The Evans Lane facility is located in an old department store. The facility is made up of three floors in two separate buildings. The SCC computing facility is located in the middle of the lower floor (at the foot of the escalators). The project labs are located at the north end of both floors. Most of the Engineers are housed in the giant room on the top floor, but there are a few on the bottom floor and in the attached building. There is also a third building, currently housing a grocery store, that will hopefully hold more engineers in the near future. The initial Ethernet design needed to allow for terminals on any potential engineer's desk within the facility. The initial design included 4 trunk lines connected via 3 repeaters. Each of three floors has a trunk and there was one trunk connecting the three. An additional length of trunk line has been designed into the system to network the grocery store building when it becomes available. The network floor plan can be seen in Figure 22. The design falls within the maximum allowable Ethernet configurations. Even though there are three repeaters connecting the network, there are only two between any two nodes. By placing the repeaters in their current locations, network slot time is minimized. The total slot time for the Emerson Ethernet is 10 us. The network was designed to allow transceiver connectivity to any part of all floors.

Initially, each trunk line was intended to be one continuous piece. It was discovered that the weight of 500 meters of Ethernet trunk line was too heavy and inflexible to weave around the ceiling. The actual cabling used trunk lengths which allow the design to bend and flow with the length and width of the buildings (70.2 and 230.3 meter segments were used). The separate sections for each major trunk line are connected via barrel connectors. Finally, it was determined the only feasible place to run this cable was in the ceiling. The terminal servers were to be placed in the ceiling (or under the raised floors in the computer center). By placing the terminal servers in the ceiling, it kept them out of the reach of users.

IV. IMPLEMENTATION

After the initial trunk lines, repeaters, and transceivers were installed, there were some difficulties with the network. The problem segments of the network could be pinpointed by selectively shutting off different repeaters. A Digital Field Service representative was brought in to troubleshoot the network. A Time Domain Reflectometer (TDR) was

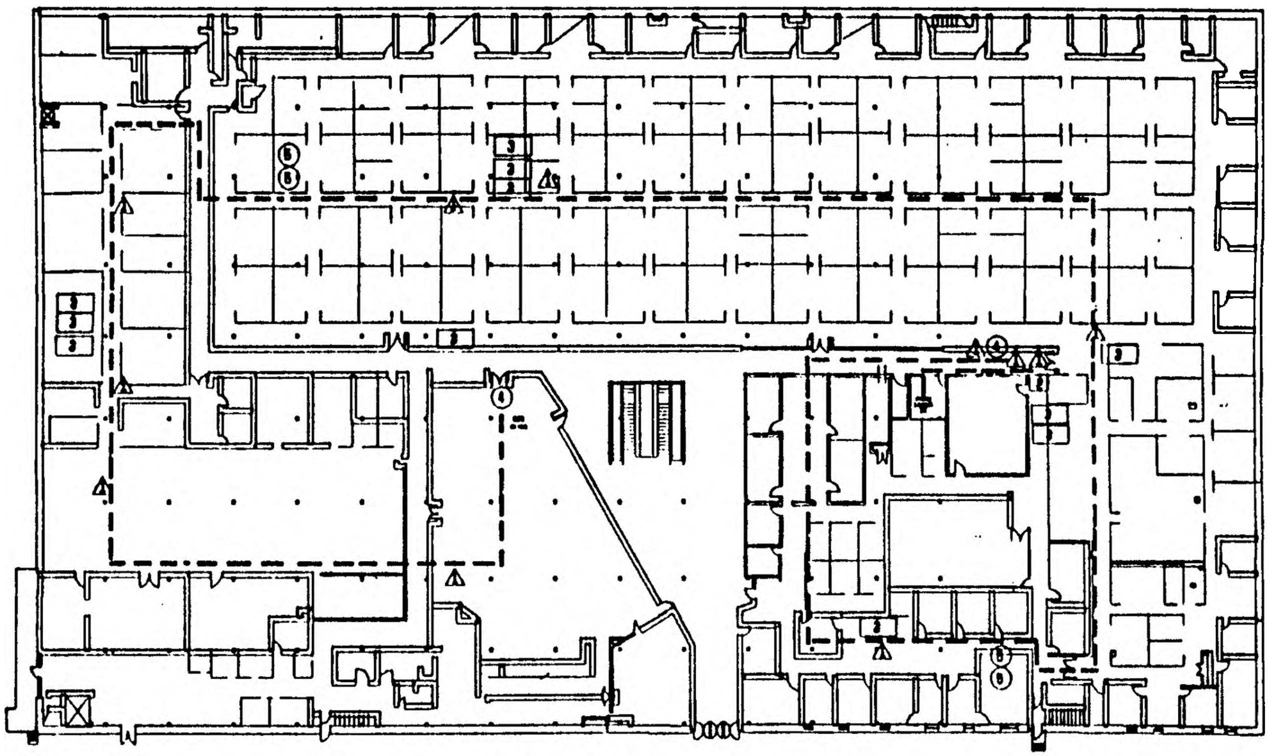

UPPER LEVEL BLDG HH

MENU

- TAP
- REPEATER
- TERMINAL SERVER
- TERMINATOR
- BARREL CONNECTOR

UPPER HH TRUNK:
1300 feet
400 meters

FIGURE 22 - EMERSON ELECTRIC NETWORK FLOOR PLAN

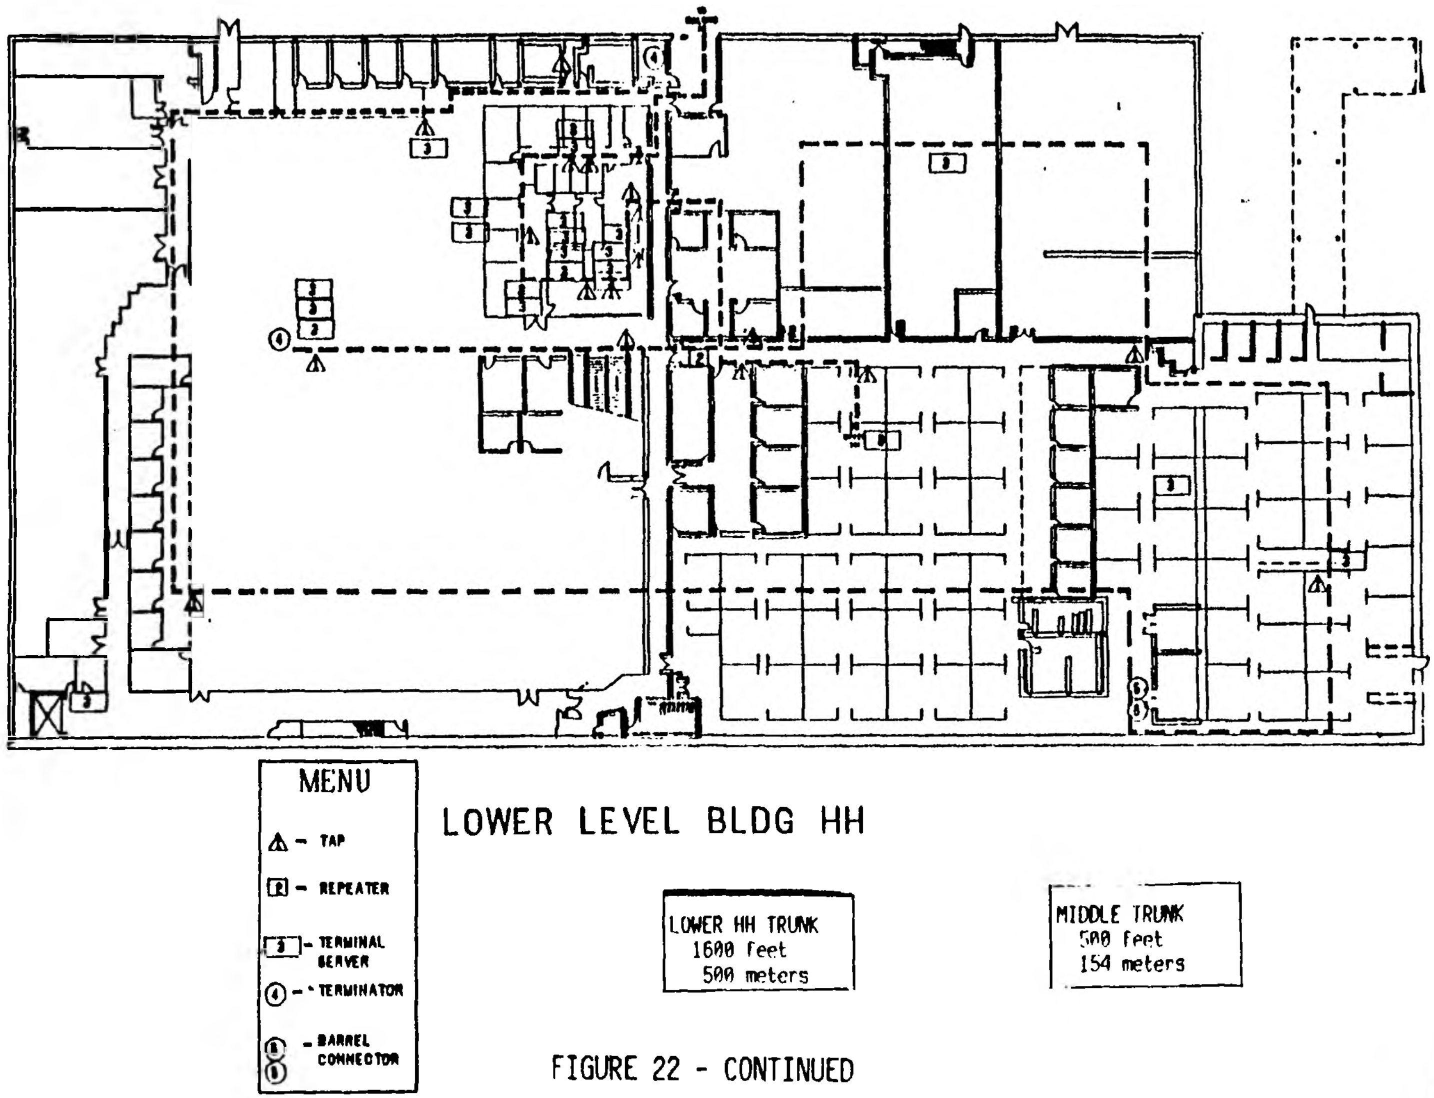

FIGURE 22 - CONTINUED

BLDG GG

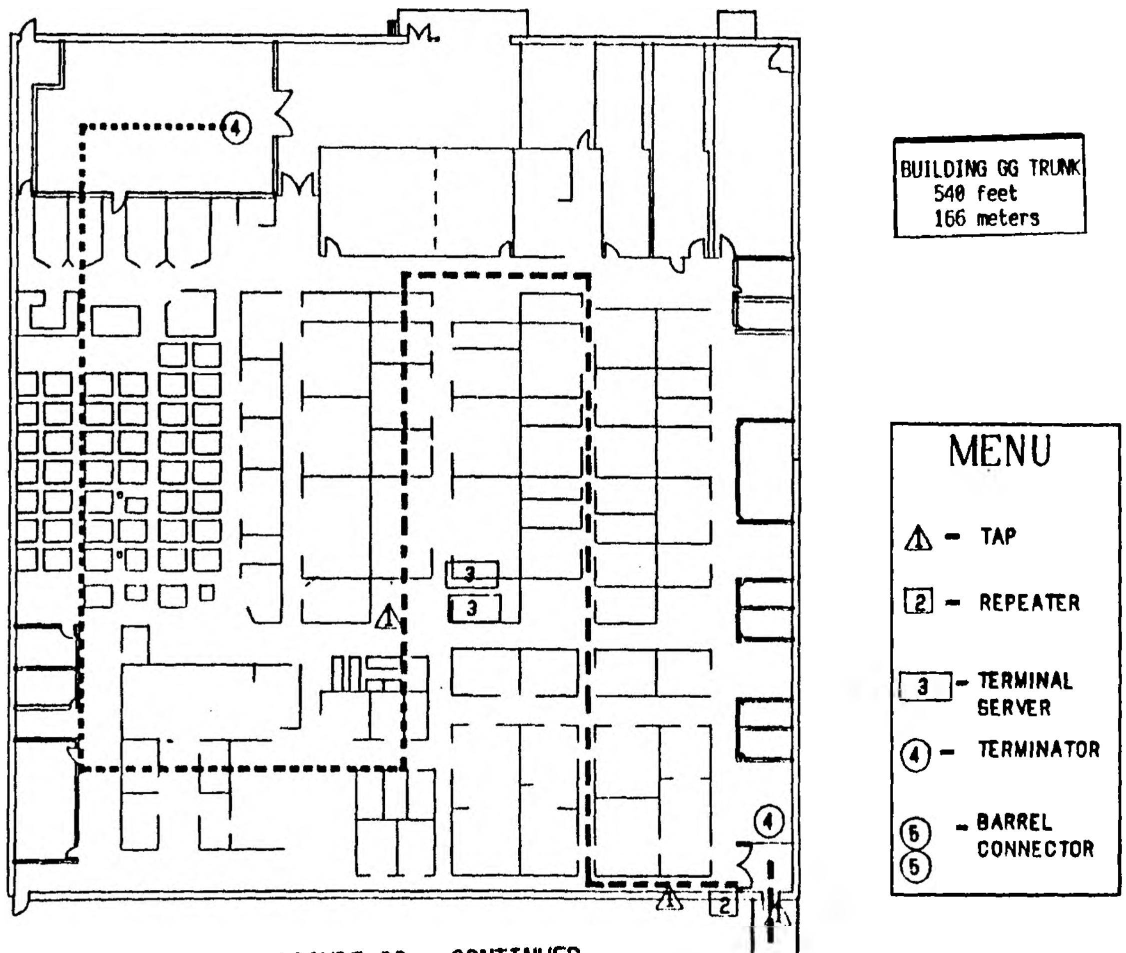

FIGURE 22 - CONTINUED

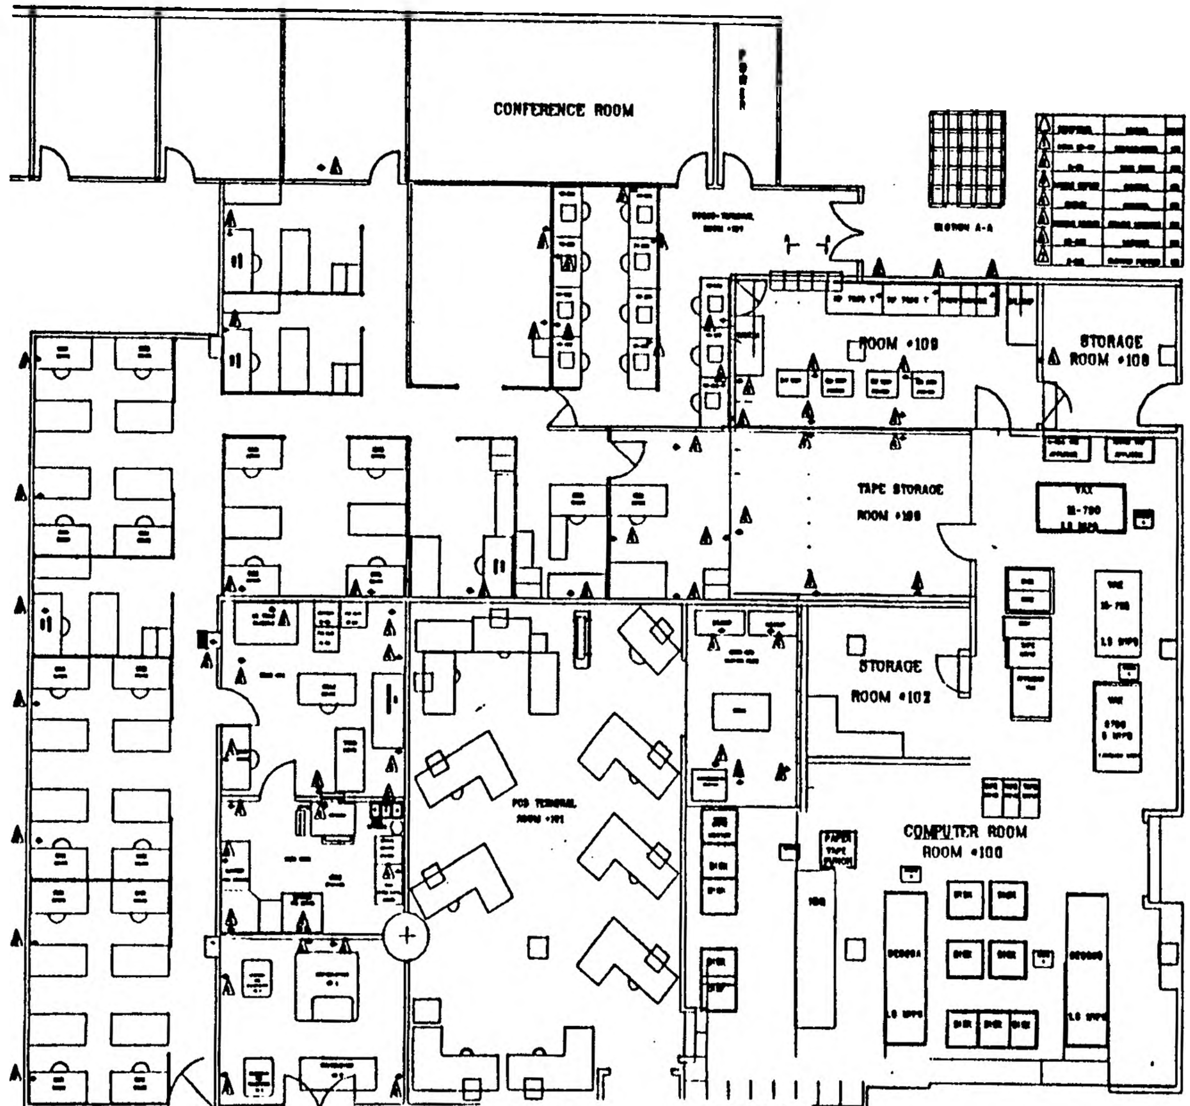

FIGURE 22 - CONTINUED

used to check the problems. It was determined that when individual transceiver taps had been drilled, some metal shavings had not been cleaned out of the holes. These were causing shorts in some cable segments. To fix the problem, all transceivers and equipment were removed from the trunk lines. Then the electrician ran 120 volts over the trunk lines. This process should have disintegrated the metal chips that were causing the problems. The problem disappeared for awhile. When it returned, the TDR procedure was attempted again. Another power surge was applied to fix the problem. The trunk line problems have not been a problem since. After the initial problems, the network transceivers were changed from DEC H4000 transceivers to the Interlan NT100 model. The H4000 transceiver required drilling for its installation while the NT100 used a pierce installation technique. This transceiver change may have cleared up the trunk line problems. Since the initial problems, the Ethernet has been very reliable. There have been some minor problems. During some power failures, repeaters have crashed. Until the repeater is rebooted, the network segments on either side cannot communicate. Some intermittent terminal server problems have been experienced. By changing the tap and transceivers, this has usually cleared up the problem. For the most part, the Ethernet Network has been trouble free.

IV. EMERSON NETWORK TODAY

The configuration worksheet for the Emerson electric E&S division Evans Lane facility can be seen in Table 2. The Ethernet network has

TABLE 2 - EMERSON ELECTRIC CONFIGURATION WORKSHEET

	#	NODE NAME	NO. OF STATIONS	HOST ?
	1	APPLICON (APPLE 1)	1	Y
	2	APPLICON (APPLE 2)	1	Y
	3	ASG15 AREA 1	8	
	4	ASG15 AREA 2	8	
	5	ASG15 LAB	8	
	6	CT2 CABINET	8	
	7	CT2 CEILING	24	
	8	CT2 LAB	8	
	9	DATSA AREA	24	
	10	DEC 2060 A	32	Y
	11	DEC 2060 B	24	Y
	12	ELEC. DESIGN LAB	16	
	13	IRAD USERS	8	
	14	IRAD VAX	16	Y
	15	LOWER HH ANNEX	8	
	16	LOWER HH PRE-ANNEX	8	
	17	LOWER HH TERM. RM.	16	
	18	MODEM CABINET	16	
	19	RELIABILITY	8	
	20	SCC STAFF	8	
	21	TPSWB	8	
	22	TRAINING ROOM	8	
	23	UPPER HH ADMIN.	8	
	24	UPPER HH TERM. RM.	16	
	25	VAX 11780	32	Y
	26	VAX 11785	33	Y
	27	VAX DMF32 PORTS	8	Y
TOTALS	27		363	

grown since its beginning. The SCC computer facility now supports seven hosts; two Dec 2060 computers, one Vax 11/785, one Vax 11/780, one Vax 8700 and two Applicon Micorvax IIs. The terminals have migrated from the terminal room to the desktop. The numbers of terminals owned by the SCC have also grown to over 500 and the number is steadily increasing. Finally, the network traffic has changed from a strictly terminal traffic based network to a terminal traffic and file transfer network. Five of the hosts are connected via DECNET. Predicted performance information for the Emerson Ethernet can be seen in Figures 6 and 15.

VI. FUTURE GROWTH

The E&S Ethernet is heading in a file transfer direction. The file transfer capabilities (now and needed) can be seen in Figure 23. The need for the users to communicate with all of our various vendor machines is becoming more important. Because of this multi-vendor connectivity question, the future terminal servers will be of a new variety. The NTS100 is Micom-Interlan's new product. It gives the user both XNS and TCP/IP protocol functionality. To change from one protocal to the other, a prom pack switch is all that is needed. The existing network is running the XNS protocal which is the more efficient of the two. The Department of Defense has made a push in the networking area for standardization. TCP/IP has become their protocol of choice. Thus, the different vendor versions of TCP/IP are much more compatible than the different versions of XNS. Though this issue is not as important to us now. As the need for mutli-vendor connectivity grows, this protocol connectivity flexibility may be needed.

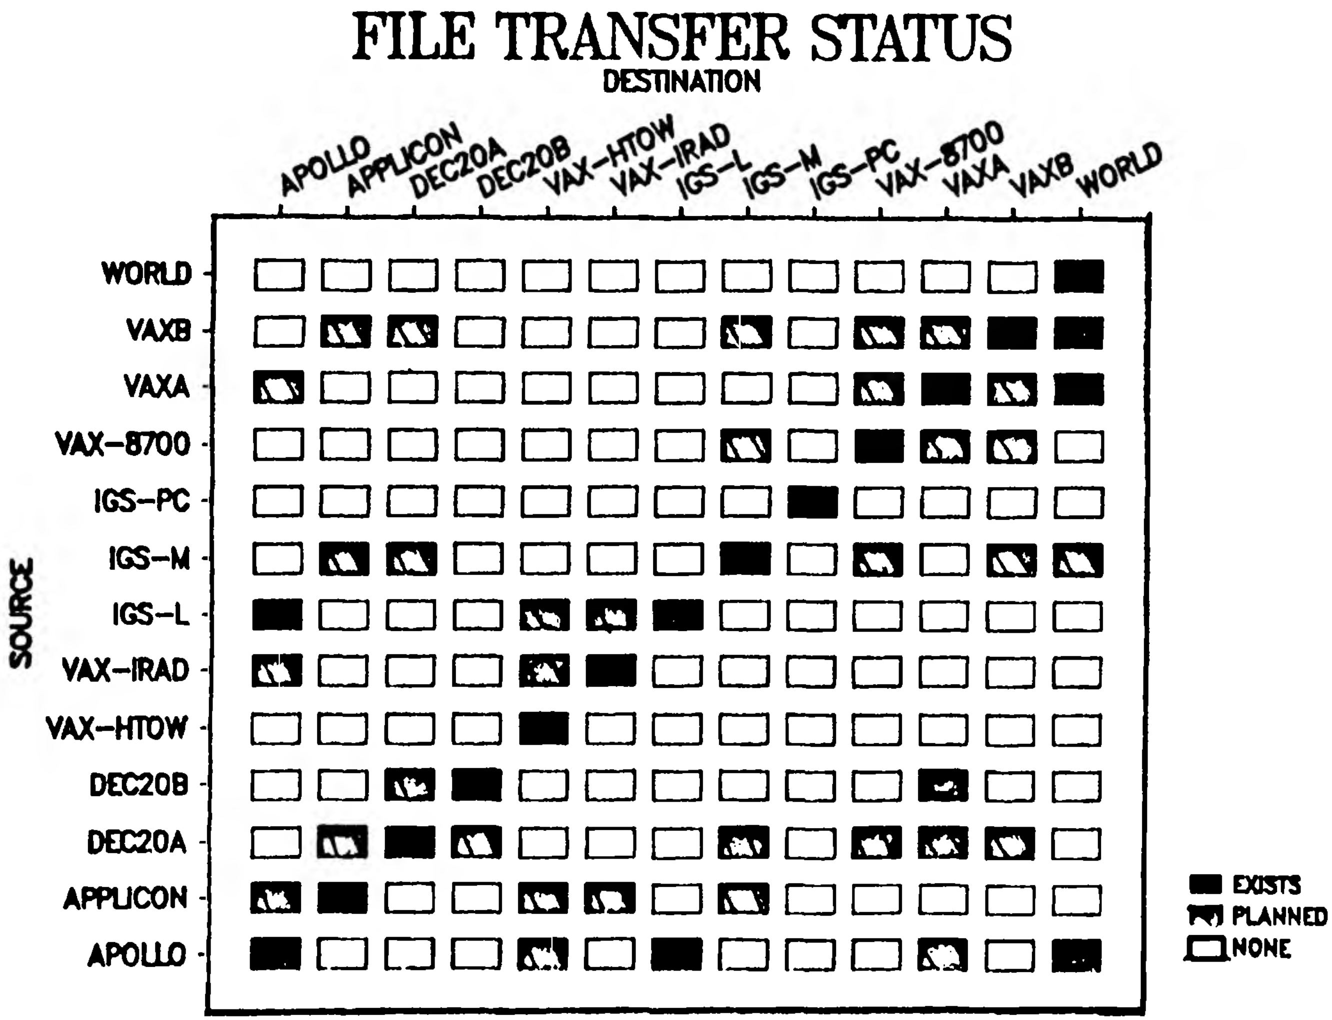

FIGURE 23 - EMERSON ELECTRIC FILE TRANSFER CAPABILITIES

Like many other companies, Emerson is experiencing an internal personal computer push. With this new PC rush, the need for network connectivity for these machines is also a must. The ideas of direct Ethernet connectivity and separate PC networks are both being examined at this time. A previous transceiver problem has just been solved. With the high concentration of transceivers needed in a computer room, the tap distance limitation becomes a problem. The DEC Delni (and similar vendor products) allows up to eight transceiver connections from one Ethernet tap. It acts like a transceiver multiplexor. This eliminates the problem. Finally, it has been determined that the ceiling is not the best place to house terminal servers. When they crash, it is difficult to climb in the ceiling and reboot them. Also,it is hard to change a broken server from the top of a ladder. The new push is to move the servers to communications closets. The cable movement problem that caused the initial networking push is again an issue. A product from Cable Management Systems will give our user's terminal connectivity from any support column in the facility. 25 twisted-pair RJ71 cables are being dropped down the support columns to connect terminal's. The terminal servers are housed in a locked closet out of the users reach, but within easy reach of the network manager.

VII. GENERAL OVERVIEW

For Emerson Electric, the Ethernet Network has been a very important tool. It has eliminated a lot of time and cost in the cabling arena. With DECNET, the computing facility has become more efficient. The manual file transfer function is now automated. It gives users the ability to connect to different hosts using a user friendly mechanism accessed from their desktop. From a managerial standpoint, the network has been easy to maintain. Aside from the initial installation problems, there have not been many problems. Because the Micom-Interlan Ethernet hardware and software are relatively easy to install and operate, it makes it easier for a network manager to learn how to repair any problems that may occur. Overall, it has been a good tool for giving our users reliable and easily attainable computing resources.

B. BANDWIDTH EFFICIENCY TABLES

TABLE 3 - EMERSON EFFICIENCY DATA TABLE

ETHERNET PERCENT EFFICIENT UTILIZATION
EMERSON ELECTRIC ETHERNET NETWORK
SLOT TIME = 10 µs
C = 10 mbps

	PACKET SIZE IN BYTES				
USERS	64	128	256	384	512
1	100	100	100	100	100
26	75	86	92	95	96
51	75	86	92	95	96
76	75	86	92	95	96
101	75	86	92	95	96
126	75	86	92	95	96
151	75	86	92	95	96
176	75	86	92	95	96
201	75	86	92	95	96
226	75	86	92	95	96
251	75	86	92	95	96
276	75	86	92	95	96
301	75	86	92	95	96
326	75	86	92	95	96
351	75	86	92	95	96
376	75	86	92	95	96
401	75	86	92	95	96
426	75	86	92	95	96
451	75	86	92	95	96
476	75	86	92	95	96
501	75	86	92	95	96
1001	75	86	92	95	96
1501	75	86	92	95	96
2001	75	86	92	95	96
2501	75	86	92	95	96
3001	75	86	92	95	96
3501	75	86	92	95	96
4001	75	86	92	95	96
4501	75	86	92	95	96
5001	75	86	92	95	96

TABLE 4 - XEROX EFFICIENCY DATA TABLE

ETHERNET PERCENT EFFICIENT UTILIZATION
XEROX PALO ALTO EXPERIMENTAL ETHERNET NETWORK
SLOT TIME = 32.53 μs
C = 3 mbps

	PACKET SIZE IN BYTES				
USERS	64	128	256	384	512
1	100	100	100	100	100
26	76	86	93	95	96
51	76	86	93	95	96
76	76	86	93	95	96
101	75	86	92	95	96
126	75	86	92	95	96
151	75	86	92	95	96
176	75	86	92	95	96
201	75	86	92	95	96
226	75	86	92	95	96
251	75	86	92	95	96
276	75	86	92	95	96
301	75	86	92	95	96
326	75	86	92	95	96
351	75	86	92	95	96
376	75	86	92	95	96
401	75	86	92	95	96
426	75	86	92	95	96
451	75	86	92	95	96
476	75	86	92	95	96
501	75	86	92	95	96
1001	75	86	92	95	96
1501	75	86	92	95	96
2001	75	86	92	95	96
2501	75	86	92	95	96
3001	75	86	92	95	96
3501	75	86	92	95	96
4001	75	86	92	95	96
4501	75	86	92	95	96
5001	75	86	92	95	96

TABLE 5 - UPDATED XEROX EFFICIENCY TABLE

ETHERNET PERCENT EFFICIENT UTILIZATION
XEROX PALO ALTO EXPERIMENTAL ETHERNET NETWORK
SLOT TIME = 65.06 μs
C = 3 mbps

	PACKET SIZE IN BYTES				
USERS	64	128	256	384	512
1	100	100	100	100	100
26	61	76	86	90	93
51	61	76	86	90	93
76	61	76	86	90	93
101	61	75	86	90	92
126	61	75	86	90	92
151	61	75	86	90	92
176	61	75	86	90	92
201	61	75	86	90	92
226	61	75	86	90	92
251	61	75	86	90	92
276	60	75	86	90	92
301	60	75	86	90	92
326	60	75	86	90	92
351	60	75	86	90	92
376	60	75	86	90	92
401	60	75	86	90	92
426	60	75	86	90	92
451	60	75	86	90	92
476	60	75	86	90	92
501	60	75	86	90	92
1001	60	75	86	90	92
1501	60	75	86	90	92
2001	60	75	86	90	92
2501	60	75	86	90	92
3001	60	75	86	90	92
3501	60	75	86	90	92
4001	60	75	86	90	92
4501	60	75	86	90	92
5001	60	75	86	90	92

TABLE 6 - DESIGN STUDY UTILIZATION DATA TABLE

ETHERNET PERCENT EFFICIENT UTILIZATION
NEW FIX INCORPORATED
PROPOSED ETHERNET NETWORK
SLOT TIME = 11.73 µs
C = 10 mbps

	PACKET SIZE IN BYTES				
USERS	64	128	256	384	512
1	100	100	100	100	100
26	90	95	97	98	99
51	90	95	97	98	99
76	90	94	97	98	99
101	90	94	97	98	99
126	90	94	97	98	99
151	89	94	97	98	99
176	89	94	97	98	99
201	89	94	97	98	99
226	89	94	97	98	99
251	89	94	97	98	99
276	89	94	97	98	99
301	89	94	97	98	99
326	89	94	97	98	99
351	89	94	97	98	99
376	89	94	97	98	99
401	89	94	97	98	99
426	89	94	97	98	99
451	89	94	97	98	99
476	89	94	97	98	99
501	89	94	97	98	99
1001	89	94	97	98	99
1501	89	94	97	98	99
2001	89	94	97	98	99
2501	89	94	97	98	99
3001	89	94	97	98	99
3501	89	94	97	98	99
4001	89	94	97	98	99
4501	89	94	97	98	99
5001	89	94	97	98	99

TABLE 7 - UPDATED DESIGN STUDY UTILIZATION DATA TABLE

ETHERNET PERCENT EFFICIENT UTILIZATION
NEW FIX INCORPORATED
REDESIGNED ETHERNET NETWORK
SLOT TIME = .3 μs
C = 10 mbps

	PACKET SIZE IN BYTES				
USERS	64	128	256	384	512
1	100	100	100	100	100
26	99	100	100	100	100
51	99	100	100	100	100
76	99	100	100	100	100
101	99	100	100	100	100
126	99	100	100	100	100
151	99	100	100	100	100
176	99	100	100	100	100
201	99	100	100	100	100
226	99	100	100	100	100
251	99	100	100	100	100
276	99	100	100	100	100
301	99	100	100	100	100
326	99	100	100	100	100
351	99	100	100	100	100
376	99	100	100	100	100
401	99	100	100	100	100
426	99	100	100	100	100
451	99	100	100	100	100
476	99	100	100	100	100
501	99	100	100	100	100
1001	99	99	100	100	100
1501	99	99	100	100	100
2001	99	99	100	100	100
2501	99	99	100	100	100
3001	99	99	100	100	100
3501	99	99	100	100	100
4001	99	99	100	100	100
4501	99	99	100	100	100
5001	99	99	100	100	100