

To the Graduate Council:

I am submitting herewith a thesis written by David Thomas Bell entitled "Strategy for the Implementation of Data Link Bridges." I have examined the final copy of this thesis for form and content and recommend that it be accepted in partial fulfillment of the requirements for the degree of Master of Science, with a major in Electrical Engineering.

Robert W. Rochelle

Robert W. Rochelle, Major Professor

We have read this thesis
and recommend its acceptance:

Samuel B. Koch

William J. McClure

Accepted for the Council:

Leuminkel

Vice Provost
and Dean of The Graduate School

STRATEGY FOR THE IMPLEMENTATION
OF DATA LINK BRIDGES

A Thesis
Presented for the
Master of Science
Degree
The University of Tennessee, Knoxville

David Thomas Bell
December 1989

DEDICATION

To my mother and father who taught me the value of a good education and to my wife, Mary Ellen, whose support and patience were invaluable in the pursuit of this degree.

ACKNOWLEDGEMENTS

A portion of the work presented in this thesis was sponsored by Martin Marietta Energy Systems, Inc., under contract number DE-AC05-84OR21400 with the U.S. Department of Energy.

The direction of my major professor, Dr. R. W. Rochelle, and the very timely review of the thesis by Dr. W. J. McClain and Dr. D. B. Koch are gratefully acknowledged. I am very appreciative of their patience and willingness to serve on my committee.

I wish to express my sincere thanks to Mr. J. E. Cleaves for his many hours of consultation and encouragement in the completion of this task, and Dr. W. J. McClain in helping me assemble this thesis into an acceptable publication.

Finally, I wish to thank my wife Mary Ellen, whose "tolerance" of this effort has been appreciated more than she could ever know.

ABSTRACT

One operation requiring the implementation of a local area network is the ability to transfer large files of data between mainframe computers attached to the network. Many companies utilizes IEEE 802 networks and the Digital Network Architecture DECnet protocol suite as the means by which large files of data are transported between processors involved in their manufacturing process. In order to support file transport over the network, a means of utilizing diverse network architectures and topologies is required. The challenge is to devise a strategy for meshing the physical cable plant, communications protocols, and DECnet software to form the communications link between processors sharing data. Solutions to these problems require a thorough understanding of the different methods used to interconnect networks, performance characteristics of the solutions, and trade-offs. One solution to this required interconnection involves the deployment of data link layer bridges to extend the IEEE 802 networks. Development of technical specifications required for data-link bridges for such an application are examined in detail. This development process takes into account rapidly evolving network standards, expanding local area network technology, and the network implementation planning that may be in

progress in a manufacturing company. The approach is (1) to begin with a functional set of requirements, (2) to examine currently available product offerings, and (3) to compare these products against the implementation plans. The result is a technical specification utilized to obtain a system of data link bridges and a management system which supports the manufacturing environment and conforms to the manufacturing company's networking strategy. This paper represents the culmination of one and one-half years of research by the author; this effort was initiated due to the lack of previously available equipment which met the functional requirements of networking in a manufacturing environment.

TABLE OF CONTENTS

SECTION		PAGE
I.	INTRODUCTION	1
	A. Networking in a Manufacturing Environment	1
	B. Problem Description	6
	C. Literature Survey	8
	D. Solution Alternatives and Selection	10
II.	OVERVIEW AND INTERNETWORKING	16
	A. Physical Layer Transmission Media	16
	B. Data Link Layer	20
	1. Network Access	20
	2. Logical Link Control	25
	3. Network Performance	27
	C. Connectivity Device	30
	D. Network Management	35
	1. Management Functions	35
	2. Capacity Planning	37
	3. Information Management	38
	4. Accounting	39
	5. Security	40
	6. Other Requirements	41
	E. Description of Digital Network	
	Architecture and DECnet	42
	1. Digital Network Architecture	42
	2. Protocol Structures of DNA	44
	A. User Layer	48
	B. Network Management Layer	49
	C. Network Application Layer	50
	D. Session Control Layer	51
	E. End Communication Layer	52
	F. Routing Layer	53
	G. Data Link Layer	56
	H. Physical Link Layer	59
III.	PREPARATION OF A TECHNICAL SPECIFICATION	61
	A. General Discussion	61
	B. Requirements Definition	62
	C. Bridge Functional Description	65
	1. Data Link Bridge Description	65
	2. Function Required of the Data	
	Link Bridge	66
	D. Network System Functional Description	68
	E. Bridge Physical Requirements	70
	F. Bridge Performance Characteristics	72
	G. Bridge Management System	82
	H. Reliability, Availability, and	
	Maintainability (RAM) Requirements	88

I.	Quality Assurance	92
1.	General Testing Requirements	92
2.	Local Bridge Testing	92
3.	Remote Bridge Testing	93
4.	Packet Filtering Testing	95
5.	Miscellaneous Testing	97
6.	RF Testing	98
7.	Application Layer Testing	98
8.	System RAM Testing	100
J.	Evaluation Criteria	103
IV.	SUMMARY AND CONCLUSIONS	106
A.	Evaluation and Product Selection	106
B.	Long Range Target Architecture	107
	BIBLIOGRAPHY	112
	APPENDIX	115
	VITA	130

LIST OF FIGURES

FIGURE	PAGE
1. Point-to-Point Versus LAN Connectivity	3
2. Layers of the OSI Reference Model	5
3. Typical Broadband CATV Network	21
4. IEEE 802.3 Frame Format	23
5. Throughput as a Function of Network Traffic	28
6. Network Delay as a Function of Network Traffic	29
7. DECnet Architecture	45
8. Bridge and Router Operation Models	47
9. Data Link Bridge Configurations	64
10. System Description	67
11. Network System Configuration	69
12. Bridge Management System Architecture	83
13. Bridge Packet Filtering Test Configuration	96
14. Application Layer Test Configuration	99
15. Bridge RAM Test Configuration	101

I. INTRODUCTION

A. Networking in a Manufacturing Environment


The concept of the Computer Integrated Enterprise can best be described by the integration of corporate-wide information systems using computer and related technologies. Within an integrated manufacturing environment, networking provides flexibility for the coordinated control of the manufacturing process from product definition to final shipment. Integration of various design and production-control strategies requires a number of computing machines which range in size from mainframes to programmable logic controllers (PLCs) and numerical control machines on the factory floor with large minicomputers, microcomputer workstations, and terminals included in this wide spectrum. The mainframe computers are usually responsible for large central data processing functions such as inventory control, production scheduling, data collection and analysis, tool management, and design functions (Computer Aided Design).

A common network topology found in manufacturing is the interconnection of computers involved in the process using a point-to-point network configuration. This configuration provides information and data flow in a limited, non-robust fashion and also provides a central point of failure. On the other hand, a high degree of flexibility and robustness


can be achieved by the implementation of local area networks. Figure 1 illustrates the differences in these two configurations. The advantages of the networking concept relative to the conventional point-to-point configuration are:

1. Traffic flow is not restricted to a specific destination based on topology
2. Information exchange between two computers is not necessarily dependent on other possible intermediate computers and therefore does not influence the performance of a computer not directly involved in the exchange
3. Since mainframes are available (directly) to serve more than one other remote processor, the network provides improved flexibility, greater utilization of computing resources, and reduced network-induced delays resulting in increased data throughput
4. Simultaneous mixing of heterogeneous protocols between machines attached to the network can be achieved

Although many other rationale exist, those listed above are the most dominant in terms of justification for change. The major design challenges in the communications systems for



(a) Point-to-point connectivity



(b) Lan connectivity

NIU = Network Interface Unit

FIGURE 1

POINT-TO-POINT VERSUS LAN CONNECTIVITY

the manufacturing environment are connectivity, flexibility, interoperability, performance, and management. The challenge in connectivity is that of establishing effective communications among all systems in the manufacturing enterprise, whether such systems are from a single vendor or from multiple vendors. The problems here are exacerbated by the need to interconnect and distribute processing among computers which are not even from the same generation. Functional flexibility is required to deal with the wide range of needs in the factory. It is difficult to design a single network topology that can satisfy the needs of all applications. Interoperability is driven by the recognition of the need for industry-wide standards in communications. Standards based on the Open Systems Interconnection (OSI) Reference Model, portrayed in Figure 2, provide a common ground from which vendors can create integrated systems and be assured that their products will be compatible with what other vendors are designing. This issue is also driven by user's desires for internetworking and coordination among different architectural approaches. The network management infrastructure manages the physical routing and troubleshooting. It also manages logical resources such as extended local area network (LAN) services for bridge traffic. Standards also play a key role in the direction that the management issue is taking.

Layer	Name	Function
7	Application	Selection of appropriate service for the application
6	Presentation	Code conversion and data formatting
5	Session	Coordination of interaction between end-application processes
4	Transport	End-to-end data integrity and quality of service
3	Network	Switching and routing of information
2	Data Link	Transfer of units of information to the other end of the physical link
1	Physical	Transmission of the bit stream to the transmission medium

FIGURE 2
LAYERS OF THE OSI REFERENCE MODEL

B. Problem Description

Typically, the user environment plays the determining role in the network implementation. The three possible environments are office, campus, and factory. Each one imposes a specific and different set of requirements and a characteristic set of problems. In assessing this environment, the factory environment is actually a combination of all the above. This paper assumes that the manufacturing environment employs point-to-point communications links as the means by which information is transferred between machines. Performance characteristics, connectivity, management, and security requirements are sources of reason for the need for a different method of transferring data. These point-to-point links utilize DECnet¹ Phase III and Phase IV, in the transfer of data from shop floor machines to the central facility hosts. These links are deployed via a dedicated twisted pair local area network, the state-of-the-art in connectivity only a few years ago.

Many manufacturing companies currently utilize broadband coaxial cable communications networks for providing terminal-to-host communications needs. This system employs

¹ DECnet is a registered trademark of Digital Equipment Corporation

Community Antenna Television (CATV) technology and transmits data in an analog format using frequency modulation (FM) methods.² Geographically, this cable plant runs adjacent to the areas requiring high speed communications (the same areas currently serviced by the dedicated twisted pair network). In addition to high speed data transfer capabilities, the network is also able to provide a more secure, robust, and tightly-managed environment. High-capacity communications is provided via multiple channels. The network also has the large area of coverage needed to support frequent moves and additions of individual manufacturing components.

As system needs have grown, point-to-point capabilities have been unable to support the manufacturing applications. Bandwidth-usage increases have been greater than can be supported on point-to-point channels. Diversity of protocol suites have also not been supported on these point-to-point channels. The broadband medium can provide a high-speed, common communications bus by which all hosts and subnets can intercommunicate regardless of their geographical location. The network can serve the immediate needs of interconnecting hosts by supporting proprietary and non-standard protocols. At the same time, it can help meet the growing requirements

² Broadband Network Technology. Sytek, Inc., 2nd ed. New Jersey: Prentice-Hall, c 1986.

for interoperability by supporting a migration to OSI protocols. From an economic standpoint, the network provides reasonable installation and operating costs. Therefore, this networking base affords an opportunity to use international standards in the initial deployment of data networks, and to design the interconnection of the networks using the current standards, products, techniques, and plans for the future. Ethernet³ on broadband provides a transport for high speed host-to-host services, but the actual services provided are a function of the higher layers of the network protocol suites (i.e. DECnet, etc.). The network must be reliable and have a high degree of availability. The network must not be a "bottleneck" to the manufacturing process as the current twisted-pair networks have proven to be.

C. Literature Survey

A literature search was performed to determine if a similar situation had been approached before. The search consisted of reviewing periodicals, books, magazines, vendor literature, and conference proceedings, as well as participation in verbal discussions with other professionals in the communications and networking field. The results of

³ Ethernet is a registered trademark of the Xerox Corporation

this search resulted in virtually no documented indication of work in this area. Conversations with other networking professionals indicated that this particular area needed this type of attention, since the prevalent attitude among current users is that they will simply use whatever products the various networking vendors will provide, instead of specifying what devices are required to meet functional requirements.

The search focused on the libraries at Martin Marietta Energy Systems, Inc. and the University of Tennessee and utilized the keyword and subject area search techniques using the computerized library card systems. Specific searches were made on the following keywords and subject areas: "Local Area Network", "Internetworking", "Network", "Bridge", "Gateway", "Ethernet", "Broadband", "IEEE", "Communications", "Data Communications", and "Data Processing". As a result, a great many references were generated for the general categories. A review of this material generated some material presented here. However, nothing specific to the area of data link bridge functional specification preparation was uncovered. Much of the material found is recent in nature, since this subject area is relatively new technology.

Review of magazines and periodicals provided the greatest volume of information. While the depth of technical discussion and detail is not great, these sources provided the most current work in this area. This area of technology is moving so rapidly that many advances and new product offerings are made from "month to month". Yet, with all the material available on related subjects, this work should provide future reference for additional efforts in this area.

D. Solution Alternatives and Selection

It is assumed that the major trends in the manufacturing plant's host computer communications environment (which involves the network and communications protocol) will impact the selection of a communications architecture to support user's needs. These trends will affect network design factors such as network interface speeds, physical connections, transmission media, connectivity requirements, interoperability requirements, end system and intermediate system protocols, and backbone bandwidth.⁴ The two major trends that are impacting this concept of networking are (1) distributed processing and applications and (2) external host-to-host communications. The trend towards distributed

⁴ Daigle, J. N. "Communications for Manufacturing: An Overview". IEEE Network, vol. 2, issue 3, p. 6-13.

processing and applications will be driven by the availability of local processing and the need to collect information from multiple sources. The external host-to-host communications reflects the trend toward electronic exchange of documents. The concept of centralized processing and hierarchical system control are giving way to distributed processing.⁵ An example of this is the use of database management systems (DBMS). While databases were originally designed for mainframe control and access, the latest trend is the movement toward distributed databases whereby portions of the database are relocated to remote locations for local processing. DBMS and other distributed applications will contribute to the need for peer-to-peer communications which often require greater bandwidth. Therefore, higher-speed interfaces will be needed to handle the increased bandwidth requirements.

The second major development impacting host computers is the trend towards ubiquitous access from across the organization. An example of this is an application such as Electronic Data Interchange (EDI). This is the standard that defines the exchange of documents between computers of different companies by defining their format and information content. Since hosts are becoming more powerful, but also

⁵ Interconnectivity: Open Enterprise Networking.
Ungermann-Bass, Inc., c. 1989.

smaller and more distributed, their applications will rely on peer-to-peer network access to distribute data files and processing among multiple, remote hosts.

The transition from centralized to distributed processing has been accelerating rapidly in recent years. Mainframe-oriented applications accessed by remote terminals and microcomputer workstations are being replaced with distributed applications. After years of logging onto a single host to access a single application, users are gaining the ability to access multiple hosts, to downline-load portions of a database for local processing, and to log onto multiple applications simultaneously. Host computers (or workstations in many cases) connected to host computers on a peer-to-peer basis create different connectivity requirements than the past hierarchical structure of terminal-to-host communications. These connections have typically higher speed network requirements (greater than 1 million bits per second) than those of terminals (less than 19,200 bits per second).

The need to interconnect hosts at high speeds will increase as processing is distributed and data is shared among local and remote hosts and peripherals. In the near future,

industry protocols (e.g. Network File System (NFS)⁶, Distributed Queueing Services (DQS)⁷, etc.) will allow hosts to share resources. In the long term, most of these host-to-host protocols will be standardized by the International Standards Organization (ISO) at the international level. At present, most host-to-host protocols are vendor proprietary. DECnet is an example of the concept of a proprietary architecture. The possible physical layer implementations of DECnet are:

1. Leased T1 lines and existing fiber optic links in point-to-point configurations
2. Ethernet over existing broadband cable plant using data link bridges
3. Ethernet over new fiber optic network
4. Point-to-point virtual links utilizing radio frequency (RF) modems on the broadband
5. Ethernet over existing broadband cable plant using repeaters
6. Ethernet over existing broadband cable plant using routers

⁶ Network File System (NFS) is a registered trademark of SUN Microsystems, Inc.

⁷ Distributed Queueing Services (DQS) is a registered trademark of Digital Equipment Corporation

These alternatives can be consolidated to a series of implementations of all the above solutions. The solution which utilizes data link bridges providing Ethernet over the broadband cable plant provides the best solution in terms of support for multiple protocols, geography, and cost benefits.⁸ In working toward this solution, a specification for data link bridges to provide DECnet host-to-host service was prepared. This specification provides Ethernet bridges and management systems for use on broadband, fiber optic, baseband coaxial cable, and T1 links. Further examination of possible network solutions for the manufacturing environment based on standards development, the networking industry's technological direction, and possible existing media installations results in the following logical groupings in order of implementation preference:

1. Existing broadband network
2. Existing baseband networks
3. Available T1 network
4. Planned fiber optic network

This concept of a functional specification to obtain data link bridges can be expanded to consider the previously

⁸ Hunter, John. "Bridges and Gateways: Supplying the Missing Link". Network World, p. 33-38. July 1988.

mentioned constraints in order to provide the following configurations:

1. Local Ethernets bridged to the broadband
2. Local Ethernets bridged together via existing fiber optic network
3. Local Ethernets bridged to other Ethernets
4. Remote Ethernets bridged via T1 links
5. Local Ethernets bridged via a future fiber optic network

Bridges serving all of the listed configurations and functions, yet managed by a single management console can be obtained via competitive procurement.

II. OVERVIEW OF INTERNETWORKING

A. Physical Layer Transmission Media

Implementation of the solution utilizing data link bridges requires various media for networking. Most of the networks mentioned are based on the Ethernet Carrier Sense Multiple Access with Collision Detection (CSMA/CD) specification designed by Xerox, Intel, and Digital Equipment Corporation and revised by the Institute of Electrical and Electronics Engineer's (IEEE) 802.3 standards committee.⁹ Until recently, these networks were implemented on baseband coaxial cables providing a single data channel. While these networks provide high speed performance, they can lead to expensive recabling when non-data network applications become necessary. However, Draft IEEE Standard 802.1 and recently-developed products allow Ethernet networks to utilize a broadband cable. The broadband cable can support additional LANs, video, and other services on different frequency bands (channels). The focus of this standard is a device that lets an Ethernet controller (typically a board in an attached computer) access a broadband CATV medium. The IEEE 802.3 specification describes a radio frequency (RF) transceiver that modulates the Ethernet signals and

⁹ ANSI/IEEE Standard for Local Area Networks 802.3: CSMA/CD Access Method and Physical Layer Specifications. IEEE Computer Society, 4th ed. New York: IEEE Press, c 1986.

places them on one channel of the multichannel broadband network.¹⁰ In addition to the benefits of using a shared physical medium, broadband Ethernet networks can cover a wider distance and area than can the baseband version, due to the fundamental difference in the way the two networks are constructed. Baseband Ethernet is employed over a single cable bus, which can be as long as 500 meters without repeaters. For distances greater than 500 meters or for greater area coverage, the baseband uses digital repeaters. According to the Ethernet IEEE 802.3 standard, a network can use three segments of 500 meters each and a one kilometer fiber optic segment, connected in series with repeaters in between.¹¹

A broadband Ethernet network (based on CATV technology), however, can extend this distance to 3.75 kilometers. The reasons for this increased distance can be found in the technology employed by broadband communications systems. A broadband system utilizes a single 75 ohm coaxial cable for the physical medium. The physical topology consists of a main trunk line that originates at the headend and traverses geographically-central areas. Branch lines extend from the

¹⁰ Abraham, Manachem E. "Running Ethernet Modems Over Broadband Cable". Data Communications, p. 199-212. May 1986.

¹¹ Leong, John. "Guide to Ethernet Installation and Interconnection". Data Communications, p. 267-276. September 1985.

trunk to more remote areas. From these branch lines are connections to devices made by means of drop cables. The headend serves at the origination point of all radio-frequency (RF) signals transmitted across the entire network, and the destination of all RF signals generated by devices attached to the network. Any device connected to the network may transmit signals over the network to other devices. However, these signals do not go directly to the destination devices; these transmitted signals go first to the headend where they are retransmitted back to the network. Utilizing this scheme allows the headend to prevent interference among signals travelling on the system.¹²

The ability of the coaxial cable to simultaneously carry two-way traffic (bidirectional) is achieved by dividing the cable's frequency spectrum into three main portions: (1) the forward band which carries signals from the headend to the attached devices, (2) the reverse band which carries signals from the the attached devices to the headend, and (3) the guard band which carries no signals and serves to separate the forward and reverse bands. Signals transmitted by an attached device propagate via the cable plant in the reverse band to the headend where they are converted to a higher

¹² Broadband Network Technology. Sytek, Inc., 2nd ed. New Jersey: Prentice-Hall, c 1986.

frequency (translated) and placed in the forward band for transmission to a receiving station. The division of the cable's frequency spectrum in megahertz (MHz) can be into one of the following configurations: (1) Subsplit (5-30 MHz Reverse and 54-400 MHz Forward); (2) Midsplit (5-112 MHz Reverse and 168-400 MHz Forward); and (3) Highsplit (5-174 MHz Reverse and 232-400 MHz Forward).

The usable bandwidth of the forward and reverse bands is further subdivided into six MHz channels. Each forward channel is paired with a reverse channel through the use of a translator and serves as the transmit and receive medium for each device attached and operating in those channels.

Because broadband technology is based on RF techniques and equipment, the broadband cable plant contains an array of active and passive elements, among these being amplifiers, splitters, taps, and filters. The amplifiers utilized in this system require ac power, which is derived from the coaxial cable plant via attached power supplies. These supplies provide 60 Volts alternating current (ac) to the coaxial cable through devices called power combiners. These devices permit the injection of power in either direction with little effect on the RF signals. In this way, amplifiers can be cascaded to allow for greater physical extension of the network. This is the primary reason

Ethernets employing broadband technology can span greater geographical distances than can their baseband counterparts. Figure 3 illustrates a typical broadband CATV network.

B. Data Link Layer

1. Network Access

The network access method governs how the devices attached to the LAN communicate with one another. If all devices on the network attempted to send data over the medium whenever they wished, confusion would be the result. In this light, rules must be established by which devices access the network media. These are called Network Access Protocols (NAPs) or Media Access Control (MAC) methods. Media Access Control is one of two sublayers that make up the data link layer of the OSI Reference Model. The second sublayer is the Logical Link Control (LLC) sublayer. The data link layer is responsible for what happens when the network is booted up, how a station joins/exits the network, which station transmits next, what happens when a station fails or performs incorrectly, the format of data frames to be transmitted, error checking, and error recovery.¹³ This layer links the upper layers with the physical layer.

¹³ Zimmerman, Hubert. "OSI Reference Model-The ISO Model of Architecture for Open Systems Interconnection". IEEE Transactions on Communications, vol. 28, issue 4, p. 425-431, April 1980.




FIGURE 3
TYPICAL BROADBAND CATV NETWORK

The sublayers work together to place data from the upper layers onto the physical layer in a structured way. There are two media access control methods: contention and polling. The solution presented in this paper focuses on the contention method. The contention method employed is Carrier Sense Multiple Access with Collision Detection (CSMA/CD). The access method is initiated on the transmitting end when the LLC requests transmission of a frame and the MAC constructs the frame from the data supplied by the LLC. The frame consists of eight elements and is illustrated in Figure 4. The Media Access Management component of the MAC sublayer then checks the possibility of contention with data traffic already on the medium. It monitors the carrier sense signal provided by the physical layer, and when the medium is clear, transmission begins. The MAC sublayer sends a bit stream to the physical layer for transmission, where this bit stream is converted into electrical signals on the medium. After transmission is complete, the MAC sublayer informs the LLC and then the MAC sublayer waits for the next transmission request.

On the receiving end, the arriving frame is detected by the physical layer at the receiving station. The physical layer synchronizes with the incoming preamble, and turns on its carrier sense signal. After discarding the overhead (i.e.

Preamble	Start Frame Delimiter	Destination Address	Source Address	Length or Type Field	Data Field	Pad Field	Frame Check Sequence Field
-----------------	--------------------------------------	--------------------------------	---------------------------	-------------------------------------	-----------------------	----------------------	---

Preamble:	7-byte field—Allows physical signaling circuitry to reach steady-state.
Start Frame Delimiter:	Uses bit pattern 10101011 to indicate start of frame.
Destination Address:	6-byte field—Specified address for which the frame is intended.
Source Address:	6-byte field—Identifies station from which the frame was initiated.
Length or Type Field:	2-byte field—Indicates number of LLC bytes in data field or protocol type utilized in packet (i.e., XNS, TOP/IP, DECNET, etc.).
Data Field:	Contains data being transmitted.
Pad Field:	Optional field which adds bytes to the data field to meet frame sizes specified by the standard.
Frame Check Sequence Field:	CRC techniques to specify if frame is damaged.

FIGURE 4

IEEE 802.3 FRAME FORMAT

preamble, start frame delimiter, etc.), the data is passed to the media access management component of the receiving station. This component collects the data bits until the carrier sense signal goes off. At this point, the data is passed to the MAC sublayer. The MAC sublayer checks the frame destination address field for the correct address and inspects the frame check sequence to determine if the frame was damaged enroute. If the address and data are correct, the frame is passed to the LLC sublayer; otherwise, it is discarded.¹⁴

When more than one station transmits simultaneously (when two MAC sublayers detect that the medium is clear), a collision occurs. The two stations' transmissions overlap, causing the signals to interfere with one another. The Media Access Management component of the MAC sublayer transmits a bit sequence, called a jam signal, to enforce the collision. This ensures that the duration of the collision was sufficient to be noticed by all stations involved in the collision. After transmitting the jam signal, the MAC sublayer terminates the transmission and attempts a retransmission in accordance with normal transmission procedures. Retransmission is repeatedly attempted if repeated collisions occur. Based on a

¹⁴ Digital's Networks: An Architecture with a Future. Digital Equipment Corp., c 1982.

mathematical algorithm, the transmission is abandoned after a number of attempts, and the transmitting station's MAC sublayer assumes that the medium has either failed or is overloaded.

The principle disadvantage to this access method is that it is not possible to calculate exactly how long a station will have to wait to transmit data. As the network traffic increases, transmission timing becomes more unpredictable (non-deterministic). This is an important point, because some manufacturing applications need fixed-time response and priority transmission capability. CSMA/CD networks cannot always meet this criteria when they are loaded up to the maximum with data traffic.

2. Logical Link Control

The second part of the data link layer is the logical link control (LLC).¹⁵ This sublayer is independent of the specific underlying MAC sublayer and provides queuing services for the applications running each station. For every application that needs to transmit and receive messages over the network via the MAC, the LLC sublayer

¹⁵ ANSI/IEEE Standard for Local Area Networks 802.2: Logical Link Control (LLC). IEEE Computer Society, 4th ed. New York: IEEE Press, c 1986.

provides a service access point where messages wait to be picked up or deposited. There are two classes (types) of service offered by the LLC sublayer:

LLC Type 1 service is called connectionless (datagram) service that permits the network layers of two applications to exchange messages without first making a specific data link connection. The receiving station does not have to be notified before the message is sent. However, this also means that the sending station nodes do not know whether or not the receiving station is capable of accepting the message. There is also no expectation of an acknowledgement from the receiving station or any guarantee that the messages will arrive in the order they were sent.

LLC Type 2 service is called connection-oriented (virtual circuit) service and provides synchronization, sequential delivery, and error recovery using acknowledgments and retransmissions. LLC Type 2 uses additional fields for sequence numbers and connection identifiers.

3. Network Performance

Network performance is usually thought of in terms of throughput and is dependent upon factors such as number of devices attached to the network, bit rate of the cabling, the application being utilized, the efficiency of the protocol implementation, and the network topology. Figure 5 shows network throughput as a function of the traffic on a CSMA/CD network. Throughput increases as traffic generating devices are added to the network, reaching a peak as the network nears saturation. Beyond the peak is a zone (Zone 4) where adding new traffic to the system will actually cause throughput to decline. This is caused by the increased traffic causing more collisions on the network. Each collision results in dead time on the media during the jam and wait to re-transmit periods; the more collisions, the lower the throughput of the system. As traffic increases, it becomes advantageous to segment the network through the use of data link bridges.¹⁶

Network delay is the amount of time required to access the network and transmit data across it to a destination node. Figure 6 illustrates the relationship between network delay and the number of devices on the network. As devices are

¹⁶ Tobagi, F. "Performance Analysis of CSMA/CD". Computer Networks, p. 245-259, October/November 1980.





FIGURE 5
THROUGHPUT AS A FUNCTION OF NETWORK TRAFFIC




FIGURE 6
NETWORK DELAY AS A FUNCTION OF NETWORK TRAFFIC

added, any single device's ability to access the network increases in difficulty. The network induced delay climbs gradually until it reaches a point near 20% of the network channel capacity (critical capacity point). Beyond this point, the network delay increases exponentially. As with the issue of throughput, partitioning of the network via data link bridges can reduce network delay.

C. Connectivity Devices

There are many products designed and marketed for the purpose of internetworking. The three most prevalent devices used in the internetworking area are repeaters, bridges, and routers. Each of these devices share the characteristics of perpetuating and handling electronic signals, yet each operates at different levels of the OSI model. Repeaters are the simplest of these. Operating at the physical layer of the network model, they boost, reconstruct, and retiming data signals to overcome distance limitations.

Within the Ethernet environment, bridges filter traffic by storing node source addresses on either side of it and forwarding only those packets that must pass through the bridge to reach their destination. The packets may be routed locally between two LANs or remotely by way of a wide

area network (WAN) link between two LAN segments. Bridges operate at the data link or media access control (MAC) layer of the OSI model and are protocol independent.¹⁷ Most are self-learning in that they require no previous knowledge of the network segments being bridged. As the bridge learns, it creates routing tables of the logical addresses of the devices on each of the attached networks. Each entry in the routing tables lists a logical address, the network on which the device with that address can be found, and a time stamp. The time stamp indicates how long it has been since a packet was received from that device. Once the address table is built, the bridge is ready to operate. When it receives a packet, it examines the source address (filtering). If the source address is not local, it will forward (copy the packet to the LAN on the other side of the bridge). High filtering and forwarding rates are critical because of the bridge's transparency. If the bridge delays the packet too long, the sending device will assume the packet never reached its destination and retransmit.

A bridged network is still a single logical network, but with optimized traffic segmentation. Bridges may operate between different media types and most offer a range of features. The most basic bridges simply filter data

¹⁷ Data Link Bridge Principles of Operation. Ungermann-Bass, Inc., Santa Clara, Ca. February 1987.

utilizing destination addresses. Some of the more sophisticated bridges offer selective filtering on any data field contained in the incoming packets, including source addresses, packet size, and packet type (identification of protocol being used).

Remote bridges utilize telecommunications links, fiber optic links, and broadband networks to connect two LANs. The use of these links in conjunction with remote bridges introduces the concept of load sharing, which is the ability to utilize two different paths between LANs simultaneously to transfer data. Remote bridges perform load sharing by utilizing backup links to increase bandwidth. These bridges send packets to a single destination along alternate paths. This load balancing evenly distributes the traffic over the two attached remote links. If one bridge fails, the other bridge carries all of the traffic.

In order to make a bridge robust enough for remote networking, the IEEE developed an algorithm called spanning tree. The IEEE 802.1 algorithm makes it possible for a bridged LAN to include loops which are illegal in Ethernet's bus topology. Spanning tree makes it possible to have parallel bridges so a backup bridge can take over in case the primary bridge fails. According to the spanning tree algorithm, one bridge is designated the primary bridge; this

bridge is the only one to pass traffic. If it fails, the traffic is automatically shunted to a parallel bridge. Bridges that implement spanning tree are unable to utilize load sharing.¹⁸

Routers utilize the internet protocol of the OSI model network layer to direct traffic to the appropriate network. Routers are inherently protocol-dependent and must support each protocol being routed. How routers perform their routing function is largely dependent on the network protocol with which they operate and the particular implementation of this protocol. In the enterprise network that is partitioned by routers, each partition is given a different network number, and each is a fully independent entity that must be managed separately. Like bridges, routers may automatically learn changes in the network's topology though this depends on the ability of the network's protocol to pass routing information between routing nodes. Unlike bridges, however, bringing up a router-connected network requires the router to know the various network numbers up front.

¹⁸ Draft IEEE Standard 802.1, Part D: MAC Bridges. IEEE Computer Society, revision B, May 1987.

Routing, the primary function of a router, can be static or dynamic.¹⁹ With static routing, the paths to reach the different networks are configured manually in the system by the network manager. But with dynamic routing, new paths can be learned from the network. Though static routing allows a great level of control by the network manager, it is not well suited for large networks where topology is often subject to change. Static routers may require being taken out of service while being reconfigured. Dynamic routing improves the overall quality of network service by minimizing or eliminating downtime for changes or upgrades while also facilitating redundancy in the network. However, it relies on protocols to convey the routing information around the network. This automatic update process may require a long period of time for a large network; additionally, if several topology changes occur at the same time, the accuracy of the routing updates may be in question for some period of time until the network reaches steady-state in terms of changes and additions. This illustrates the complexity of the dynamic routing process and further illustrates the desirability of bridges in linking and extending networks.

¹⁹ Boule, Richard. "Inside Routers: A Technology Guide for Network Builders". Data Communications, p. 53-66. September 1989.

The performance of bridges is generally higher and offers greater throughput than the router. A router introduces delay. It, like a bridge, is a store-and-forward device; the major difference is that it must process several layers of protocol to forward packets toward their destinations. Because it is a resource shared by multiple stations, it can be a bottleneck in the network, especially if it communicates with a remote network through a slow interconnection medium.

D. Network Management

1. Management Functions

A major concern in a network that utilizes bridges is that of management. Since the bridges extend the LAN, how can the functionality be controlled and managed? In general, management of a network which utilizes bridges must be evaluated in terms of network operation, capacity planning, information management, accounting, and security.²⁰

The functions of network operation for bridged networks consist of the following areas: name service, configuration, fault isolation, network monitoring, network restoration,

²⁰ Sylor, Mark W. "Managing Phase V DECnet Networks: the Entity Model". IEEE Network, p. 30-36, vol. 2, issue 2, p. 30-36, March 1988.

network tuning, and information transfer. Name service for the network relates a logical name to a specific network address or set of addresses. It is administered by the network manager and allows easy construction of routing tables in bridges. The configuration function controls the operations of the network in terms of baud rate, active/inactive link for a remote bridge, etc. The function also controls the address table at each bridge and determines what capability the nodes which utilize the bridge will have to access the network. Fault isolation relates to insuring that a fault occurring in one segment of a LAN does not bring down the entire network. The bridge should automatically isolate the fault and allow the network manager to act on problems before they inconvenience the user. The monitoring function focuses on a bridge's ability to gather statistics on network activity. This information is typically forwarded to the bridge management system. The network manager has the ability to program threshold levels for various parameters in the bridge; if the bridge or traffic through the bridge exceeds the level, the bridge management system alerts the network manager who can take corrective action. Network restoration involves the identification of the problem, and its subsequent resolution. This resolution is in the form of rerouting circuits, substituting hardware and/or software, and other functions required to bypass faults or possible faults.

Monitoring of the network provides information as to whether the bridges are being underutilized or overutilized. If the bridge management system determines that the bridges are not operating at their optimum levels, the network manager should take remedial action in order to optimize the network. Information transfer is for those large networks that consist of multiple bridged networks, there will be a need to transfer information in a timely and efficient manner. This includes the ability to download routing tables across the bridged networks; establish reporting techniques and bridge tuning across bridged networks; and the ability to communicate with another vendor's management system.

2. Capacity Planning

The function of capacity planning consists of the following: monitoring, statistics gathering, trend analysis, and modeling. As an organization's network evolves, management needs information-gathering tools to determine how the network is being utilized. The monitoring functions of the bridge management system should center around response time, effect of new hardware and software, and the effect of additional users or channels on a network. This management system should have the ability to gather, analyze, and report bridge and network statistics. These statistics

include all information which is relevant to capacity planning for channels, bridged networks, and the network backbone. The statistics gathered must include information on channel loading, session loading, traffic loading per session, and collision-related information.

The bridge management system must have the capability to track trends in the bridged network. The statistics gathered from the daily operation can provide historical performance information for the bridges and the attached networks. This information can detail the types of problems plaguing a network and how long it takes to isolate and solve them. The bridge management should also have some modeling capabilities that can help the network manager define the limits of a network. This allows a user to determine how the network should be designed and also to predict the bridged network's performance in terms of service information such as percent of bandwidth usage, response time delay, and worst case routing.

3. Information Management

The function in information management consists of the following: user information, component information, topology diagrams, and reporting. The database of the bridge management system should provide both administrative

and technical data on the bridges installed in the network; this information must be accessible by selected persons, sortable, and customized for reports. As the network grows, the network manager will need to display, retrieve, and modify information regarding the bridges installed on the network. This includes both administrative and technical data that encompasses the network backbone, bridged networks, and individual bridges that make up the network. As the network becomes more complicated, graphic displays of the network backbone, local area networks, and wide-area communications links will become very important. This type of feature provides a one-screen overview that aids in planning the network and is useful for trouble-shooting purposes. The bridge management system should be capable of generating extensive reports from its database, threshold levels, and performance modeling functions. These reports should be available in a detailed written format for analysis by the network manager.

4. Accounting

The accounting function of the bridge management system should consist of the following capabilities: audit trails, billing, and capital management. The bridge management system should display an audit trail for every network connection. The major areas displayed by device must

include destination, source, user, start and stop time, packets transmitted and/or received, errors, and network resources used. The management system must automatically and transparently use the information stored in the audit trail and may have the capability to charge by user, department, and/or billing code. This system should also manage the assets of the network, including the capitalization of the network equipment, depreciation, and reports on both apparent and hidden costs of establishing and maintaining a network.

5. Security

The bridge management system must provide security for the network manager, in such areas as access control, session control, and multi-level security. Access control should be the first level of support provided. The basic functions provided are login and password protection on the network backbone, its attached networks, and the bridge management system itself. The session formation must be controllable from a central console. If a host port does not have a session, the network device should have the capability of alerting the host to logoff the user. If a session is left active at a node, the network node should timeout, close the session, and logon the next user to the device.

The bridge management system should have the ability to confine packets to a given perimeter. This system has the ability to subdivide the network and ensure that predefined sessions and/or packets do not leave a given area. The ability to access the network must also be controlled by the bridge management system; this is accomplished by defining either by user, device, and/or privilege. The management system must also have the ability to limit access to its database.

6. Other Requirements

Other major requirements of the bridge management system include communication with multiple bridged networks, reliability, cost effectiveness, ease-of-use, and modularity. Large networks consist of a backbone network with multiple networks attached to the backbone via bridges. These networks may all be products of either a single vendor or of multiple vendors; the important issue is that the different bridge management systems should have the ability to communicate with each other. If a bridged network does not have its own management system, the predominant bridge management system must provide this service. The management system should be highly reliable and the system should also have the capability of managing itself. The network management price should not outweigh the features it

provides to the network. The total expenditure of the management system should cost no more than a small percentage of the total network equipment cost. From the network manager's perspective, the management system should be menu-driven and not require any system expertise to operate the system; from the user's perspective, the system's operations should be transparent and as unobtrusive as possible. The most important feature of the management system is that the day-to-day operations should be automated so that the management system runs the network predominately by itself. It should also be modular so that different features can be obtained and installed as they are required.

E. Description of Digital Network Architecture and DECnet

1. Digital Network Architecture

The Digital Network Architecture (DNA)²¹ is a layered family of architectures which defines how DECnet networks operate. This layering is the same structure as defined in the OSI reference model. DNA has evolved over a series of phases, each incorporating new technologies, such as Ethernet LANs (Phase IV), packet switching (Phase III), etc. As a network architecture, DNA outlines a broad range of hardware, software, and protocol implementations designed to connect

²¹ DNA is a registered trademark of Digital Equipment Corporation

Digital Equipment Corporation (DEC) computers together, as well as multivendor's devices.²²

DECnet is the implementation of DNA in a series of software packages for various DEC processors. DEC's current implementation of DNA is Phase IV, which incorporates Ethernet local area network technology into DECnet. DECnet provides transparent networking capabilities to the end user. Any DECnet node (a processor with DECnet software) can initiate and participate in terminal-to-terminal, program-to-program, task-to-task communications with another DECnet node.

Most nodes are capable of downline-system or task loading for execution at remote sites. This system also allows for a host node to provide resources and services for one or more remote nodes. A host may store programs in a local file system which are executed on the remote host, and later restore them to the remote's main memories. The Ethernet capabilities of Phase IV allow DECnet nodes to communicate via coaxial cable at a maximum data rate of 10 million bits per second (Mbps). Using the DECnet facilities, nodes on an Ethernet network can perform any of the high-level network functions (program-to-program or host-to-host file

²² DECnet Digital Network Architecture Phase IV. Digital Equipment Corp., c May 1982.

transfer). A DNA network consists of two or more processors loaded with DECnet software. The DECnet software, as illustrated in Figure 7, is a layered architecture similar to the OSI model. A very important feature of the DECnet network is the peer relationship each node has with one another. This allows each node within the network to communicate with every other node without having to consult a central controlling node. The advantage of this decentralization is that it allows each node to be equally responsive to user requests. A network user at a node can easily access applications and facilities on other nodes within the network. Without having to go through a central node for data exchange, resource sharing, and remote access operations, communications overhead is reduced and the network's performance increases in efficiency.

2. Protocol Structures of DNA

The DNA structure, being hierarchically layered, supports communication between adjacent layers within a single system by using architecturally-defined interfaces. In addition, communication takes place between computer systems by exchanging messages following architecturally-defined protocols. Each DNA layer contains a protocol specification which defines the messages to be exchanged between layers and the procedures for sending and receiving those messages.

ISO LAYERS	DNA LAYERS	DNA FUNCTIONS	DNA CAPABILITIES
Application	User	Contains user supplied functions and programs.	Network Applications
	Network Management	Uses Network Information and Control Exchange (NICE) protocol for tasks such as downline loading, upline dumping and testing.	Peer-to-peer Network Management
Presentation	Network Application	Uses Data Access Protocol (DAP) for remote file access and transfer. Uses loopback Mirror Protocol for loopback testing.	Generic Network Functions, File Access, and Terminal Access
Session	Session Control	Uses Session Control Protocol for sending and receiving logical link data and disconnecting logical links.	Network/Node Addressing/Naming, Access Control, Name-to-Address Translation.
Transport	End to End Communication	Uses Network Services Protocol (NSP) to handle all system-independent aspects of creating logical links, such as data flow control.	Task-to-Task Communication, Logical Link handling, End-to-End Error control, Flow control, and Segmentation.
Network	Routing	Uses Transport Protocol to route user data in packets to its destination and control congestion on the lines.	Dynamic Message Routing, Packetization Congestion Control
Data Link	Data Link	Uses DDCMP, X.25, or Ethernet protocol. Makes the connection between adjacent nodes and ensures accuracy and sequencing of messages.	LAN/WAN Interconnect, data integrity.
Physical Link	Physical Link	Manages the physical transmission of data over a channel. Functions include monitoring channel signals, clocking on a channel, and handling interrupts from the hardware.	Communications facilities interface physical transmission of data.

FIGURE 7
DECnet ARCHITECTURE

From an architecture perspective, a computer network is divided into a set of computer systems, each system being divided into layers. Each layer contains one or more protocol modules. Figure 8 illustrates the case of a two node network in which communication between modules takes place horizontally, using these protocols, a bridge, and a router. A message sent from one application to another passes through each layer twice; once in departing the local processor and again upon entering the remote processor.

The DNA architecture specifies the functional interfaces between layers. These interfaces define the services each layer provides to higher layers and thus firmly partitions the architecture into layers. The interfaces defined by the standard are abstract specifications, concentrating on the functions to be provided, not on the form of the interface implementation.

The details of interface realization are left to the implementers of each DECnet product. These details typically depend on the structure and conventions established in each operating system. DNA provides a precise specification of the protocols between layers. This precision is necessary to ensure that separate DNA implementations can interoperate. In some cases the details of protocol operation are left to the implementers. If





FIGURE 8
BRIDGE AND ROUTER OPERATION MODELS

incompatible choices are possible, the standard specifies rules for negotiation to select compatible options. The DNA specifications include model protocol implementations with which all valid implementations must communicate correctly.

Protocols in the DNA architecture are strictly layered. These are peer protocols that define the relationships between modules in the same layer of two computer systems. The protocols in each layer operate independently, with communications between layers taking place only across the defined interfaces. This structure allows the independent replacement or addition of protocols in each layer as the architecture evolves. It should be pointed out that strict layering is an architectural, not necessarily an implementation, concept. Experience in implementing DNA has shown that significant performance improvements are possible by collapsing the implementation of several layers into fewer layers. However, this performance improvement is obtained at the cost of reduced modularity and flexibility of that particular implementation. Each layer is discussed below in terms of its functions, interfaces, and protocols.

A. User Layer

The user layer is where the message originates as a result of an application program or as the result of direct input

from a user. This data may be simply a command or an extensive file destined for a remote processor running another application. This layer contains user-defined functions, such as applications programs. Application programs in the user layer can directly access the session layer for program-to-program communication. They can also use the application layer to gain access to common network services, such as remote file access, data base management, or network management.

B. Network Management Layer

The network management layer provides methods that define the functions used by the network manager to plan, control, and maintain the operation of the network. Unlike the other layers of DNA, this layer may interface with any or all of the layers below it in order to gather data regarding network performance or to change operational parameters. The network management modules within a node are responsible for two functions: (1) they coordinate the management of that particular node, and (2) they communicate with peer management modules in remote nodes as needed to accomplish decentralized management. The network management module uses the services of the lower layers to provide its functions. This layer also supports the following functions: (1) down-line loading of files to a remote

system, (2) testing of links at both the data link and the logical link levels, and (3) setting and displaying the states of links and nodes.

C. Network Application Layer

The network application layer appends control information regarding the message's function to the user data. This layer is used to provide common network services, such as remote file access, remote file transfer, remote system loader, and a network command terminal function. The remote file access and file transfer functions performed in this layer are accomplished through the Data Access Protocol (DAP) which allows these functions to occur independent of the input/output (I/O) structure of the operating system being accessed. This module has a special interface to the data link layer so that simple management functions can be provided between adjacent nodes. The functions include remote restarting, down-line and up-line loading, and loopback testing.

The network application layer provides additional generic services to the user and network management layers. These services include remote interactive terminal access and gateway access to non-DNA systems. Modules in this layer operate independently and asynchronously. A single DECnet

node may support many different network applications modules, which communicate using many different protocols. Two cooperating network application layer modules exchange messages using the DNA session control service; the user module at the node requests the file operation, and the server module acts on the user's behalf at the remote node.

D. Session Control Layer

This layer provides system-dependent, process-to-process communication functions for processes residing in the user, network management, and network application layers. These functions bridge the gap between the pure communication functions provided by the end communication layer and the functions required by processes running under an operating system. The communication service provided by the session control layer is connection-oriented; an initiating process requests a connection to a destination process. The session control layer manages these connections. Once a connection is established, data flows between the processes without further intervention by the session control layer, using the facilities provided by the end communications layer. When establishing a connection, the higher layer specifies the destination process in two parts: (1) by destination node and (2) by process within the destination node. This layer also performs the following additional functions: (1) the

establishment of a node name mapping table that defines the correspondence between a node name and a node address, (2) the creation of new processes or activation of existing processes to handle incoming connection requests, and (3) the validation of incoming connection requests.

E. End Communication Layer

The end communications layer, residing immediately above the routing layer, provides a standardized communication service, used by the higher layers of the DNA software. This layer provides a system-independent process-to-process communication service which is reliable, sequential, and connection-oriented regardless of their locations in the network. This layer, in conjunction with the session control layer, isolates the higher layers from any transient errors or reordering of data introduced by lower layers. It also provides a multiplexing function, enabling multiple connections to be established between pairs of nodes or between a node and multiple nodes (called logical links). This layer utilizes the network services protocol (NSP) to provide the logical link service to the session layer, exchanging protocol messages using the routing layer. The originating session control module requests the local NSP module set up a logical link to a remote session control module. If the remote node is accessible via the routing

layer and has resources to support an additional connection, the remote NSP module will indicate its desire to connect to the remote session layer. The NSP module will transfer session control protocol data, such as user identifiers, passwords, and the DECnet object type. The remote session control module can either accept or reject the link. Only when the logical link is finally established can it support the flow of data.

F. Routing Layer

The routing layer handles message addressing and provides a network-wide message delivery service. This layer accepts messages from the end communications layer in a source node, extracts the address and routing information necessary to forwards the packets, possibly through intermediate nodes, to a destination node from a predefined network parameter table. The routing layer implements a datagram service, which delivers packets on a best-effort basis. This layer makes no absolute guarantees against packets being lost, duplicated, or delivered out of order. Such guarantees are made by the end communications layer. To provide this network-wide service, the routing layer determines routes by using an adaptive, distributed algorithm that responds to changing network configurations and uses these routes to forward packets. In the forwarding process, the routing

layer must attempt to avoid or at least control any congestion that results from overloading the network with excessive traffic. The routing layer also ensures that packets do not wander around for too long before being delivered. Routes between all sources and destinations are automatically calculated or recalculated by the routing layer whenever new nodes or links are added to or removed from the network. These changes to the network topology can either be planned by the network operators or may result from the unplanned failures and recoveries of network nodes and links.

Routes are calculated on the basis of link costs, which typically are inversely proportional to link speed. The route to each destination node is along a path having the minimum total path cost, which is the sum of the link costs along the path. Link costs can be set either by network managers, if desired, or by using default values. This feature, called adaptive routing, is a key aspect of the DNA software and makes it easy to operate large DECnet networks. Without adaptive routing, network operators and users would have to calculate paths between all pairs of data sources and destination, including alternative paths to handle failures.²³

²³ DECnet Digital Network Architecture Phase IV. Digital Equipment Corp., c May 1982.

The routing layer forwards packets based on a uniform addressing scheme. Each node in a DECnet network is assigned a unique address, used by the routing algorithms to calculate routes. The addresses indicate each packets's source and destination and guide the forwarding decision. A uniform addressing scheme allows the higher layers and network applications to treat the network as a uniform resource.

The routing layer supports a variety of communications facilities for communicating between adjacent nodes. A complete path from a source node to a destination node can use a mixture of link types. Three main types are supported: (1) dedicated links using the Digital Data Communications Message Protocol (DDCMP)²⁴, (2) packet-switched networks, and (3) Ethernet LANs. The dedicated links using the DDCMP protocol are treated by the routing layer as point-to-point virtual circuits; hence, in these networks, end nodes have a particularly simple task on these types of links since end nodes make no decision when sending a packet; they simply send it on the link to the adjacent node.

²⁴ DDCMP is a registered trademark of Digital Equipment Corporation

End node routing is more complex on Ethernet LANs since each station can send to any other station on the Ethernet; therefore, the end nodes attached to an Ethernet must make a routing decision. When sending to nodes remote from their Ethernet, the end nodes must send to a router. When sending to nodes on their Ethernet, the end nodes send directly to the destination node. End nodes follow a simple procedure to determine which path to follow. If a router is present and the end nodes do not know about a particular destination, they forward their packets to the router. End nodes learn that particular nodes are on their Ethernet by receiving packets directly from those nodes or by being informed by the router.

G. Data Link Layer

The data link layer creates a communications path between adjacent nodes. This layer frames messages for transmission on the channel connecting the nodes, checks the integrity of received messages, manages the use of channel resources, and, when required, ensures the integrity and proper sequence of transmitted data. Currently, there are three protocols residing in the DNA data link layer: (1) DDCMP, (2) packet switching (X.25), and (3) Ethernet.

DDCMP operates over synchronous or asynchronous communications links. It can operate in point-to-point configurations. DDCMP is a byte-oriented protocol which provides a count of the number of bytes that will be sent in the data portion of each data message sent and uses a 16-bit cyclic redundancy check (CRC-16) to detect errors in headers or user data. On half-duplex channels, DDCMP executes link allocation procedures to ensure that two stations do not conflict in their use of the channel. These techniques are based on polling in which one station extends permission to the other to transmit.

The Ethernet LAN provides a communications facility for high-speed communication among computers located within a moderately sized geographic area, such as a building or a campus. This LAN includes a data link layer and a physical layer, which can transmit data at a maximum rate of 10 Mbps. The Ethernet has a maximum station separation of 2.5 kilometers with a maximum of 1024 stations. A shielded coaxial cable is used as the physical medium.

Ethernet also uses a branching, nonrooted tree topology. Messages, called frames, are transmitted over the physical channel in a broadcast fashion. Each frame contains a source address and a destination address, with each station assigned a logical address. A frame can be addressed to an

individual station or to a group of stations, using a group address (called a multicast address). A special multicast address is used to denote the set of all stations on an Ethernet and is typically used for maintenance purposes. In the DNA architecture, the multicast capability is used for network configuration purposes by the routing and network management layers. End nodes advertise their availability to routing nodes by periodically broadcasting "hello" messages to the multicast address.

The Ethernet data link protocol frames messages using the properties of the Manchester coding scheme employed by the physical channel to mark the beginning and end of each frame. In addition to source and destination addresses, frames employ a protocol type field to identify the higher-level protocol carried in the frame. The protocol type field values are assigned in blocks to all vendors who manufacture Ethernet, thus permitting different proprietary and public protocols to coexist in a single Ethernet station. Ethernet frames also contain a 32-bit CRC to ensure that frames received in error are detected and discarded. Since the Ethernet physical channel can transmit data only from one station at a time, the Ethernet data link protocol must allocate the single channel among all the stations. This allocation is accomplished by utilization of CSMA/CD. As described previously, this contention-based

protocol forces stations to "listen" before transmitting (carrier sense) and defer their transmissions to other stations already transmitting. Should several stations begin transmitting simultaneously, a collision will occur, preventing correct reception of any transmission. In this case the physical channel hardware in each colliding station will detect the collision (collision detection) and each station will reschedule transmission after a randomly selected delay. To ensure efficient, stable operation of the network under both low and high load conditions, this random delay is adjusted on subsequent collisions by the back-off algorithm. This causes each station to reduce the load presented to the network under overload conditions.

H. Physical Link Layer

The physical link layer transmits bits of information between adjacent nodes. The functions in this layer include encoding and decoding signals on the connecting channel, performing clock recovery of received signals, and interfacing the communications channel to any processor and memory used to implement higher-level protocol functions. Implementations of this layer encompass hardware interface devices and device drivers in operating systems, as well as communications hardware such as bridges, transceivers, and the physical channels themselves.

Protocols for the physical layer are rudimentary, emphasizing the specification of electrical interfacing parameters. No physical layer specifications have been developed for DNA. Instead, it relies on industry-standard methods for the physical layer, thereby ensuring that DECnet products can operate over available communications technologies and infrastructures. Physical layer standards include Electronic Industry Association (EIA) RS-232 and RS-423, baseband implementations, and a broadband implementation.

III. PREPARATION OF A TECHNICAL SPECIFICATION

A. General Discussion

A technical specification is a mechanism by which a product to be procured is described in detail. This product can be material or a service. The technical specification is an integral and necessary part of a purchase agreement between two or more parties. However, in itself, it is not a contract for procurement. In this case, the specification is used to define the requirements of non-standard or standard equipment that is of sufficient complexity to require a basis for the evaluation of a prospective vendor's proposal, capabilities, and determination of conformance to requirements. The specification provides the following functions:

- 1.) Protects the requesting Companies interests
- 2.) Defines precise technical requirements
- 3.) Minimizes costs by introducing competition
- 4.) Guarantees quality of workmanship and materials used
- 5.) Provides a competitive environment for business

It is critical for the specification to communicate to prospective vendors the soliciting company's requirements for the product. A well-prepared, organized, and planned

specification can reduce the likelihood of accepting nonconforming equipment and getting the desired product resulting in reduced monetary and time commitments. The specification must contain functional requirements which allow prospective vendors to design the product, testing requirements for determining compliance with the functional requirements, and acceptance criteria used in measuring that compliance. The following sections reflect the preparation of a technical specification that may be utilized in the procurement of data link bridges.

B. Requirements Definition

The definition of requirements for the specification begins with an introduction of what is expected in the way of the final product. The data link bridge shall provide the protocol-independent bridging for packets between multiple local area networks. The bridge is a store-and-forward switch that isolates traffic to only those LANs on which the traffic must appear. The bridge examines each packet above the media access control protocol layer. As stated earlier, the bridge examines each packet received from each bridged network and forwards the packet if the destination address is not on the source network.

The bridge knows the network topology through a table of logical addresses stored in the bridge for user nodes on the network. In the learning mode of operation, as new user nodes come on line, the bridge shall "learn" of their existence from the source address field in the packet and add the user node to the table of addresses.

At the option of the operator or on command from the Bridge Management Control Center (to be defined later), the "learning" function of the bridge shall be disabled and the address table shall be downline loaded. Any node can establish a data link connection with any other node on any of the bridged IEEE 802 networks even though it must traverse a broadband, fiber optic, or T1 backbone to do so. The fact that a data link may traverse one or more bridges over one or more local area networks (LANs) must be completely transparent to the end user. The equipment and services required by this specification shall include the following:

- a. All data link bridge configurations illustrated in Figure 9
- b. All hardware required to interface to a single cable mid-split IEEE 802.3 broadband system, an IEEE 802.3 baseband Ethernet system, a fiber optic cable system employing Ethernet signaling




FIGURE 9

DATA LINK BRIDGE CONFIGURATIONS

techniques, and a T1 system employing the DS1 framing format.

- c. Provide specifications for any special equipment required for remote bridge operation (i.e. broadband CATV translator, fiber optic repeater, etc.)
- d. Specific hardware and all software required for bridge management
- e. All software required for operation, including address table construction, routing algorithms, and diagnostics
- f. Design verification testing

C. Bridge Functional Description

1. Data Link Bridge Description

The bridge must be a passive filtering device; i.e. each bridge shall "listen to" (filter) all of the traffic being broadcast on each of the attached media. If the bridge determines that the frame is addressed to a Media Access Control (MAC) node that is located on the "same side of the bridge" as the source MAC node (i.e. the packet is local), then the bridge shall discard/ignore the packet. If the bridge determines that the packet is addressed to a MAC node that is not attached to the source network (the MAC node is on the "other side of the bridge"), then the bridge will

transmit the packet, making any adjustments to the content and format of the packet necessary to reconcile any differences in the framing requirements of the source and destination networks.

2. Function Required of the Data Link Bridge

The functioning existence of the bridge shall be transparent to all nodes attached to the associated networks, although the bridge may have a MAC node address assigned to it for the purposes of network management. The operation of each bridge must also be transparent to other bridges functioning on the networks, although some bridges may exchange network management data with each other; however, all interbridge communications must be transparent to other bridges and nodes on the associated networks.

The functional configuration required by this specification can be seen in Figure 10. The data link bridge shall connect any two logical IEEE 802.3 Ethernet networks together and allow communication between devices residing on either local area network. The bridge automatically buffers the data to match speed differences between the attached networks. This is necessary because the implementation on the broadband network is 5 Mbps versus the baseband Ethernet specification of 10 Mbps. The bridge must compensate for




FIGURE 10
SYSTEM DESCRIPTION

the difference. If the bandwidth of all the individual data links in a network is not uniform, congestion at the bridges may occur, resulting in loss of packets. Traffic arriving on a data link at a rate exceeding the bandwidth of the outgoing data link requires the bridge to buffer the excess. If the quantity of data to be buffered exceeds the capacity of the bridge, then data will be lost.

D. Network System Functional Description

Figure 11 illustrates the current network topology requirements under which the bridges must operate. The local network described is a single cable, mid-split 5-450 MHz broadband (CATV) system which consists of active and passive broadband components, and has a maximum trunk length (maximum distance from the headend to the outermost node) of three miles. The existing Ethernet baseband systems that will be attached to the broadband network via the bridges conform to the IEEE 802.3 specification and are based on coaxial, T1, and fiber optic cable media. The data link bridge forms a general purpose link between two networks and allows any vendor's Ethernet equipment to pass data across the bridge. Industry standard protocols Xerox Network Systems (XNS)²⁵, International Standards Organization (ISO),

²⁵ XNS is a registered trademark of Xerox Corporation




FIGURE 11
NETWORK SYSTEM CONFIGURATION

and Transmission Control Protocol/Internet Protocol (TCP/IP), as well as the proprietary protocol Digital Network Architecture DECnet, shall be supported. This is essential in that support of multiprotocol environment is a fundamental requirement in the implementation.

E. Bridge Physical Requirements

Local bridges are defined as those bridges which make an Ethernet-to-Ethernet connection directly or make an RF channel-to-RF channel connection on a broadband network. In the latter case the protocol is still the Ethernet protocol modulated onto a carrier. Remote bridges are those which make a connection from an Ethernet to an intermediate network, in this case a broadband channel, T1 network, or fiber optic network, to another remote bridge and Ethernet network. The modulation method, translation and/or remodulation at the headend of the broadband network, must be specified. The type of fiber optic repeater, if required for bridge operation in the case of the fiber optic network, must also be specified to satisfy the physical networking requirements.

The data link bridge shall be physically structured from no more than three physically separate devices which shall conform to the criteria in Table 1.

TABLE 1
BRIDGE PHYSICAL REQUIREMENTS

<u>Category</u>	<u>Specification</u>
<u>Mounting</u>	Industry-standard 19" rack mountable
<u>Power</u>	120 VAC (input variations to the unit will not exceed $\pm 10\%$), 60 Hertz at a maximum amperage of 3.0 amperes utilizing a NEMA 5-15 plug.
<u>Environmental</u>	0 deg. C to 40 deg. C (operating) -55 deg. C to 85 deg. C (storage) 0 to 95% relative humidity
<u>Interface</u>	Broadband: 75 ohm F-type connector (RG-59B/U coaxial cable compatible) Ethernet: 15-pin subminiature "D" Series connector with locking posts in accordance with IEEE 802.3 Fiber Optic: Amphenol 906 (SMA) or equivalent T1: V.35 34-pin female connector
<u>Data Rate</u>	Broadband: 5 Mbps per 6 MHz channel Ethernet: 10 Mbps Fiber Optic: 10 Mbps T1: 1.544 Mbps
<u>Bit Error Rate</u>	Less than 10^{-9}
<u>Status/Indicating Lights</u>	Power, Transmitted Data, Received Data, Collision Detect (Transmitted/Received Data and Collision Detect shall be separate for each cable segment) All LEDs shall be visible from the front panel
<u>Switches</u>	Power on/off, bridge reset, bridge self-test

F. Bridge Performance Characteristics

The performance characteristics of the data link bridge must meet the following criteria:

- a. IEEE 802.2 LLC Class I service shall be provided by the bridge.
- b. The number of network user addresses available in the bridge address table shall be 250 addresses on each side of the bridge in both "management" and "learning" modes.
- c. The number of packets forwarded through the bridge from one attached network to the other (packet throughput) shall meet or exceed 5,000 packets per second. The forwarding function is based on reliable receipt and retransmission of all anticipated intersegment traffic. It is related to packet size, as packets are copied from one network interface to another. This rate is established in relation to the average and maximum packet sizes for internetwork traffic anticipated. The 5,000 packets per second rate is for the minimum packet size (64 bytes).

- d. The bridge shall filter packets between attached networks at a minimum rate of 7,000 packets per second. The maximum anticipated usage of the network is often a small fraction of its theoretical limit. The 10 Mbps Ethernet network can carry minimum size (64 bytes) at a maximum theoretical rate of 14,880 packets/second. Larger packets significantly decrease the packet rate. The 7,000 packet/second steady rate represents a 100 percent load using 157 byte packets or a 47 percent load using 64 byte packets. This steady rate should be sufficient in the network topology on which these bridges will operate.
- e. The end-to-end delay of all IEEE 802 packets over the bridged networks (including the longest path with the most hops) shall not exceed three seconds. Because some protocols have not been designed for internetworking, or assume a fast delivery of packets, the bridge delays may affect the traffic enough to cause problems. Typical symptoms include network and bridge congestion as transmitting stations retransmit new copies of packets before the initial copies have even reached their destinations. Some protocol implementations will adapt dynamically to the delay incurred in the internetwork while others

may be configured to adjust their tolerance level. Timing dependant high-level protocols (i.e. DECnet, TCP/IP, and XNS) must be properly tuned in order to tolerate the maximum packet latency introduced by the worst case concatenation of bridges. The characteristics of this delay were illustrated in Figure 6.

- f. The radiated electromagnetic susceptibility of the bridge shall be 10 volts per meter over the frequency range of 20,000 Hz to 1,000 MHz. The retransmission rate for this characteristic shall be less than 10% of the available packets. The concern of electromagnetic susceptibility is a direct result of the environment in which the bridges must operate. In the typical factory environment, large motors, induction heating devices, fluorescent lighting fixtures, etc. generate electromagnetic interference which can adversely affect the performance of electronic components. This requirement requires the bridge's electronic components to be shielded, but not to the extent that all interference is stopped. The typical retransmission capability of CSMA/CD can compensate for some of the effects of the interference.

- g. The power line conducted susceptibility of the bridge precludes the exhibition of any malfunction, degradation of performance, or deviation from specified performance characteristics beyond the tolerances indicated in this equipment specification for a 400 Volt direct current (dc) impulse with a pulse width of ten microseconds. The retransmission rate under this condition shall be less than 10% of the available packets. This requirement is similar to the requirement stated in Subsection (f) above; the difference is that the power line supplying the bridge the 120 Volts ac required for operation must be shielded to arrest the propagation of interference which may adversely affect the performance of the bridge. The normal retransmission characteristics of the CSMA/CD protocol are an acceptable mechanism to compensate for a portion of this interference.
- h. Electromagnetic and radio frequency emissions shall not exceed those set forth in the Federal Communications Commission standard for Class B equipment (Part 15, Subpart J). This is the industry standard for most commercial electronic products.

- i. The operation of the bridge shall not be interrupted by voltage spikes on the broadband cable plant of 60 dBmV (or less) and shall not cause a loss of packets. Retransmission as a result of perturbations caused by this condition is acceptable. Other devices attached to the broadband system can generate voltage spikes when attached to the network prior to power-on. This requirement is that the bridge must be capable of handling these spikes.
- j. Cycling power on the bridge from off-to-on and vice versa shall not affect the signal levels present on the attached networks by more the 10% of their nominal operating levels. The only exception to this is the diagnostics invoked when the bridge is powered on. This is the exact inverse of the requirement stated in Subsection (i) above.
- k. The bridge shall operate on the broadband network under the given backbone trunk length and physical configuration. A bridge pair (in the case of the fiber optic backbone) shall operate over a 3000 foot trunk length of fiber optic cable given by the following parameters: 100 micron core/140 micron cladding, 50 micron core/125 micron cladding, and 62.5 micron core/125 micron

cladding, one transmit fiber and one receive fiber, a 0.20 numerical aperture, 500 MHz bandwidth, and a 5 dB/km nominal attenuation characteristic over all available wavelengths.

1. The Ethernet-to-broadband bridges shall meet the following frequency allocation requirements:
 1. Transmission on the broadband media shall be no less than 5 Mbps for every 6 MHz channel used
 2. Bridges shall operate on channel pairs 5 and S, 6 and T, and at least two other forward and reverse channel pairs
 3. No single bridge shall require more than 12 MHz in the forward and reverse directions. The IEEE has issued a draft standard which allows bridges to utilize three 6 MHz channels in each of the forward and reverse bands in order to support 10 Mbps. Several vendors support this draft standard, however, the requirement of the bridges supplied by this specification is that no more than two 6 MHz channels be utilized in each direction.
 4. Channel-to-channel bridges or some other configuration sufficient to allow the

Bridge Management Control Center to manage all bridges shall be provided. This requirements allows a management system to reside on a pair of 6 Mhz channels (one forward and one reverse) and manage bridges attached to other channel pairs; this support the implementation of more than one Ethernet channel on a broadband system.

- m. The Ethernet-to-broadband bridges shall function on the broadband cable plant within the RF specifications listed below:

- System Amplitude Response 150-400 MHz
forward
5-112 MHz
reverse
- Channel Flatness ± 2.5 dB
- Adjacent Channel Isolation 60 dB minimum
- Bridge Transmit Level 50 dBmV ± 2
dBmV
- Bridge Receive Level 3 dBmV ± 2 dBmV
- RF Impedance 75 ohms ± 2
ohms
- Frequency Stability $\pm 0.01\%$ minimum

- n. Destination filtering of all packets passing through the bridge from one attached network to another is required. The destination filtering shall be performed by means of a table of addresses residing in the bridge. The filtering method shall keep local data traffic restricted to the local network and shall allow remote data to be transmitted only to the local network having the destination address.
- o. The error rate, as defined by subtracting the number of packets seen at the receiving analyzer from the number of packets forwarded - with the results divided by the number forwarded, shall be 10^{-9} or less.
- p. The algorithms for detecting transmission errors shall be equal to or better than CRC-16 (Cyclic Redundancy Check-Modulo 16). Error correction shall be by retransmission of lost or damaged packets.
- q. The packet pick-up rate, as defined by the number of packets per unit of time actually placed on the broadband channel, fiber optic backbone, or T1 backbone, shall be within 1% of the number of packets forwarded through the bridge.
- r. In the "learning" mode, the bridge shall automatically construct the address table from the

existing network topology. The table of addresses shall be updated dynamically while in "learning" mode.

- s. The address table shall be constructed to allow the manual entry of prohibited destinations to be entered into the bridge ("management" mode) from the Bridge Management Control Center.
- t. The bridge shall operate at media speeds and shall provide the buffering necessary to transition from one network type to another while operating at the data rates as specified previously.
- u. The bridge shall be able to receive commands and upload network management data by communicating with the Bridge Management Control Center over the IEEE 802 networks. Some systems support out-of-band network management communications in lieu of utilizing the IEEE 802 networks they are managing. This system shall not use the out-of-band signaling technique.
- w. The bridge shall be equipped with a self-test algorithm which can be invoked either locally (at the bridge) or from the BMCC.
- x. The bridge shall accommodate loops (bridges operating in parallel) in the network topology by performing the following functions:
 - Automatically detect all loops

- In the case of different media, select the lowest speed link to deactivate
- In the case of the same media, select one link to deactivate
- Return to the spanning tree topology
- Remain cognizant of the deactivated link
- Automatically reactivate the link when a failure of another link or bridge causes isolation of a subtree in the network topology

As discussed previously, Ethernet networks are designed for a bus or tree topology and cannot tolerate loops. This makes load sharing between bridges difficult. A loop cannot be tolerated because packets would be regenerated indefinitely on the network by the next bridge in the loop. If a loop is created, as new packets are transmitted on the network, the traffic level will increase continuously until saturation of one or several elements of the loop. Loops must be quickly detected and corrected. The spanning tree algorithm defined by the IEEE 802 committee overcomes these problems by allowing a hot standby bridge to automatically take over in case of failure of an associated internetwork path. This is especially useful with remote bridges, where

the reliability of the interconnecting medium is far lower than that of the bridges.

- y. The overhead associated with all packets forwarded from the baseband Ethernet network onto the broadband, T1, and fiber optic media shall be less than or equal to the overhead associated with the packets on the Ethernet network. This requirement is established to insure against "flooding" the network with overhead associated with the transition between unlike network media types.

G. Bridge Management System

The bridge shall receive management commands and information required for operation from a Bridge Management Control Center (BMCC). Figure 12 illustrates the functional requirements of the BMCC. All bridge management commands and control information shall originate in the BMCC. The BMCC shall reside on a remotely attached International Business Machines (IBM)²⁶ compatible Personal Computer (PC) (Model AT or PS/2), a SUN Microsystems Model 3/160M, or a Digital Equipment Corporation (DEC) VAX²⁷-based computer or

²⁶ IBM is a registered trademark of International Business Machines, Inc.

²⁷ VAX is a registered trademark of Digital Equipment Corporation




FIGURE 12

BRIDGE MANAGEMENT SYSTEM ARCHITECTURE

workstation which shall be specified. Any additional hardware required to support the bridge management system and comply with the requirements of this section shall be provided, in addition to, all software necessary for full system operation including (but not limited to) address table construction, routing algorithms, system diagnostics, and reporting. The BMCC shall collect data from the bridge(s) through event logging and network manager requests. Event logging shall be used to obtain bridge status and characteristics and traffic/error counters automatically, while network manager requests shall be utilized to obtain specific information pertaining to system operation at any time, including status changes and traffic/error counters. The BMCC shall have the following attributes:

- a. The BMCC shall present the data either as a screen display or as a printed report.
- b. The BMCC shall have full control over which nodes are involved in the event logging process.
- c. The BMCC shall provide on-line graphical displays of selected error statistics, traffic statistics, and status information. The graphic displays shall include histograms, linear scales, and a network map.

- d. The BMCC shall, on command, disable the bridge, restrict the flow of packets through the bridge based on destination address, and control all other operational aspects of the bridge using software resident in the BMCC.
- e. The filtering/forwarding mechanism for all bridges shall be programmable from the BMCC. The BMCC software shall provide a mechanism for overriding the address learning function (disable "learning" mode) of the bridge so that filtering/forwarding tables can be directly viewed and manually updated by the BMCC ("management" mode). This requirement allows the direct intervention of the network manager in deciding how to restrict the flow of packets in the bridged network.
- f. In "management" mode, the bridge shall not "learn" of any new addresses on the network without the direct intervention of the BMCC. The addresses in the address table, created in "management" mode, shall not "age" and be discarded unless specified by the BMCC.
- g. The BMCC shall display bridge and data link counters, bridge status (i.e. power on/off, transmitting data, receiving data), and the bridge forwarding database.

- h. The BMCC shall modify all bridge parameters, including operational state, forwarding database, and spanning tree characteristics.
- i. The BMCC shall invoke the bridge self-test and remote reset functions.
- j. The BMCC shall provide for the association of alphanumeric names with specific bridge physical addresses for ease of use in management commands.
- k. The bridge shall track local and internetwork traffic for each attached network by recording the number of packets filtered and forwarded (in both directions), the number of packets discarded by the bridge, and the number of packets received in error.
- l. All collected statistics shall be expressed in terms that can be translated into packets per second, and the sampling interval shall be at most 15 minutes.
- m. The bridge shall collect all bridge-related traffic statistics and report such information to the BMCC. The traffic statistics reported shall include, but not be limited to, the following:
 - 1. Number of transmission attempts per unit of time
 - 2. Number of packet collisions per unit of time

3. Number of CRC errors per unit of time

n. The bridge shall generate alerts on the BMCC screen, printer, and as an audible tone for the following conditions:

- Failure of the bridge or link due to problems in the bridge, or in the case of remote bridge connections, because of link degradation or failure
- The bridge or link is overloaded and cannot filter and/or forward the packets being generated due to speed limitations of the link or buffers being full.
- A degradation of bridge or link performance which may still allow packets to be filtered/forwarded, but the bridge is no longer able to function at its full rated capacity.
- The bridge detects violations of the IEEE 802 protocols.
- The bridge is not able to filter or forward a packet because it cannot reconcile the MAC addresses of a received packet with the bridge's filtering tables.

- o. The bridge shall include trouble reports with each alert. Table 2 illustrates possible trouble situations with the minimum status information required for each condition.
- p. Each bridge shall provide alerts and trouble reports to the BMCC within 30 seconds of the occurrence of the alarm condition.

H. Reliability, Availability, and Maintainability (RAM) Requirements

This section describes a method of determining the percentage of time that the bridges and management system must be operational. Each system (consisting of the data link bridge(s), software, and management system) shall have an availability (denoted by A) of at least 0.999.

Availability shall be defined as follows:

$$A = (T^{op} - \sum_{i=1}^f \sum_{j=1}^n T_{ij}^{dn} \times P_j) / T^{op}$$

where

T^{op} = Operational time over which the availability calculation applies

TABLE 2
ALARM REPORT/STATUS INDICATION MATRIX

<u>Alarm Condition</u>	<u>Minimum Status Information Required</u>					
	<u>A</u>	<u>B</u>	<u>C</u>	<u>D</u>	<u>E</u>	<u>F</u>
Bridge Failure					X	X
Link Failure		X	X	X		X
Bridge Overload		X	X	X		
Link Overload	X	X	X			
Bridge Degredation		X	X	X		X
Link Degredation		X	X	X		X
Protocol Violation			X	X	X	
Unknown Address		X			X	
Power Failure					X	X

Key: Status Information Learned

A = Number of Transmission Attempts per Unit of Time

B = Number of Packet Collisions per Unit of Time

C = Number of CRC Errors per Unit of Time

D = Type of Protocol Error

E = Transmitting or Receiving Data

F = Power On/Off

T^{dn} = Downtime or time equipment is not available for service functioning in accordance with this specification. This includes any time required for preventive maintenance, startup, programming, configuration, and other activities not associated with data communication

P = Average portion of total system service lost for each failure:

For software management systems, $P = 1$

For bridges on fiber optic Ethernet, $P = 0.05$

For bridges on the broadband, $P = 0.1$

For bridges on coaxial cable Ethernet, $P = 0.1$

For bridges on T1 links, $P = 0.1$

This equation illustrates that availability can be increased by decreasing downtime (T^{dn}). Downtime can be decreased by providing easily accessible spare parts, easy access to equipment, and by minimizing the troubleshooting time required to identify the faulty part.

Meeting the necessary availability requires the following approach: (1) system design that provides high reliability and low mean-time-to-restore (MTTR); (2) maintenance planning that provides rapid troubleshooting techniques, sufficient preventive maintenance activity, and readily available spare parts; and (3) operation techniques that

provide needed user help and assistance and appropriate measurements and logs. Data link bridge mean-time-to-restore (MTTR) shall be less than two hours and shall consist of the following:

1. Time to determine which bridge has failed
2. Time to replace each 19-inch rack mountable unit (assuming trained and qualified technicians are performing the work)
3. If multiple pieces of equipment are used, the time to determine which equipment must be replaced (in meeting this requirement, it will be assumed that only one piece of equipment will fail due to random causes)
4. After replacement, time to set all controls, energize equipment, configure and/or download any software, and place the unit into operation

All bridge connections utilize connectors and appropriate cables such that repair can be made by replacing units. All power switches, reset switches, operational lights and/or indicators, and any other controls used during normal operation must be easily accessible from the front or rear of the unit(s). Controls or test jacks used for troubleshooting must be accessible from either the front or rear of the unit(s). These factors, along with efficient

maintenance when failures do occur, are the most important considerations in maintaining high availability.

I. QUALITY ASSURANCE

1. General Testing Requirements

In order to assure that the bridge will perform satisfactorily in service, a comprehensive test plan must be developed and then utilized to develop step-by-step test procedures required to ensure that the entire bridge and bridge management system meets all technical, operational, and performance requirements of this specification. Certified test results must be furnished prior to shipment of any part of the system.

2. Local Bridge Testing

Local data link bridges to be furnished must be tested to determine the maximum number of packets that it can pick up and forward (total packet throughput) to the destination network. Two Ethernet analyzers, one on the source network and one on the destination network, are used to generate and receive the packets. The test must be conducted in both directions. The packet size is varied from the minimum (64 bytes) to the maximum (1518 bytes) Ethernet packet sizes.

The following range of packets sizes are used as a minimum; 64, 256, 512, 1024, and 1518 bytes per packet.

The following parameters are measured, recorded, and plotted as a function of the packet size:

1. Supplied packet rate
2. Error Rate
3. Total Packet Throughput
4. Megabit/Second Throughput - total packet throughput multiplied by the bytes per packet multiplied by the bits per byte

At all prescribed packet sizes, the throughput rate must meet or exceed that specified.

3. Remote Bridge Testing

Remote bridges must also undergo testing to determine their performance over all specified backbone networks. The test configuration consists of a local Ethernet network (the source network), a source Ethernet analyzer, and the proposed remote bridge interfacing to backbone (broadband, fiber optic, or T1) networks. A second remote bridge connects the backbone network to a second local Ethernet network (the destination network), to which a destination

Ethernet analyzer is attached. These same tests as outlined previously are repeated under the following conditions for the remote bridges:

1. Both local area networks and bridges are located within 300 feet of the broadband headend equipment.
2. One of the local area networks is a distance of 3000 feet from the broadband headend equipment.
3. Both of the local area networks are at a distance of 3000 feet from the broadband headend equipment.
4. Both of the local area networks are interconnected using 3000 feet of fiber optic cable.

The parameters measured, recorded, and plotted are the same as for the local bridges; in addition, the packet pick-up rate (defined as the number of packets per unit of time actually placed on the broadband channel) and the packet transmitted on the broadband, fiber optic backbone, or T1 backbone is examined. At all prescribed packet sizes, the throughput rate and the packet pick-up rate must meet or exceed that specified and the packet overhead must meet that specified.

4. Packet Filtering Testing

Figure 13 illustrates the test to reveal the ability of the bridge to examine local network traffic, determine which packets are to be forwarded on to destination network, and then forward the appropriate packets on to the bridged LAN. The test is performed for a variable number of packet sizes and for a variable mix of local/remote packets. The range of packet sizes used is 40, 256, 512, 1024, and 1550 bytes per packet.

The mixture of remote/local traffic must include, as a minimum: 0%, 20%, 40%, and 80% (100% would correspond to the tests previously defined). The acceptance criteria for this test is as follows:

1. When Node A is transmitting packets destined for Node B, Node B must receive all of these packets. Node C must not receive any of these packets.
2. When Node A is transmitting packets destined for Node C, Nodes B and C must receive all of these packets.
3. When Node B is transmitting packets destined for Node C, Nodes A and C must receive all of these packets.




FIGURE 13
BRIDGE PACKET FILTERING TEST CONFIGURATION

Any baseline testing required before actual bridge testing is started must be performed. A test showing the capability of the source analyzer to deliver the required packet rates to the destination analyzer as a function of packet size must be conducted with the results recorded and plotted.

5. Miscellaneous Testing

This testing must include a demonstration of the features required by the specification. Specification features that must be demonstrated include, but are not limited to, the following:

- Loading and operation of bridge operating and management software
- Downloading software to the bridge
- Collection and display of network statistics

Physical attributes which must be demonstrated include:

- Status indicators for power on/off, transmit/receive data, and collision detect
- Self-test and reset
- Power on
- Packets received/forwarded

6. RF Testing

RF testing must be conducted in accordance to the procedures set forth in the National Cable Television Association (NCTA) Standard 008-0477 and Federal Communications Commission (FCC) Rules and Regulations Part 76 and then meet the specifications stated previously. The radiated and power line conducted susceptibility testing must be conducted in accordance to the procedures set forth in MIL-STD-462 for Radiated Susceptibility, RS03 and Power Line Conducted Susceptibility, CS06 and then meet the specifications outlined in this paper.

7. Application Layer Testing

Bridge operation will be insured using the application layer protocols of DECnet. Figure 14 depicts the test configuration. The tests will include both task-to-task communications and node-to-node file transfers using DECnet software. Testing will be conducted bi-directionally between computers and will test both local and remote bridges. Testing will be done with data block sizes ranging from one (1) byte to 4096 bytes formatted in accordance with the IEEE 802.3 specification. The test software is structured to demonstrate a node-to-node file transfer capability equal to 50 kilobytes per second. These tests




FIGURE 14
APPLICATION LAYER TEST CONFIGURATION

insure compatibility with the DECnet protocol suite and compliance with the IEEE 802.3 specification for packet construction. The test packets must be passed within the error rate specified and at the file transfer rate specified previously.

8. System RAM Testing

System RAM requirements must be tested by operating five data link bridges and the software management system of each system in the configuration shown in Figure 15 based on the following parameters and the availability formula previously defined:

$$T^{op} = 30 \text{ days}$$

$$T^{dn} = T^r + T^{pm}$$

where

T^r = Time to repair failures

T^{pm} = Time to perform any vendor recommended preventive maintenance

Failure = Data link bridge or management system not meeting any requirement of this specification

The testing shall be conducted as follows:




FIGURE 15
BRIDGE RAM TEST CONFIGURATION

1. Deliver all operating equipment plus one spare data link bridge of each type.
2. Install all equipment and test to insure proper operation.
3. Log time T^{op} will begin.
4. After five days of operation, randomly deenergize one data link bridge. Insure proper management system logging and alarm and record time to restore unit to full operation.
5. For each equipment failure, record the time required to return the equipment to full operation.
6. After 30 days calculate and record total T^{dn} and A (the action of step 4 shall not be included in these calculations).

The acceptance of the system shall be based, in part, on the test results defined by the following criteria:

1. $A \geq 0.999$
2. Step 4 $T^{dn} < 2$ hours
3. $T^{dn} < 2$ hours for any data link bridge failure
4. Proper software management system recording of all bridge failures.

J. EVALUATION CRITERIA

The evaluation of proposals allows for the selection of an established, financially stable, and technically reliable firm to fulfill the requirements of the specification. The proposals were evaluated in the following specific areas: price, technical excellence, and corporate experience. In the pricing arena, the lowest total offer of those firms which are in the competitive range receive the maximum price score with higher offers receiving proportionately lower scores. The technical evaluation of the bridges is based on meeting all of the criteria of the specification prior to any further evaluation. In defining the acceptability of the technical portion of the proposal, a descriptive narrative is required to define how compliance with each requirement is achieved. Further, while not specific requirements, the following items provide a basis from which to select a technically superior product:

1. Type of Filtering - Destination-field filtering is defined as the minimum technical requirement by the specification. However, customizable filters are highly desirable. The filters should be able to filter packets based on source and destination addresses, packet type, and any other field defined in an 802.3 packet.

2. Number of Channels Required - This applies to the Ethernet-to-Broadband bridges only. The highly desirable feature is a bridge which uses a pair of 6 MHz channels to support 5 Mbps of Ethernet traffic as compared to a pair of 12 MHz channels to support 10 Mbps of Ethernet traffic. This allows for the segmentation of individual Ethernet networks on the broadband cable plant.
3. Bridge Modularity - The ability to change media by simply changing an interface card in a modular enclosure is highly desirable as compared to adding individually-housed units each time a media change is made. This is utilized to reduce the number of chassis involved in the bridge configuration, and thereby improve the RAM capabilities of the bridge.

The final portion of the evaluation examines the corporate experience of prospective vendors. The proposals must contain brief histories of the companies and the number of years they had been supplying data link bridges. In addition, the proposal must include the number of years that

the bridges and software have been in production. This is to attempt to reduce the likelihood of "development" or "beta" equipment being supplied on the contract.

IV. SUMMARY AND CONCLUSIONS

A. Evaluation and Product Selection

Five proposals were received, one each from the following vendors: Ungermann-Bass, Bridge Communications, Digital Equipment Corporation, Retix, and CEI. Ungermann-Bass was selected to provide data link bridges based on the successful completion of the evaluation criteria established previously.

Testing must now be performed on these bridges to assure compliance with the specification. The summary of the results of these tests are found in Appendix 1. Much of the testing had never been done before by Ungermann-Bass, so the test procedures and data gathered provided some valuable insight into the performance aspects of their products in addition to assuring conformance to the specification. As expected, all testing demonstrated either conformance to the requirements of the specification or surpassed those requirements. Overall, the testing proved the products supplied by Ungermann-Bass were superior to the requirements outlined in the specification.

B. Long Range Target Architecture

Influence is being exerted on vendors by users to move to "wiring closet" LAN topologies and providing multiport bridges (similar to multiport repeaters). Most of these will be able to change "personalities" by plugging interface cards for various LAN and remote links into a backplane. Users will thus be able to migrate from a star-coupled Ethernet backbone to a Fiber Distributed Data Interface (FDDI) backbone with minimal disruption and cost, as both fiber and a bridge will be used. In addition to better reflecting wiring-closet topology, multiport bridges will amortize the cost of connecting multiple networks to high-priced, high-performance FDDI backbones. Remote multiport bridges will act as quasi-routers, forwarding traffic to multiple destinations attached to different ports.

Migration from Ethernet to FDDI is consistent with the implementation of fiber optic bridges and the installation of plant-wide fiber optic networks that many companies are in the process of performing.²⁸ FDDI is not a technological afterthought like Ethernet. Supporting up to 500 stations on one network and able to span more than 100 kilometers,

²⁸ Long Range Planning Study. Network Strategies, Inc., Fairfax, Va. January 1988.

FDDI's future as a standard seems bright. Developed by the American National Standards Institute (ANSI) X3T9.5 Accredited Standards Committee, FDDI has been endorsed by many of the computing community's major players.²⁹ Completion of the FDDI standard, expected late 1989, coincides with the demand by network system implementors for both improved performance and multivendor interoperability. Systems complying with the FDDI specification will meet both requirements. The data rate supported by FDDI is 100 Mbps and is ten times faster than Ethernet's 10 Mbps. It is a fiber optic, ring-based, token-passing network that specifies a protocol set up to the station management level of the data link layer. Most implementors will use 802.2 LLC over the FDDI media-access control services. In contrast to Ethernet's bus design, FDDI's ring topology consists of point-to-point connections to all stations on the ring and link only from adjacent station to adjacent station. This point-to-point method allows flexibility, from fiber size to cable type and connector style, in network wiring. FDDI requires no special repeaters, amplifiers, or other signal conditioning equipment since it employs commonly available fibers and electro-optical components. It is based on a dual backbone of primary and secondary rings that each transmit counter rotationally.

²⁹ Moore, Michael V. "FDDI: A Federal Government LAN Solution". Telecommunications, p. 35-44. September 1989.

the primary ring handles data flow, while the secondary ring provides backup or additional data flow. The dual ring design allows FDDI to continue to operate if a station on the ring fails or a point to point fiber optic segment becomes disabled. FDDI networks employ a token passing protocol that allows network users to share the media via a sequential access scheme. The token grants permission for an individual station to send data on the network. A message, either data or the token itself, passes through the FDDI ring by entering every station on the inbound fiber and exiting via the outbound fiber. When one network station sends information to another, it creates a data packet that includes the address of the destination station. That data packet cycles through the ring sequentially. Each station examines the specified address in the data packet and accepts the data only if the packet is intended for it. A concentrator connects the two counter rotational rings to single rings that are functional but subordinate elements of the primary ring. Designed for network segments that do not require redundancy, the single rings employ single fibers, rather than the two fibers required by the dual counter rotating ring.

Several factors have combined to drive the demand for FDDI-based products. In particular, the number of nodes per network has risen sharply over the past few years. Networks

featuring hundreds of nodes are common, and the typical throughput requirements of the nodes has risen. User performance requirements are stressing Ethernet technologies, while the capabilities of hosts, the growth in the number of networks and workgroups and the reality of multivendor environments are similarly presenting challenges for large network implementors. FDDI-based systems will significantly improve the performance of future networks. The development efforts are focusing on FDDI-based backbone networks matched with Ethernets (specifically twisted pair Ethernet networks based on IEEE Standard 10BASET) as one of the economical and effective ways to implement FDDI technology. Bridges containing FDDI network interfaces on one side and IEEE 10BASET network interfaces on the other would promote cost effective segmentation of networks. Traffic on a LAN would be isolated on low cost twisted pair system, while internetwork traffic would be moved via the fiber based backbone. Host-to-host connection will also be among the key implementations of FDDI. With the growth in potential throughput associated with CPUs and buses, the development of technologies such as FDDI will be essential to accelerating data delivery.

The architecture of these developing fiber optic backbone networks will deploy strategically placed "hubs" around a plant and interconnect these with multifiber optical cables.

These hubs will initially consist of Ethernet star concentrators with data link bridges providing the necessary traffic isolation between the backbone network and the local attached Ethernets. However, as the FDDI to Ethernet bridges become commercially available and competition forces prices down to a reasonably level (\$10,000 per node versus the current \$25,000 - \$30,000 per node), this Ethernet-based backbone network will be converted to FDDI. This plan is consistent with the development of the technology.

The procurement of data link bridges under this specification sets the stage for future migration to higher speed backbone networks based on international standards. This migration takes into account all of the functional requirements stated for the present procurement of bridges (i.e. single management system, etc.).

BIBLIOGRAPHY

BIBLIOGRAPHY

Sylor, Mark W. "Managing Phase V DECnet Networks: the Entity Model". IEEE Network, p. 30-36, vol. 2, issue 2, p. 30-36, March 1988.

Schnaidt, Patricia. "Bridge Deluge". LAN Magazine, p. 70-78. February 1989.

Leong, John. "Guide to Ethernet Installation and Interconnection". Data Communications, p. 267-276. September 1985.

Green, Paul E. "An Introduction to Network Architectures and Protocols". IEEE Transactions on Communications, vol. 28, issue 4, p. 413-424, April 1980.

Abraham, Menachem E. "Running Ethernet Modems Over Broadband Cable". Data Communications, p. 199-212. May 1986.

Zimmermann, Hubert. "OSI Reference Model-The ISO Model of Architecture for Open Systems Interconnection". IEEE Transactions on Communications, vol. 28, issue 4, p. 425-431, April 1980.

Bertine, H. V. "Physical Level Protocols". IEEE Transactions on Communications, vol. 28, issue 4, p. 433-444, April 1980.

Daigle, J. N. "Communications for Manufacturing: An Overview". IEEE Network, vol. 2, issue 3, p. 6-13, May 1988.

Internet Data Strategy. Network Strategies, Inc., Fairfax, Va. June 1988.

Hunter, John. "Bridges and Gateways: Supplying the Missing Link". Network World, p. 33-38. July 1988.

Data Link Bridge Principles of Operation. Ungermann-Bass, Inc., Santa Clara, Ca. February 1987.

Long Range Planning Study. Network Strategies, Inc., Fairfax, Va. January 1988.

Brusil, Paul J. "Toward a Unified Theory of Managing Large Networks". IEEE Spectrum, p. 39-42. April 1989.

Schnaidt, Patricia. "Span the LAN". LAN Magazine, p. 27-30. February 1988.

ANSI/IEEE Standard for Local Area Networks 802.3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications. IEEE Computer Society, 4th ed. New York: IEEE Press, c 1986.

Draft IEEE Standard 802.1, Part D: MAC Bridges. IEEE Computer Society, revision B, May 1987.

Introduction to Local Area Networks. Digital Equipment Corp, c 1982.

Digital's Networks: An Architecture with a Future. Digital Equipment Corp., c 1982.

Interconnectivity: Open Enterprise Networking. Ungermann-Bass, Inc., c. 1989.

Broadband Network Technology. Sytek, Inc., 2nd ed. New Jersey: Prentice-Hall, c 1986.

McQuillan, J. M. "Routers as Building Blocks for Robust Internetworks". Data Communications, p. 28-33. September 1989.

Boule, Richard. "Inside Routers: A Technology Guide for Network Builders". Data Communications, p. 53-66. September 1989.

Hawe, William R. "The Extended LAN Architecture and LANBridge 100". Digital Technical Journal, p. 54-72. September 1986.

Tobagi, F. "Performance Analysis of CSMA/CD". Computer Networks, p. 245-259. October/November 1980.

Moore, Michael V. "FDDI: A Federal Government LAN Solution". Telecommunications, p. 35-44. September 1989.

DECnet Digital Network Architecture Phase IV. Digital Equipment Corp., c. May 1982

ANSI/IEEE Standard for Local Area Networks 802.2: Logical Link Control (LLC). IEEE Computer Society, 4th ed. New York: IEEE Press, c 1986.

APPENDIX

1. Baseline Testing

Ethernet-to-Ethernet (Baseband)

To initiate the testing procedure, two Ethernet network analyzers will be set up on a single Ethernet coaxial cable segment. One analyzer is programmed to transmit packets to the network and the other analyzer is programmed to receive the packets. Packet sizes will be 64, 256, 512, 1024, and 1518 bytes. The following table contains the results:

Ethernet-to-Ethernet Baseline Testing

<u>Throughput</u>	<u>Network Loading</u>	<u>Packets per Second</u>	<u>Bytes</u>
2.14	26 %	4463	64
6.15	65 %	3000	256
8.73	90 %	2132	512
9.20	94 %	1124	1024
9.46	96 %	780	1518

Ethernet-to-Ethernet (Broadband)

To initiate the testing procedure, two Ethernet network analyzers will be set up on a single broadband coaxial cable segment. One analyzer is programmed to transmit packets to the network and the other analyzer is programmed to receive the packets. Packet sizes will be 64, 256, 512, 1024, and 1518 bytes. The only variation between this test and the previous baseband test is that the maximum traffic load that can be generated on the broadband network is 5 Mbps. Consequently, maximum packet rates will be 50% of their baseband maximums for any loadings limited by media capacity. The following table contains the results:

Ethernet-to-Broadband Baseline Testing

<u>Throughput</u>	<u>Network Loading</u>	<u>Packets per Second</u>	<u>Bytes</u>
2.12	26 %	4425	64
4.33	81 %	2114	256
4.61	90 %	1125	512
4.77	94 %	582	1024
4.80	96 %	398	1518

Ethernet-to-Ethernet (Fiber Optic)

To initiate the testing procedure, two Ethernet network analyzers will be set up on a single Ethernet fiber optic cable segment. One analyzer is programmed to transmit packets to the network and the other analyzer is programmed to receive the packets. Packet sizes will be 64, 256, 512, 1024, and 1518 bytes. The following table contains the results:

Ethernet-to-Fiber Optic Baseline Testing

<u>Throughput</u>	<u>Network Loading</u>	<u>Packets per Second</u>	<u>Bytes</u>
2.14	26 %	4461	64
6.16	65 %	3006	256
8.73	90 %	2131	512
9.21	94 %	1124	1024
9.46	96 %	781	1518

Ethernet-to-Ethernet (T1)

To initiate the testing procedure, two Ethernet network analyzers will be set up on a single T1 segment. One analyzer is programmed to transmit packets to the network and the other analyzer is programmed to receive the packets. Packet sizes will be 64, 256, 512, 1024, and 1518 bytes. The following table contains the results:

Ethernet-to-T1 Baseline Testing

<u>Throughput</u>	<u>Network Loading</u>	<u>Packets per Second</u>	<u>Bytes</u>
1.41	68 %	2100	64
1.45	84 %	650	256
1.52	90 %	355	512
1.40	85 %	168	1024
1.40	85 %	115	1518

2. Performance Throughput Testing

Ethernet-to-Ethernet (Baseband)

To initiate the testing procedure, one Ethernet network analyzer will be set up on an Ethernet coaxial cable segment; a second Ethernet network analyzer will be set up on another Ethernet coaxial cable segment, and the two Ethernet segments will be connected via a data link bridge. One analyzer is programmed to transmit packets to the network and the other analyzer is programmed to receive the packets. Packet sizes will be 64, 256, 512, 1024, and 1518 bytes. The following table contains the results:

Ethernet-to-Ethernet Throughput Testing

<u>Throughput</u>	<u>Network Loading</u>	<u>Packets per Second</u>	<u>Bytes</u>
1.45	19 %	3023	64
3.62	38 %	1770	256
4.50	46 %	1105	512
5.20	53 %	638	1024
5.40	55 %	448	1518

Ethernet-to-Ethernet (Broadband)

To initiate the testing procedure, one Ethernet network analyzer will be set up on an Ethernet coaxial cable segment; a second Ethernet network analyzer will be set up at the end of 3000 feet of a broadband coaxial cable segment, and the the Ethernet segment will be connected to the broadband segment via a data link bridge. The analyzer on the Ethernet network is programmed to transmit packets to the network and the analyzer on the broadband network is programmed to receive the packets. Packet sizes will be 64, 256, 512, 1024, and 1518 bytes. The following table contains the results:

Ethernet-to-Broadband Throughput Testing

<u>Throughput</u>	<u>Network Loading</u>	<u>Packets per Second</u>	<u>Bytes</u>
1.45	18 %	3024	64
3.57	38 %	1770	256
4.16	43 %	1020	512
4.57	47 %	560	1024
4.61	47 %	383	1518

Ethernet-to-Ethernet (Broadband)

To initiate the testing procedure, one Ethernet network analyzer will be set up on an Ethernet coaxial cable segment; a second Ethernet network analyzer will be set up at the end of 3000 feet of a broadband coaxial cable segment, and the the Ethernet segment will be connected to the broadband segment via a data link bridge. The analyzer on the broadband network is programmed to transmit packets to the network and the analyzer on the Ethernet network is programmed to receive the packets. Packet sizes will be 64, 256, 512, 1024, and 1518 bytes. The following table contains the results:

Ethernet-to-Broadband Throughput Testing

<u>Throughput</u>	<u>Network Loading</u>	<u>Packets per Second</u>	<u>Bytes</u>
1.45	18 %	3024	64
3.48	37 %	1700	256
4.17	43 %	1017	512
4.57	47 %	560	1024
4.71	48 %	390	1518

Ethernet-to-Ethernet (Fiber Optic)

To initiate the testing procedure, one Ethernet network analyzer will be set up on an Ethernet coaxial cable segment; a second Ethernet network analyzer will be set up at on a second Ethernet coaxial cable segment which is at the end of 3000 feet of fiber optic cable, and each Ethernet segment will be connected to the fiber optic cable segment via a data link bridge. The analyzer on the first Ethernet segment is programmed to transmit packets to the network and the analyzer on the second Ethernet network (at the end of 3000 feet of fiber optic cable) is programmed to receive the packets. Packet sizes will be 64, 256, 512, 1024, and 1518 bytes. The following table contains the results:

Ethernet-to-Fiber Optic Throughput Testing

<u>Throughput</u>	<u>Network Loading</u>	<u>Packets per Second</u>	<u>Bytes</u>
1.45	19 %	3024	64
3.48	37 %	1700	256
4.17	43 %	1017	512
4.60	47 %	560	1024
4.74	48 %	391	1518

Ethernet-to-Ethernet (T1)

To initiate the testing procedure, two Ethernet network analyzers will be set up on a single T1 segment. One analyzer is programmed to transmit packets to the network and the other analyzer is programmed to receive the packets. Packet sizes will be 64, 256, 512, 1024, and 1518 bytes. The following table contains the results:

Ethernet-to-T1 Throughput Testing

<u>Throughput</u>	<u>Network Loading</u>	<u>Packets per Second</u>	<u>Bytes</u>
1.41	68 %	2100	64
1.55	89 %	691	256
1.52	90 %	355	512
1.51	90 %	180	1024
1.50	92 %	122	1518

3. Filter Rate Testing

To initiate the testing procedure, a test configuration identical to Figure 14 is established. One analyzer is programmed to transmit packets to the network and the other analyzer is programmed to receive the packets. Packet sizes will be 64, 256, 512, 1024, and 1518 bytes. Initially, none of the packets will be addressed to Node C. Data will be gathered by Nodes A and B and each bridge's buffers examined to determine the filtering capacity of the bridge. If the bridge experiences a buffer overrun, the packet rate will be reduced until a sustainable rate is found for the packet size, rate, and forwarding combination under test. Upon completion of 0% testing, the percentage of packets addressed to Node C will be increased to 20% and the test will be repeated. The percentage of packets addressed to Node C will then be increased to 40%, 60%, and 80%, with the test repeated at each level. The following table contains the results:

Bridge Filter Testing

<u>Throughput</u>	<u>Network Loading</u>	<u>Packets per Second</u>	<u>Bytes</u>
2.15	43 %	7385	64
6.42	81 %	3700	256
8.70	90 %	2100	512
9.21	94 %	1124	1024
9.57	96 %	782	1518

4. RF Testing

Power Line Conducted Susceptibility (Method CS06)

A $\pm$ 400 volt, 10 microsecond pulse is applied to each of the following input lines:

117 Vac, 60 Hz, High Line

117 Vac, 60 Hz, Return Line

During application of the pulse, an oscilloscope is placed across the input terminals to ensure that the waveform required is adhered to. Lines tested are exposed to the following test conditions:

- Nonsynchronized repetitive transient of 6 to 10 pulses/second
- Single pulse
- Synchronized with test bridge ac source and phase positioned at each 90 degree point in the ac waveform. In addition, its position was varied from 0 to 360 degrees on the ac waveform. The transient pulse was series injected into each ac line. During the application of the spike(s), the

bridge was monitored for performance characteristics previously described.

The results of this test were no change in indication of the bridges, malfunction or degradation in bridge operations were observed.

Radiated Susceptibility (Method RS03)

A data link bridge is subjected to a radio-frequency (RF) field as follows:

<u>Level V/m</u>	<u>Frequency Spectrum</u>	<u>Modulation</u>
10	14,000 Hz to 1,000 MHz	CW

The following antennas are utilized to produce the necessary fields:

<u>Spectrum</u>	<u>Antenna</u>
14,000 Hz to 200 MHz	Parallel Plate
200 MHz to 1,000 MHz	Double Ridge Guide

Power is applied to the transmitting antenna while it is directed at the data link bridge. The transmitting and receiving antennas are placed one (1) meter apart with the data link bridge between. During the scanning process, the

bridge is monitored for performance characteristics defined previously.

The results of this test were no change in indication of the bridges, malfunction or degradation in bridge operations were observed.

VITA

David Thomas Bell was born in Oak Ridge, Tennessee on April 28, 1959. He attended Oak Ridge High School and was graduated in June 1977. The following September he entered The University of Tennessee, and in June 1982 he received a Bachelor of Science degree in Electrical Engineering, graduating with Honors. After graduation, he was employed by the Carolina Eastman Company, a division of Eastman Kodak, where he worked as a design engineer.

In March 1984, he left the Carolina Eastman Company to begin employment with Martin Marietta Energy Systems, Inc. in Oak Ridge. At this time, he also enrolled in The Graduate School of The University of Tennessee to begin the pursuit of the Master of Science degree in Electrical Engineering.

At present, Mr. Bell is the principal engineer of network and communication systems for Martin Marietta Energy Systems, Inc., a position he has held since 1988. He is a member of the Communications and Computer Societies of the IEEE.

Mr. Bell is married to the former Mary Ellen Frost of Oak Ridge, Tennessee. They currently reside in Roane County, Tennessee.