

DOCUMENT RESUME

ED 167 172

IB 007 009

AUTHOR Cotton, Ira W., Ed.
TITLE Computer Science & Technology: Local Area Networking.
INSTITUTION National Bureau of Standards (DOC), Washington, D.C. Inst. for Computer Sciences and Technology.
REPORT NO NBS-SP-500-31
PUB DATE Apr 78
NOTE 82p.; Report of a Workshop (Gaithersburg, Maryland, August 22-23, 1977) ; Parts may not reproduce clearly due to print quality
AVAILABLE FROM Superintendent of Documents, U.S. Government Printing Office, Washington, D.C. 20402 (SD Stock No. SN003-003-01918-6; \$2.40)
EDRS PRICE MF-\$0.83 HC-\$4.67 Plus Postage.
DESCRIPTORS *Electronic Data Processing; Information Processing;
IDENTIFIERS *Networks; *Telecommunication; Time Sharing
 *Computer Networks; Protocols

ABSTRACT

The different technologies applicable to computer networks serving limited geographic areas, e.g., a single campus, factory, or office complex, are discussed in a number of short presentations made by active researchers and implementers in this field. Intensive discussion by participants in working sessions is reported for six topics: subnet architecture, protocols for local area networks, local network applications, network architecture, network operating systems, and analysis and performance evaluation. A list of attendees and a bibliography on local area computer networks is included. (Author/RAO)

 * Reproductions supplied by EDRS are the best that can be made *
 * from the original document. *

ED167172

COMPUTER SCIENCE & TECHNOLOGY:

Local Area Networking

U.S. DEPARTMENT OF HEALTH,
EDUCATION & WELFARE
NATIONAL INSTITUTE OF
EDUCATION

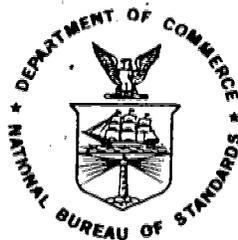
THIS DOCUMENT HAS BEEN REPRODUCED EXACTLY AS RECEIVED FROM THE PERSON OR ORGANIZATION ORIGINATING IT. POINTS OF VIEW OR OPINIONS STATED DO NOT NECESSARILY REPRESENT OFFICIAL NATIONAL INSTITUTE OF EDUCATION POSITION OR POLICY.

Report of Workshop Held at the
National Bureau of Standards,
Gaithersburg, Maryland

August 22-23, 1977

Ira W. Cotton, Chairman & Editor

Institute for Computer Sciences and Technology
National Bureau of Standards
Washington, D.C. 20234



U.S. DEPARTMENT OF COMMERCE, Juanita M. Kreps, Secretary

Dr. Sidney Harman, Under Secretary

Jordan J. Baruch, Assistant Secretary for Science and Technology

NATIONAL BUREAU OF STANDARDS, Ernest Ambler, Director

Issued April 1978

2

JK001007

Reports on Computer Science and Technology

The National Bureau of Standards has a special responsibility within the Federal Government for computer science and technology activities. The programs of the NBS Institute for Computer Sciences and Technology are designed to provide ADP standards, guidelines, and technical advisory services to improve the effectiveness of computer utilization in the Federal sector, and to perform appropriate research and development efforts as foundation for such activities and programs. This publication series will report these NBS efforts to the Federal computer community as well as to interested specialists in the academic and private sectors. Those wishing to receive notices of publications in this series should complete and return the form at the end of this publication.

National Bureau of Standards Special Publication 500-31

Nat. Bur. Stand. (U.S.) Spec. Publ. 500-31, 82 pages (Apr. 1978)

CODEN: XNBSAV

Library of Congress Cataloging in Publication Data

Main entry under title:

Local area networking.

(Computer science & technology) (NBS special publication 500-31)

Supt. of Docs. no.: C13.10:500-31

1. Computer networks—Congresses. / 2. Data transmission systems—
Congresses. I. Cotton, Ira W. II. United States. National Bureau of
Standards. III. Series. IV. Series: United States. National Bureau of
Standards. Special publication ; 500-31.

QC100.U57 no. 500-31 [TK5105.5] 602'.1s [001.644'04] 78-606029

U.S. GOVERNMENT PRINTING OFFICE
WASHINGTON: 1978

For sale by the Superintendent of Documents, U.S. Government Printing Office, Washington, D.C. 20402
Stock No. 003-003-01918-6 Price \$2.40

(Add 25 percent additional for other than U.S. mailing).

U.S. DEPT. OF COMM. BIBLIOGRAPHIC DATA SHEET		1. PUBLICATION OR REPORT NO. NBS SP 500-31	2. Gov't Accession No.	3. Recipient's Accession No.
4. TITLE AND SUBTITLE COMPUTER SCIENCE & TECHNOLOGY: Local Area Networking			5. Publication Date April 1978	
Report of a Workshop Held at the National Bureau of Standards, Gaithersburg, Maryland, August 22-23, 1977			6. Performing Organization Code	
7. AUTHOR(S) Ira W. Cotton, Chairman & Editor			8. Performing Organization Report No.	
9. PERFORMING ORGANIZATION NAME AND ADDRESS NATIONAL BUREAU OF STANDARDS DEPARTMENT OF COMMERCE WASHINGTON, D.C. 20234			10. Project/Task/Work Unit No.	
12. Sponsoring Organization Name and Complete Address (Street, City, State, ZIP) Same as item 9			11. Contract/Grant No.	
15. SUPPLEMENTARY NOTES Library of Congress Catalog Card Number: 78-606029			13. Type of Report & Period Covered	
16. ABSTRACT (A 200-word or less factual summary of most significant information. If document includes a significant bibliography or literature survey, mention it here.) This is the report of a workshop convened at the National Bureau of Standards on August 22-23, 1977, to discuss the different technologies applicable to computer networks serving a limited geographic area, such as a single campus, factory or office complex. A number of short presentations were made by active researchers and implementers in this area, afterwards the group broke up into a number of working sessions for intensive discussion of specific topics. A recorder at each session prepared a session report with the session chairman. The sessions were as follows: <ol style="list-style-type: none"> 1. Subnet architecture 2. Protocols for local area networks 3. Local network applications 4. Network architecture 5. Network operating systems 6. Analysis and performance evaluation A list of attendees and bibliography on local area computer networks is included in the report.			14. Sponsoring Agency Code	
17. KEY WORDS (six to twelve entries, alphabetical order, capitalize only the first letter of the first key word unless a proper name, separated by semicolons) Computer communications; computer networks; data communications; operating systems; performance evaluation; protocols			19. SECURITY CLASS (THIS REPORT) UNCLASSIFIED	
18. AVAILABILITY <input checked="" type="checkbox"/> Unlimited <input type="checkbox"/> For Official Distribution - Do Not Release to NTIS <input checked="" type="checkbox"/> Order From Sup. of Doc., U.S. Government Printing Office Washington, D.C. 20402, SD Stock No. SN003-003 - <input type="checkbox"/> Order From National Technical Information Service (NTIS) Springfield, Virginia 22151			20. SECURITY CLASS (THIS PAGE) UNCLASSIFIED	
			21. NO. OF PAGES 82 22. Price \$9.40	

ABSTRACT

This is the report of a workshop convened at the National Bureau of Standards on August 22-23, 1977, to discuss the different technologies applicable to computer networks serving a limited geographic area, such as a single campus, factory or office complex. A number of short presentations were made by active researchers and implementers in this area, afterwards the group broke up into a number of working sessions for intensive discussion of specific topics. A recorder at each session prepared a session report with the session chairman. The sessions were as follows:

1. Subnet architecture
2. Protocols for local area networks
3. Local network applications
4. Network architecture
5. Network operating systems.
6. Analysis and performance evaluation

A list of attendees and bibliography on local area computer networks is included in the report.

Key words: Computer communications; computer networks; data communications; operating systems; performance evaluation; protocols.

ACKNOWLEDGEMENTS

Special appreciation is due to the session chairmen and recorders, without whose assistance the Workshop could not have been held nor this Workshop Report prepared. The session chairmen were David Mills, Stuart Wecker, Philip Stein, Richard Sherman, Stephen Kimbleton and Ashok Agrawala. The recorders were Paul Meissner, Robert Rosenthal, Gary Donnelly, Robert Carpenter, James Hanks and William Franta.

Note: reference to any commercial products in this report is for the purpose of identification only and does not imply endorsement by NBS.

CONTENTS

1. INTRODUCTION	1
2. ABSTRACTS AND SHORT PAPERS	
A Comparative Evaluation of the Performance of Alternative Communications Technologies - A. K. Agrawala	5
Real-time Network for the Control of a Very Large Machine - J. Altaber	5
A Local Network for the National Bureau of Standards - R. J. Carpenter & R. Rosenthal	7
Application of Hyperchannel - G. Christensen	10
Mitrenet -- Introduction and Overview - J. P. Hanks	10
Data Ring at Computer Laboratory, University of Cambridge - A. Hopper	11
Local Mission-Oriented Network - R. L. Larsen	17
Interconnection of Local Networks Using Satellite Broadcast Technology - D. L. Mills	17
Ethernet: Distributed Packet Switching for Local Computer Networks - R. M. Metcalfe & D. R. Boggs	18
The ARPA Local Network Interface - P. V. Mockapetris, M. R. Lyle & D. J. Farber	18
Computer Cells -- High Performance Multi Computing - D. L. Nelson	19
The MIT Laboratory for Computer Science Network - K. T. Pogran & D. P. Reed	20
Current Summary of Ford Activities in Local Networking - R. H. Sherman, M. Gable & G. McClure	22
Local Area Networks at Queen Mary College - A. R. West	23
DECNET: Issues Related to Local Networking - S. Wecker	26

3. WORKSHOP REPORTS

Subnet Architecture 32
Chairman: D. Mills
Recorder: P. Meissner

Protocols for Local Area Networks 40
Chairman: S. Wecker
Recorder: R. Rosenthal

Local Network Applications 44
Chairman: P. Stein
Recorder: G. Donnelly

Network Architecture 47
Chairman: R. Sherman
Recorder: R. Carpenter

Network Operating Systems 55
Chairman: S. R. Kimbleton
Recorder: J. P. Hanks

Analysis and Performance Evaluation 59
Chairman: A. K. Agrawala
Recorder: W. R. Franta

4. BIBLIOGRAPHY ON LOCAL AREA COMPUTER NETWORKING 65

5. WORKSHOP ATTENDEES 70

INTRODUCTION

The NBS Workshop on Local Area Computer Networking is part of an effort to develop standards and guidelines for Federal agencies on the implementation and utilization of local area data communications networks. This is felt to be an area that will be receiving increased attention over the next several years, and one for which adequate guidance does not yet exist. For example, we have been conducting an investigation into the best way to meet NBS needs for interconnecting large numbers of simple terminals and minicomputers and a modest number of full sized host computers on the NBS Gaithersburg campus. While surveying available technologies for accomplishing the desired interconnection, two things became clear:

1. Many other organizations, including Government and civilian laboratories, office complexes and factories, felt the same need as NBS for local area data communications capabilities; and
2. People in these other organizations who were also investigating local area networking technologies, and in many cases even building prototype systems, were extremely interested to find out what their counterparts elsewhere were doing.

All of the people we contacted during our technology survey* were enthusiastic at the idea that NBS host a Workshop on the subject to be attended by the active investigators in this new area of networking. In addition to the primary goal of eventual standards-making in this area, NBS interest was sustained by the obvious benefits to advancing the state-of-the-art through information interchange among leading-edge researchers and system developers and by the parochial desire to ensure that we had not overlooked any significant candidate solution to NBS local computer networking needs. Accordingly, this workshop was organized and held on August 22-23, 1977 at the National Bureau of Standards Headquarters in Gaithersburg, Maryland.

The workshop was attended by approximately 50 of the most active workers in the local area networking field, including representatives from other Government agencies, universities and industry in the U.S. and abroad. We were very pleased at our success in attracting the right set of

* To be issued as an NBS Special Publication.

attendees, particularly workers with whom we were not familiar or had failed to contact during our survey.

The two days of the Workshop were organized into three plenary and two working group periods. The first plenary session was devoted to short presentations of work in progress by attendees wishing to make such presentations, and by discussion of the working group topics. Three parallel working groups met during each of the two periods allotted; thus six specific topics were covered quite intensively. Chairmen of the first three working groups gave short reports after the sessions on the first day. A Workshop dinner was followed by a "Blue Sky" session in the evening. The final plenary session was devoted to reports from the second three working groups and to general discussion on all topics.

Fifteen short presentations were made in the first plenary session. The abstracts or short paper provided by the presenters along with copies of some of the transparencies used are included as the first section of this report. The presentations were all somewhat abbreviated due to the large number that had to be fit into the morning allotted, but they did serve to portray the various approaches that were being taken and the status of on-going projects. It was evident that a wide variety of different approaches are being tried, spanning a cost domain of at least three orders of magnitude.

Following the short presentations, the afternoon working group session topics were discussed, and the following decisions made as to topic, chairman and recorder:

1. Subnet Architecture
Chairman: David Mills, COMSAT
Recorder: Paul Meissner, NBS
2. Protocols
Chairman: Stuart Wecker, DEC
Recorder: Robert Rosenthal, NBS
3. Applications
Chairman: Philip Stein, NBS
Recorder: Gary Donnelly, NSA

At the end of the afternoon each of the chairmen gave a short report of the discussion and results of the working group. The written reports prepared afterwards by the recorders in coordination with the chairmen are included as the second part of this report.

Before adjourning for the day, the topics, chairmen and recorders for the following morning's working groups were also selected, as follows:

1. Network Architecture & Interconnection
Chairman: Richard Sherman, Ford
Recorder: Robert Carpenter, NBS
2. Network Operating Systems
Chairman: Stephen Kimbleton, NBS
Recorder: James Hanks, Mitre
3. Performance Analysis
Chairman: Ashok Agrawala, University of Maryland
Recorder: William Franta, University of Minnesota

In the evening, most of the day's participants met for dinner at a nearby motel. After dinner, a "Blue Sky" session was chaired by Robert Metcalfe of Xerox. This session, which had been planned to permit discussion of "far out" or "half-baked" ideas which people might be reluctant to suggest in formal sessions, devoted itself to consideration of the parameters for a "standard local area network interface chip." The ensuing discussion was only partially tongue-in-cheek!

In the morning, each working group convened directly to consider its chosen topic. Following lunch, all attendees reassembled in the final plenary session, which began with the Chairmen's reports for the morning working groups. As with the first set of working groups, the recorders' written reports for this set are included in this report.

In the discussion following the reports, there was general agreement on the following points:

1. The technical problems involved in designing local area computer networks are not very different from the problems in designing global networks.
2. A major distinction between local and global networks is the higher degree of control that a single organization is likely to have over the design and operation of a local network.
3. The most pressing problem in the local network field is not technical but rather information dissemination.

4. Some sort of standardization is needed in order to make the production of a "local area network interface" chip economically attractive to a semiconductor manufacturer. This is vital to reducing costs.

There was rather a lively controversy over the expected size of future local area networks. One vendor predicted few networks larger than 3-5 interconnected hosts on the basis that users could not deal with the complexity of larger networks. Other vendors and users countered that the network should be viewed as a resource pool which could be shared by a large number of terminals (simple terminals up to host computers) each interworking with only a few other terminals at any one time. There was agreement that reliability issues had not been adequately addressed by the Workshop, and that reliability problems could limit the network complexity that could be achieved.

Participants expressed satisfaction with the size, duration and general organization of the Workshop. There was some feeling that the short presentations were overly short; it was suggested that general presentations be retained at an opening session but that specialized presentations be moved to the appropriate working group session. It was agreed that another Workshop in about a year's time would be extremely useful, since many local networks currently under development will have reached operational status by then.



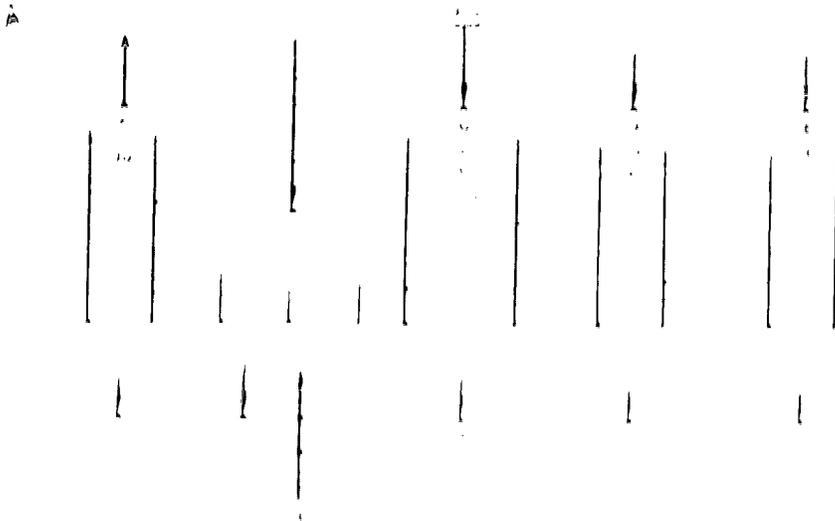
A LOCAL NETWORK FOR THE NATIONAL BUREAU OF STANDARDS

Robert J. Carpenter
Robert Rosenthal
Computer Systems Engineering Division
Institute for Computer Sciences and Technology
National Bureau of Standards
Washington, D.C. 20234

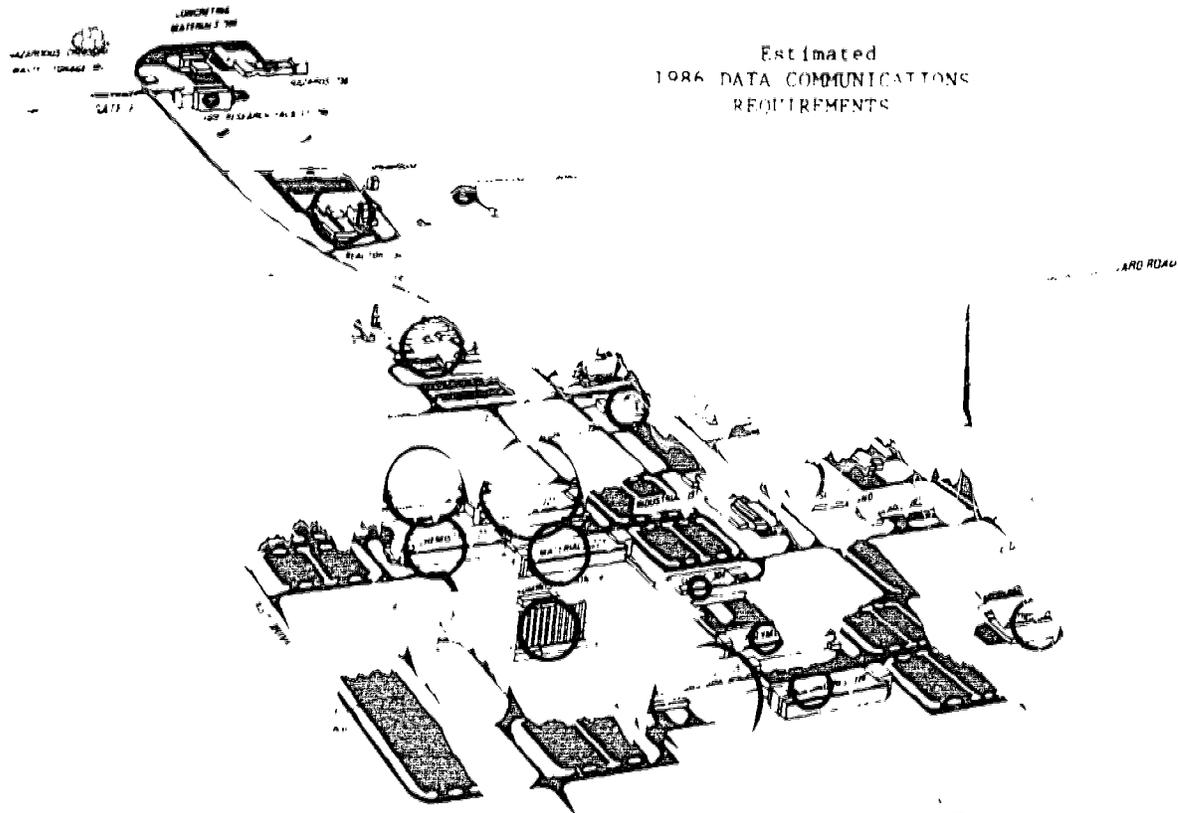
The National Bureau of Standards (NBS) is currently conducting a study of the requirements for a local area network (LAN) to be used by the various divisions of the Bureau. The study is being conducted in order to determine the requirements for a LAN that will be able to support the various types of computers and peripheral devices that are currently in use at the Bureau. The study is being conducted in order to determine the requirements for a LAN that will be able to support the various types of computers and peripheral devices that are currently in use at the Bureau.

The study is being conducted in order to determine the requirements for a LAN that will be able to support the various types of computers and peripheral devices that are currently in use at the Bureau. The study is being conducted in order to determine the requirements for a LAN that will be able to support the various types of computers and peripheral devices that are currently in use at the Bureau.

A connection retry protocol is proposed that provides a software "rotary" allowing this standard TIE to be used on host computer ports with a single connection address, automatically incremented to bypass busy or broken ports.



Estimated
1986 DATA COMMUNICATIONS
REQUIREMENTS



Building	Area	Microcomputers	Other
102			
105			
120			
121			
122			
123			
124			
125			
126			
127			
128			
129			
130			
131			
132			
133			
134			

APPLICATIONS OF HYPERCHANNEL (TM)

Gary S. Christensen
Network Systems Corporation
Brooklyn Center, Minnesota

Present examples of actual and planned applications of
HYPERCHANNEL local computer network. These applications
include:

1. ...
2. ...
3. ...

DATA RING AT COMPUTER LABORATORY, UNIVERSITY OF CAMBRIDGE

A. Hopper
Computer Laboratory
University of Cambridge
Cambridge, England

The data ring at Cambridge was designed to provide a high speed, low error rate communication path between computers and other devices in the Computer Laboratory. These devices are connected through the ring on an individual basis, and yet there are no global high level protocols to provide automatic call establishment. The primary use of the ring is for file sharing and file dumping.

The data ring is a single bus system. It consists of a central controller and a number of peripheral devices. The central controller is a PDP-11 computer. The peripheral devices are a number of other computers and devices. The data ring is used for file sharing and file dumping. The data ring is a single bus system. It consists of a central controller and a number of peripheral devices. The central controller is a PDP-11 computer. The peripheral devices are a number of other computers and devices. The data ring is used for file sharing and file dumping.

The data ring is a single bus system. It consists of a central controller and a number of peripheral devices. The central controller is a PDP-11 computer. The peripheral devices are a number of other computers and devices. The data ring is used for file sharing and file dumping. The data ring is a single bus system. It consists of a central controller and a number of peripheral devices. The central controller is a PDP-11 computer. The peripheral devices are a number of other computers and devices. The data ring is used for file sharing and file dumping.

The data ring is a single bus system. It consists of a central controller and a number of peripheral devices. The central controller is a PDP-11 computer. The peripheral devices are a number of other computers and devices. The data ring is used for file sharing and file dumping. The data ring is a single bus system. It consists of a central controller and a number of peripheral devices. The central controller is a PDP-11 computer. The peripheral devices are a number of other computers and devices. The data ring is used for file sharing and file dumping.

When a station has a packet ready for transmission in its shift register it waits until the beginning of the next slot. It now reads the full/empty bit and at the same time writes a one at the output. If the full/empty bit was zero it transmits the packet, if however the full/empty bit was a one the slot is already occupied and the algorithm is repeated for the next packet. This scheme minimises the number of bits delay at each station.

The control bits are sent in a separate channel and are not subject to collision. If the packet is not sent, the station is marked as busy. If the packet is sent, the station is marked as free. The station also computes the number of packets in the ring and the time to wait before it can transmit.

The control bits are sent in a separate channel and are not subject to collision. If the packet is not sent, the station is marked as busy. If the packet is sent, the station is marked as free. The station also computes the number of packets in the ring and the time to wait before it can transmit.

destination.

If one of the SOP bits is corrupted or the full/empty bit becomes full then this will be detected and corrected by the monitor station. If the full becomes empty then the packet might be ignored at the destination but this will be detected by the source. Similarly the transmitter will detect if the monitor station bit becomes corrupted in such a way that the slot is marked empty. An error in the address fields may cause the packet to be delivered incorrectly or be assigned to the wrong station. An error in the response bits might have more serious effect as it will not be detected by the transmitter which might repeat the packet. It seems it was received correctly here as was not the case. Under some circumstances such errors can propagate to other stations if they are detected by the monitor station with a delay. The delay is a function of the rate.

has no ambiguity about the start of a digit (unlike phase modulation). A change on both pairs indicates a one and a change on only one pair indicates a zero, each pair being used alternately. The advantages of the four wire system can be summarised as follows:

1. It is a balanced system

2. It is a differential system

3. It is a long distance system

Discussion

The Cambridge ring was designed in an environment where many different types of machines exist and where the disruption to their operating systems has to be minimal. This differs significantly from systems where the designer has a free hand to develop host software according to his wishes, and especially from systems which connect a large number of identical machines. Furthermore, it was a design task to make the system as inexpensive as possible, and it was thus kept simple. Nevertheless, many options are left and some of these are discussed below.

The Cambridge ring is a multi-processor system. This can be seen from the multiplication of a set of different processors. If the system is to be used as a system for real time control, the speed of the processor and the time to receive and to send messages are of primary importance. For a given processor, the speed characteristics can be improved by employing a micro-processor. However, the number of processors is limited at different times. The number of processors can be associated with each processor. This number is updated when the relay has been updated. The number of processors can also be updated when the processor is updated.

The Cambridge ring is a multi-processor system. This can be seen from the multiplication of a set of different processors. If the system is to be used as a system for real time control, the speed of the processor and the time to receive and to send messages are of primary importance. For a given processor, the speed characteristics can be improved by employing a micro-processor. However, the number of processors is limited at different times. The number of processors can be associated with each processor. This number is updated when the relay has been updated. The number of processors can also be updated when the processor is updated.

LOCAL MISSION-ORIENTED NETWORK

Ronald L. Larsen
National Aeronautics & Space Administration
Goddard Space Flight Center
Greenbelt, Maryland

In the past few years we have been studying the feasibility and design of a local mission oriented network for the support of ground based support operations for spaceflight missions. Within the context of this study have been investigations of the workload to be placed on the system and its proper distribution among its components, as well as its system architecture and configuration trade-offs.

One of the major tasks in the design of a local mission oriented network is the determination of the workload to be placed on the system. This workload is determined by the number of missions to be supported, the duration of the missions, and the complexity of the operations. The workload is then distributed among the components of the system, such as the ground stations, the communication links, and the data processing centers. The distribution of the workload is determined by the relative capabilities of the components and the requirements of the missions. The workload distribution is then used to determine the system architecture and configuration trade-offs.

ETHERNET: DISTRIBUTED PACKET SWITCHING FOR LOCAL COMPUTER NETWORKS

Robert M. Metcalfe
David R. Boggs
Xerox Corporation
Palo Alto, California

Carrying digital data packets among locally distributed computing stations, the packet transport mechanism provided by Ethernet has been used to build systems which can be viewed as either local input network or local output network multiprocessors. An Ethernet's shared communication facility is essentially passive to each station with no central control or limitation of access to the other for packet broadcast or distribution among the stations transmitting packets using controlled statistical arbitration. Multiple packet transmissions on the Ethernet network are handled by a simple algorithm which permits each station to transmit packets at will and implement a distributed control system. The system can operate with up to 100 nodes and is highly scalable. A number of interesting applications of Ethernet are discussed in this paper.

been operational for three years. The major characteristics of the new communications system are as follows:

- o A design capable of implementation using a single chip LSI transmission controller (LNI) and incorporating a flexible, process oriented addressing structure.
- o A transmission system which, in its normal configuration, is a ring architecture utilizing a single unidirectional twisted pair operating at a transmission speed of more than one megabit. (The expected rate is in the range of two to four megabits).
- o A unit failure bypass mechanism to enable continued system viability in the event of unit failure.
- o An addressing structure and acknowledgement mechanism supportive of a distributed processing environment.
- o A possibility of operating in a variety of communications topologies including the "Ethernet" protocol, a contention ring, and others.

This research is supported by the Advanced Research Projects Agency under Contract N00014-76-C-0954.

[1] Mockapetris, Lyle and Farber, On the Design of Local Network Interfaces, IFIP 77.

COMPUTER CELLS--HIGH PERFORMANCE MULTI COMPUTING

David L. Nelson
Prime Computer, Inc.
Framingham, Massachusetts

Current research activities at Prime Computer, Inc., for the development of high performance multicomputer networks are described. Certain methodologies for the design of such systems have been previously presented [1] and are herein extended to include engineering and manufacturing trends. These considerations suggest future architectures that will be comprised of moderately coupled, highly regular, local homogeneous multicomputer networks which are characterized as high performance data flow computers. Described are the current designs for

interprocess communication (pipelines), system packaging and interconnection (fiber optics ring), process to processor binding, and thoughts regarding problem decomposition.

- [1] Eckhouse, R. H. and D. L. Nelson, "Distributed Operating Systems -- An Approach to Greater Flexibility", Proceedings of the 5th Texas Conference on Computing Systems, University of Texas at Austin, October 18-19, 1976, pp. 157-159.

THE MIT LABORATORY FOR COMPUTER SCIENCE NETWORK

K. T. Pogran

and

D. P. Reed

MIT Laboratory for Computer Science
Cambridge, Massachusetts

The MIT Laboratory for Computer Science is developing a local area network which will initially serve the needs of our laboratory and which, we hope, will form the basis of an eventual campus-wide network. The immediate objective of the LCS Network is two-fold: first, to provide an intercommunication capability for the ever-growing collection of minis, micros, and larger-scale systems within the Laboratory, and, second, to provide a vehicle for the Laboratory's research in the area of distributed computing.

In developing the LCS Network, we have tried to take a "total system" approach, concerning ourselves from the outset not only with architecture and hardware issues, but with protocols as well, and with such issues as: interfacing the network to already-existing systems, large and small; the impact of a high-bandwidth network on small systems, and providing economical access to a high-bandwidth network for terminals which are, by comparison, low-speed devices.

Technology and Architecture

We began two years ago by studying some of the technologies then available for local networks. Both the Ethernet and the Farber Ring Network offered the attributes of high bandwidth and completely distributed control, and we restricted our study to these two technologies. We realized that both offered essentially the same functional capabilities; in addition, we realized that, with properly designed interface hardware, a network using either basic technology could present the same logical interface to a

host. Finally, we concluded that the same basic interface hardware could be used with either network technology, with only minor modifications to its control structure and internal data flow.

Therefore, in the fall of 1976 we decided to join forces with Dave Farber's group at UC-Irvine to develop a single "Local Network Interface" which could be used for either a Ring Network or an Ethernet. Implementation of the initial "ring-only" version of the LNI, running at 1 Mb/s, is nearly complete; this fall we will be developing the modifications required for Ethernet use. We are hopeful that the Ethernet LNI will be able to operate in the 4-8 Mb/s range.

The LNI has been designed from the start with an eye toward Large Scale Integration. Once its design has been finalized, it should be possible to implement most of it on a single chip, thus making the eventual LNI a very inexpensive device.

The LCS Network will be composed of a number of "sub-networks," some using Ethernet technology and some using Ring technology, all using identical protocols, and sharing a single "address space." The sub-networks will be interconnected by means of relatively simple hardware "bridges"; the network as a whole will be connected to the ARPANET via a PDP-11 "gateway" system. This "sub-network" architecture will enable us to evaluate the relative merits of the Ethernet and Ring Net technologies; it will allow us to try out new technologies within our overall network, and it will provide us with a straightforward method of coping with future traffic growth.

Protocol Issues

The LCS Network will not exist in a vacuum. As was mentioned above, our plans already include interconnection to the ARPANET. For this reason, a primary goal in the design of protocols for the LCS Network was to incorporate at an/early stage the necessary flexibility to have each host computer, microprocessor, or terminal connected to the LCS Network participate in communications with systems outside the local network in the same way that communications occur within the net. We are thus seriously involved in the internetworking game.

In looking at the protocols available, only TCP (Transmission Control Protocol, Cerf & Kahn) seemed to attack most of the problems of addressing, technology matching, etc. Unfortunately for us, though, TCP seemed

somewhat more complicated than it had to be, so we have developed a variant called the Data Stream Protocol (DSP) that we believe is simpler than TCP. DSP is still under evolution, as is TCP, and it is our hope that they will eventually merge into a truly simple but general internetworking protocol.

We are currently trying to look at very flexible addressing schemes within the networks to allow both generic addressing of services by name in an internetworking environment where services are dynamically created and destroyed, and to improve routing of packets in an internetwork environment where gateways may choose not to participate in "optimal routing" negotiations.

While we are not currently developing new higher level protocols (we expect to use existing ARPANET TELNET and File Transfer Protocol software as our initial higher level protocols) we expect to evolve much more effective protocols to deal with distributed data as time goes on.

An important goal in our participation in an internetworking environment is to secure our communications against unauthorized prying. Our experience in designing the Multics system leads us to believe that protection is an absolute requirement, even within a university environment. Consequently, we will be experimenting with the use of end-to-end encryption, probably with the NBS algorithm, integrated into our end-to-end protocols. We feel that a protocol with features such as those of DSP or TCP is the right sort of protocol for use with end-to-end encipherment.

CURRENT SUMMARY OF FORD ACTIVITIES IN LOCAL NETWORKING

R. H. Sherman
M. Gable
G. McClure
Ford Research and Engineering Staff
Dearborn, Michigan

Ford is designing a communication network named CYBERNET to support local decentralized computing for real time data acquisition and control in the manufacturing system. The decentralized broadcast media is similar to that employed by the Xerox Ethernet. Communication connections in CYBERNET, however, are made between processes, not hosts and terminal oriented devices. The communication media is cable television (CATV) coax. Connections to the coaxial cable may employ terminated

resistive taps or a low loss cable impedance matched daisy-chain. In trunk branching situations a lossy tap fabricated from CATV power splitters is permitted. The transceiver electronic module is 15 conventional chips and serves the function of modulation/demodulation, signal amplification, idle detection, collision detection, and retry timers. The ability to detect collision under signal attenuation has been demonstrated. The modulation is baseband PCM (+-3.5 volts) using a D.C. balanced, self synchronizing, encoding of the data bits which requires six times the data bandwidth. The transceiver can be used with micro(or mini) computer serial ports provided all of the serial ports on the network use the same baud rate. For a higher performance network with synchronous, bit stuffed, 1.3 megabaud data rate, a fast microprocessor based adapter is combined with the transceiver for functions of packet switching and error control. The message protocol includes free formatted destination, source, control and data fields. The prototype network is being implemented in Research for laboratory automation. Stations will include a PDP-10 computer, an engine dynamometer test facility, a numerically controlled machine tool and an operator station.

The system is being designed to allow interconnection of networks using gateways in order to provide full support of resources. The network protocol is designed to make these interconnections as simple and reliable as possible. The gateways need not contain routing tables associated with the network topology since the message header contains the complete route (pathname) from the source to the destination. This pathname is dynamically constructed during the communications process by each gateway concatenating its name to the source name field and removing its name from the destination field of the message.

LOCAL AREA NETWORKS AT QUEEN MARY COLLEGE

Anthony R. West
Computer Systems Laboratory
Queen Mary College
Mile End Road
London E1 4NS, England

The group in the Computer Systems Laboratory has two main areas of research interest: the first is in the design of low-cost computer systems to promote a high degree of user-interaction; the second is in the architecture of distributed computer systems. In both these fields, the use of a number of low-cost micro- or minicomputers which

cooperate with one another to tackle user applications is not feasible unless some convenient "glue" exists for "sticking" systems together. We see this glue as taking the form of a high-bandwidth local-area computer network, and have been considering network designs for some time now. Although we had been interested in ring networks similar to that used in DCS at the University of California at Irvine, our interest really took shape with the appearance of the paper on the Ethernet Network in use at the Xerox Palo Alto Research Center [Metcalf 76]. We set out to design a similar bus-oriented contention network based on current 8-bit microprocessor technology.

At that time (about August 1976), we came into contact with a group working at the Rutherford High Energy Physics Laboratory who had a requirement for a flexible, extensible, fast, local network to improve the facilities for resource sharing at their site. It was felt that the best way to satisfy their requirement and ours was to embark on a joint development project to construct an Ethernet-like network, which we may decide to call the ENET. This work has been in progress for six months now, and three prototype node controllers based on the Motorola M6800 micro are nearly ready. Testing should take place starting in September. The data rate down the coaxial cable is 3 megabaud and the cable can be up to 2 km long (at present).

In the meanwhile, whilst Rutherford are working on the hardware, we decided to hack together out of the standard building bricks of our M6800 development system a similar network (but of much lower bandwidth) to investigate the software structures and problems inherent in such a network. This network, the CNET (C stands for Cheap!), is based on standard M6800 Asynchronous Communications Controller Circuits with open-collector line-transceivers to interface to the shared coax. The data rate of 76.8 kilobaud is fairly low, but all the software is interrupt driven instead of requiring DMA facilities like the ENET (E stands for Expensive?!). This CNET was intended to give us an accurate model of the future ENET controller (except for DMA) so as to give us a chance to investigate protocol questions in advance of the availability of the Rutherford Hardware. Many of the questions we propose to study are described in [West 77].

The ENET uses synchronous communications and there is a possibility that, after the prototypes have been tested, the production circuits will use HDLC interface circuits. This is highly dependent on the performance of the forthcoming chips and the as yet unknown desirability of using HDLC in both Rutherford's and QMC's contexts. The College Computer

Centre currently plans to purchase an ICL 2980 with two HDLC ports, one for an X.25 link to a London University X.25 exchange, the other for a local college network, which could be an ENET packet gateway.

Both networks use passive coaxial cable for a shared transmission medium, and, in order to protect this medium from being corrupted by unintentional (or intentional) pollution from nodes, a self-testing facility has been proposed. At regular and frequent intervals (e.g., every 10 seconds?) controllers queue a packet which is sent out onto the network and received by that same controller again. If the entire transmission and reception paths (in both software and hardware) check out, the microprocessor refreshes the timer on a relay. If this relay is not refreshed within some time interval (like 15 seconds?) it opens, disconnecting that node from the net and initiating a restart in the microprocessor software. The node then checks itself (by sending itself a packet without being connected to the network) and if it is functional, connects back to the Ether again. If a failure occurs immediately, the node repeats the process once more before deciding that the Ether is unusable and sounding an alarm.

Our laboratory also houses some undergraduate teaching facilities for Computer Science students. At the moment, these take the form of a PDP 11/40 running the UNIX operating system, and a number of satellite microcomputers supporting intelligent terminals, etc. We are about to acquire a PDP 11/34 and several ISI-11's for research into some Man-Machine Interface questions (like text processing in the distributed office environment) and we intend to start by connecting these to the CNET (and later to the ENET) in order to bootstrap software and share resources. The possibility of interfacing UNIX, local-area networks and long-haul X.25 networks is attractive.

Rutherford has also received approval to build a satellite ground station for a broadcast satellite network to link various research establishments in Europe. We propose to study ways to enable ENET users at Rutherford to send data via satellite to other sites (which may grow interested in local ENET's later).

References

- Metcalfe 76 "Ethernet: Distributed Packet Switching for Local Computer Networks," R. M. Metcalfe and D. R. Boggs, Communications of the ACM, Vol. 19, No. 7, July 1976.

West 77 "A Broadcast Packet Switched Network Design Progress Report," A. R. West, Computer Systems Laboratory, Queen Mary College.

DECNET: ISSUES RELATED TO LOCAL NETWORKING

Stuart Wecker
Research and Development Group
Digital Equipment Corporation
Maynard, Massachusetts

Abstract

The Digital Network Architecture (DNA) is the framework for DECnet implementations. Its goal is to provide efficient and flexible networks for both global and local environments. This paper presents some of the issues and tradeoffs made during the design of DNA which relate specifically to local networks.

Introduction

The Digital Network Architecture (DNA), the framework for the DECnet family of implementations, creates a general networking communication base within which programs and data can be easily accessed and shared. DNA is designed to provide this general resource sharing and distributed processing across a broad range of hardware and software components. It is designed to be efficient in network structures ranging from small local networks of 2 or 3 minicomputers in a single room, to large geographically distributed networks of many large mainframes.

The general approach taken in the architecture is to partition the system functions into: (1) communications, (2) networking, and (3) applications. These are then implemented in a layered structure, each layer creating a richer environment for the layers above it, providing them with a set of functions upon which they can build. A more detailed discussion of the architecture and its components can be found in [1, 2, 3, 4].

The layers and their functions are each reflected in a protocol which provides the communications and synchronization between corresponding layers in the distributed computer systems. The layers of DNA and their functions are:

Communications. The goal of this layer is to create a

sequential error-free data link for the movement of data over a communication channel. Here we are concerned with detecting and correcting bit errors introduced by the data channel and with the management of multipoint and half-duplex channels. The protocol used in DECnet within this layer is DDCMP (Digital Data Communications Message Protocol). Other protocols providing similar functionality are SDLC[5], HDLC[6], and ETHERNET (contention level protocol) [7].

Networking. The goal of this layer is to create a process-to-process communication mechanism that is sequential and flow controlled for the movement of data between communicating nodes. Here we are concerned with routing data between nodes, creating logical data paths between users and providing integrity, sequentiality, and flow control on these data paths. The protocol used in DECnet within this layer is NSP (Network Services Protocol). Other protocols providing similar functionality are the packet level protocols of SNAP (X.25) [8] CYCLADES [9], and the ARPANET Host-to-Host protocol [10].

Applications. The goal of this layer is to create a mechanism for the movement of application data between communicating processes and/or resources. Here we are concerned with communicating with, for example, I/O devices, disk files, system loaders, and the distributed programs of a user application. There may be many application protocols executing in a DECnet network. Some are user created protocols; others are DEC provided such as DAP (Data Access Protocol) used to access files in the network. Other protocols with similar functionality are the ARPANET FTP (File Transfer Protocol) and TELNET (Terminal Access Protocol).

Local Networking

In general, the characteristics of applications and their demands on the network are very similar in both local and global networks. User programs want to communicate with other user programs, access I/O devices and files, and interact with terminals located at other nodes in the network. However, the topologies and physical components used in local networks (systems located within a small geographic area) may be different from those used in large geographically distributed ones. For example:

1. No backbone communication network. Many local networks consist of a small number of host systems directly connected without front-end communication computers or a separate backbone communication network. The hosts perform all

routing, networking, and communication functions.

2. Use of many specialized link types. Local networks tend to choose data links based on interface costs, link length, and performance. Thus, they are able to use asynchronous links, parallel links, and high speed loops, not always available over large geographic areas.

3. Direct communications. Many local networks are topologically configured with direct point-to-point connections between the end communicating hosts. Typical topologies are stars, trees, and multipoint links, such as loops and ethers, where any node can communicate with any other node, forming completely connected networks.

There are non-geographical differences as well:

1. No central maintenance control. Many small local networks operate without any node being in control of the topology or maintenance of the network, as is usually the case in large networks.

2. Simple routing requirements. Many local networks have no routing requirements at all since they are directly connected. Those that do are usually very simple (either the operator makes changes via commands, or plugs in alternate cables).

Many of these factors were considered in the design of DNA and its protocols. The result is that many features have been included to enhance DECnet's efficiency in local networking environments. Some of these design features are:

1. Common network level protocol. All nodes in a DECnet network are equivalent at the network level. The characteristics of a node (host, front-end, router) depend on its functional use and physical location in the network. Host computers use the same NSP protocol to communicate with other host computers as they do to communicate with front-ends and switching nodes. In local networks the hosts may be directly connected without intervening communication computers. The addition of communication computers and/or a backbone communication network is transparent to the hosts, since they use the same protocol to communicate with the communication computers as they do to communicate with other hosts. Thus, some nodes may be "front-ended" to off-load some communication functions, while others, with excess processing capability, may directly communicate in the network without using a front-end. This commonality of protocol gives the user the flexibility to configure the network based on application requirements and computing node

capability, rather than on networking structure requirements.

2. Subset level of network protocol. Directly connected networks, where the end communicating systems are connected via a direct channel, are common in local configurations. Here, some functions of the NSP protocol may be omitted to eliminate the duplication of function with other levels, and increase the efficiency of the network. On directly connected links, the link level protocol provides an error-free sequential end-to-end path. The normal end-to-end functions of timeout and retransmission are omitted from NSP for simplification in the hosts. In more geographically distributed networks, communication computers can be added to perform the routing, timeout, and retransmission functions needed in these topologies. These functions are added via an intercept function in the communication computers, which accepts the subset NSP protocol from the hosts and adds these features, creating a superset protocol, suitable for use in these larger networks. This interception is totally transparent to the host. A host may participate in both a local (directly connected) and global network using the same network protocol, the host protocol code always being optimal for the environment in which it executes.

3. Extensible fields. For efficiency in small networks many of the NSP protocol fields have been made variable-length. This allows efficient use of short fields within small local networks, while allowing expansion for use in larger ones. Some examples are the node address, logical link address, process name, and accounting fields. In addition, a hierarchical addressing scheme has been used, dividing node addresses into node areas and addresses within areas. In local networks all nodes may be in the same area, reducing the addressing in the system.

4. Independence of link level protocol from physical link characteristics. The DDCMP protocol was specifically designed to be as independent as possible from the specific characteristics of the data link. Synchronization is defined specific to each type of data link used. A byte count field is used to locate the end of a message, detaching it from any specific characteristic of the link. It has been implemented on synchronous, asynchronous, and 16-bit parallel channels.

5. Optional routing header and changeable algorithm. The NSP routing header may be omitted in directly connected systems. In these configurations the receiver assumes that it, itself, is the destination and the sender is the source.

This increases the overall bit efficiency of the protocol when used in such configurations. The routing algorithm is independent of the NSP protocol and may be changed based upon the requirements of the configuration. In local networks, simple algorithms, such as change on operator command, may be implemented while complex adaptive algorithms may be used in large global networks.

6. Layered structure. In layered structures, each layer performs specific functions while hiding the techniques and protocol used within that layer. The layer is only visible through the interfaces with which it communicates to the layers above and below it. This allows clean replacement of layers with functionally equivalent layers. Implementation of DECnet in a ring or other structured environment only requires replacement of the DDCMP protocol with a suitable ring or contention link level protocol. This would result in a fully connected network, and allow use of the directly connected subset of the network protocol described above, eliminating duplication of function.

7. No central maintenance. The maintenance features of the network have been made independent of the basic structure, as is done with the routing algorithm. The network will operate with each node executing independently, coordinating with the other nodes via the protocols. Any maintenance features can be added at higher levels in the structure to provide overall control of the network. In small local networks such features may not be necessary, and may be omitted.

8. Routing in a star topology. The intercept function, described in (2) above, includes the capability for routing between end nodes connected via a single intermediate node. For example, a star configuration where the central node performs the intercept routing for the network. In this case the end nodes use the directly connected subset of the network protocol. They may have messages routed to other nodes without the addition of communication computers or use of the superset protocol. The central node may participate in application level functions as well. These topologies are very common in local area networking structures.

Conclusions

The requirements of local networks are not significantly different from those of large geographically distributed networks. Differences exist mostly in their physical topologies and data link characteristics. The protocols designed for local networks must perform the same functions as those designed for larger global networks. By

taking into account the geographical differences in the design of the protocols comprising the DECnet architecture, the same protocols and structure are used very effectively in local networking situations, without compromising their effectiveness in large geographical applications.

References

1. Wecker, S., "The Design of DECnet-A General Purpose Network Base," IEEE Electro 76, Boston, MA, May 1976.
2. Conant, G. and Wecker, S., "DNA: An Architecture for Heterogeneous Computer Networks," Third International Conference on Computer Communications, Toronto, August 1976.
3. Wecker, S., "DECnet: A Building Block Approach to Network Design," National Telecommunications Conference, Dallas, TX, November 1976.
4. Passafiume, J. and Wecker, S., "Distributed File Access in DECnet," Second Berkeley Workshop on Distributed Data Management and Computer Networks, Berkeley, CA, May 1977.
5. Synchronous Data Link Control - General Information, GA27-3093, IBM Corp., March 1974.
6. High Level Data Link Control Procedures, ISO/TC97/SC6, International Standards Organization, October 1975.
7. Metcalfe, R. and Boggs, D., "Ethernet: Distributed Packet Switching for Local Computer Networks," Communications of the ACM, July 1976.
8. Standard Network Access Protocol Specification, Trans Canada Telephone System, March 1976.
9. Pouzin, L., "Presentation and Major Design Aspects of the Cyclades Computer Network," Third Data Communications Symposium, IEEE, November 1973.
10. Crocker, S., et. al., "Function-oriented Protocols for the ARPA Computer Network," AFIPS 1970 SJCC, 1970.

Session: SUBNET ARCHITECTURE

August 22, 1977

13:15 - 16:30

Chairman: David Mills, COMSAT

Recorder: Paul Meissner, National Bureau of Standards

Participants: not recorded

The following topics were identified as candidates for discussion:

- Transceiver design
- Cable technology and the "tap" Isolation
- Optical transmission
- Interfaces
- Architecture
- Virtual circuits versus datagrams
- Guaranteeing performance
- Availability
- Standards

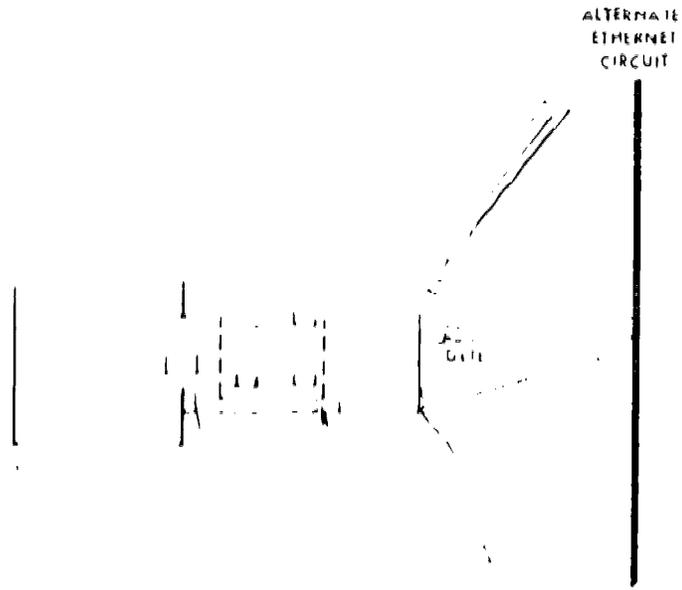
Participants requested that their names not appear in the session notes. The session opened with a discussion of standards, the general feeling being that standards would be premature for most aspects and would inhibit innovation. Later in the discussion it was suggested that the implementation of certain functions in LSI would be aided by standardizing these functions in order to achieve a sufficient production base to offset the LSI design and set-up costs.

It was observed that much of the existing design effort has been done without the benefit of extensive RF engineering, and participants were invited to comment on their experience in this regard. Some implementations were developed around Jerrold Electronics CATV equipment, including facilities for taps. An example was given of 3-megabit operation with 3-volt signals. The Jerrold equipment was modified to reduce losses introduced by taps. Coaxial cables were contrasted with twisted pairs for transmission lines. Coax was cited as advantageous in that a single cable could be run throughout an installation and used for a variety of services through multiplexing. It was noted that proper grounding can be a problem, since there may be voltage differences in the grounds between different facilities and this can result in heavy currents in the coax shield.

Tradeoffs between implementing functions in hardware or software were discussed. If speed is the governing factor, it is generally necessary to resort to hardware. Designer preference is likely to play a substantial part in the decision. An example was given showing OK hardware for receiving because of the large volume of received data but doing the RC in software for transmits because of the less data generated at the station.

Also discussed was the possibility of using a microprocessor for addressing, etc. A name table was described which would allow the system to be set up for the use of a microprocessor in software.

Another participant's approach is shown below:



Availability was discussed briefly. Some examples were given of loopback features which enable a station to perform self-testing. However, this requires two buffers, and a minimal system cannot afford this. The use of monitoring on the Packet Radio Network was cited as a means of determining circuit loading.

The following table shows the results of the tests conducted on the Packet Radio Network. The table shows the number of stations which were able to receive and transmit data. The table also shows the number of stations which were able to receive and transmit data at the same time. The table shows the number of stations which were able to receive and transmit data at the same time. The table shows the number of stations which were able to receive and transmit data at the same time.

The following table shows the results of the tests conducted on the Packet Radio Network. The table shows the number of stations which were able to receive and transmit data. The table also shows the number of stations which were able to receive and transmit data at the same time. The table shows the number of stations which were able to receive and transmit data at the same time.

Session: PROTOCOLS for LOCAL AREA NETWORKS
August 22, 1977
13:15 - 16:30

Chairman: Stuart Wecker, Digital Equipment Corporation

Recorder: Robert Rosen, National Bureau of Standards

Panelists:

Jacques	
Gary Christensen	Netw
David Nelson	Prime Computer, Inc.
David Reed	MIT/L S
Edward Row	National
Richard She	Ford Motor Company
David Wilner	Lawrence Berkeley Lab.
John Wood	Lawrence Berkeley Lab.
Jeff Yeh	Lawrence Livermore Lab.

ERIC
Full Text Provided by ERIC

- 3. Network Level (local)
 - routing
 - congestion control
 - virtual circuit creation
 - end-to-end acknowledgment
 - flow control
 - sequencing

I/O device control
 file access

OS
 ST
 as any, (ed)

SS1
 ...

...

...

1. high bandwidth (so high that the communications network itself is not a bottleneck)

2. low delay (measured in seconds in the network, very long)

3. reliability (measured in terms of network links)

networks

networks

networks

networks

bandwidth allows you to stay simple and use the network as an extension of a given local capability.

Why not use a standard protocol for all local networks? Because in many cases we are not solving a general problem with local nets. We build local networks to solve specific intercommunication problems. Usually the solutions require very high bandwidth.

Security is a very much application dependent problem. In general, we do not have little problems. Especially since the network itself might be physically contained in a small geographic area under one management. Usually, physically large networks have security problems.

Security is a very much application dependent problem. In general, we do not have little problems. Especially since the network itself might be physically contained in a small geographic area under one management. Usually, physically large networks have security problems.

Security is a very much application dependent problem. In general, we do not have little problems. Especially since the network itself might be physically contained in a small geographic area under one management. Usually, physically large networks have security problems.

Session: LOCAL NETWORK APPLICATIONS
August 22, 1977
13:15 16:30

Chairman: Phillip... Bureau of Standards

Recorder: Gary... Agency

- Phil...
- Gary... NSA
- George McCl... FOR
- Jim Banks... MIT
- Louis Pou... Inst. of R
at d'Aut...
- NB

... the ...
... 10 ...
... this ...
... specifically ...
... existing ...
... application ...
... that the network ...
... meaningful ... use

... the ...
... the ...

... the ...
... the ...
... the ...
... the ...
... the ...
... the ...

it was stated that the user doesn't care if the mean response time is 1.5 seconds with a deviation of 0.5 seconds. What he really cares about is that no task take more than two seconds. Therefore, one measure that is useful to this user is a guaranteed maximum response time.

The following is a list of parameters that are set for an application set of the network to know while the definitions are not complete, it was felt that they represent a minimal set.

1. The definition of the network topology. The topology description file is a text file that contains the network topology. (This is a standard file name on the system.)

2. The definition of the network parameters. The network parameters file is a text file that contains the network parameters. (This is a standard file name on the system.)

3. The definition of the network resources. The network resources file is a text file that contains the network resources. (This is a standard file name on the system.)

4. The definition of the network services. The network services file is a text file that contains the network services. (This is a standard file name on the system.)

5. The definition of the network users. The network users file is a text file that contains the network users. (This is a standard file name on the system.)

7. What is the long term data rate that the network will support? Over what period of time may this rate be specified?

In addition to the... always generated with the end user demand, it was also felt that computer networks suffer from a lack of human factors engineering in their design. Networks must be considerate of the human interfaces to it as well as the machine-to-machine interfaces.

In computer environment, it has been... of individual computers when the network should appear to be totally transparent to the user. There should be no hint of characters for the net, and perhaps some out of band signalling scheme should be used by the network.

... deterministically... should give predictable... is not the case in general... networks... use of a deterministic network... the user may... and its... despite the... including... all with the... all... important... a micro... of a human... the... in... 4

... 4

Session: NETWORK ARCHITECTURE

August 21, 1977

0900-1230

Chairman: Richard Sherman, Ford Motor Company

Recorder: Robert Carpenter, National Bureau of Standards

Participants:

Robert Carpenter, NBS

Ed Hart, APL

Gregory Hopkins, Mitre

Bruce Lucas, NBS

Andrew Pilipchuk, University of Maryland

Ken Pogran, MIT

Richard Sherman, Ford

David Reed, MIT

Sig Rogers, Lawrence Berkeley Lab

Robert Rosenthal, NBS

Dale Zobrist, Eldec Corp.

The session began with a discussion of the subjects to be covered. The topics decided upon were:

Internetworking

A vertical cut through protocol layers.

Local network to local network interface.

Local network to global network interface.

Network Architecture

Distributed computing

Naming

Monitoring

How to tell the user what is happening.

Language-driven approaches

Data and commands for existing operating systems.

Approaches transferring grammar nearer user.

INTERNETWORKING

Vertical Cut

It was quickly agreed that the interest of this group spanned the full vertical cut through the layers of protocol from the hardware up to the user. This is to be compared with the interests of the Network Operating Systems session.

Local to local network interconnection

The discussion of the interconnection of local networks led to the division of the problem into two categories; networks or network segments with essentially identical protocols (or at least packet designs), and networks with substantially different designs. In the first case the interconnecting means could be logically simple, Pograd proposing the name Bridge. The second case was felt to be similar to connection to global networks through a Gateway capable of complex protocol transformation.

What Does a Bridge Do?

The Bridge, as described, is essentially a store-and-forward packet repeater with address filtering. The drawing in Fig. 1 formed the basis of the discussion of the bridge and the contrast between a bridge and a gateway. The portion of Fig. 1 to the left of the broken vertical line may be considered as a single local area network, made up of four Segments. In the example, these segments were assumed to operate differently, a high-speed ethernet, a ring, a low-speed ethernet, and a segment of unspecified-but-compatible-design. These segments are connected by bridges.

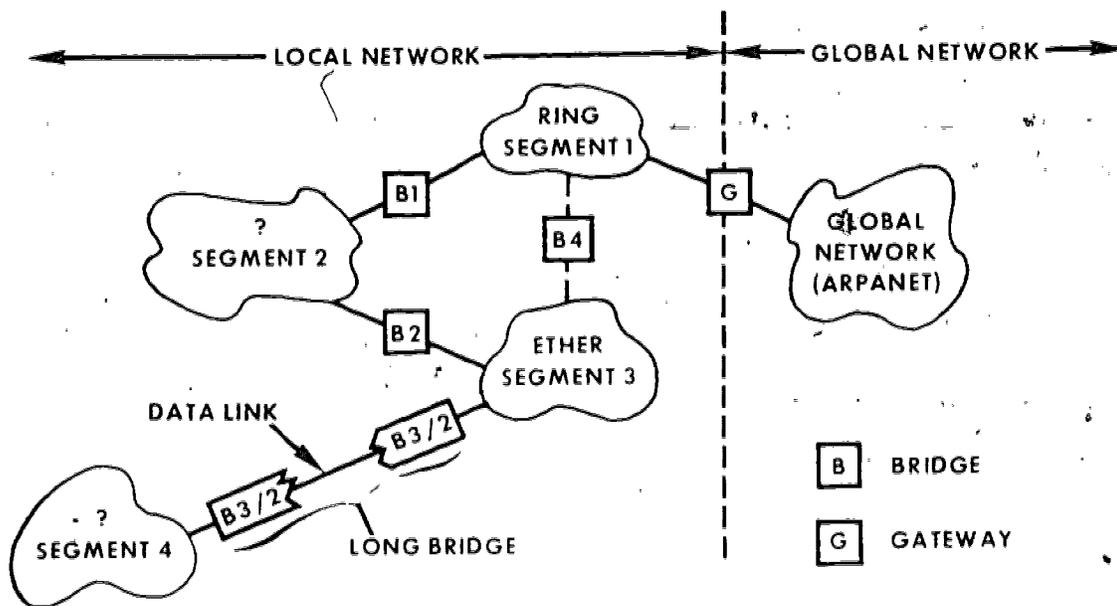


FIGURE 1. MULTISEGMENT LOCAL NETWORK EMPLOYING BRIDGES

The discussion identified the following positive statements about the network and the bridges:

A common address space encompassed all segments of the network(s) interconnected by bridges.

The packet design of all interconnected segments must be the same in all internal address, control and data fields. There may be local control and synchronizing bits local to each segment, that are stripped on entry to the bridge and added on exit on the new segment.

The packet repeater has a limited packet buffering ability and merely ignores further packets when all its buffers are full.

The bridge performs an address-filtering function. That is it examines the destination address in each packet to determine if it need be repeated through onto the other network. If not the packet is not repeated (discarded). This is a powerful function for reducing network load if much traffic is localized on the individual segments.

There must be an end-to-end protocol. The bridge itself does not issue acknowledgements in an ethernet situation and only does such acknowledgement in a ring as is necessary to avoid repeat transmissions.

Since there is storage in the bridge, it effectively breaks the contention area in an ethernet, thus increasing the efficiency of the resulting segments (which is related to the delay between the most distant stations participating in the contention). This is of particular interest in the case of the "Long Bridge" in Fig. 1, where the use of the bridge removes the very long delay and consequent loss of efficiency which would result if the distant segment were connected through a conventional non-buffered repeater amplifier.

There was an extended discussion of the ramifications of alternate or redundant paths as would be represented by bridge B4. This device might be installed to increase network reliability. Each packet from segment 1 would reach segment 3 directly, and another copy by the indirect path through segment 2, thus needlessly doubling network traffic. The situation would become worse with additional redundant paths. It was felt that there might be some oscillatory situations in more complex networks. One possible solution to truncate excessive copies would be to append a hop count

each time a packet was repeated. A limit would be set on the number of times a single copy was repeated.

A bridge would have an extremely rudimentary routing table. In most cases this would consist of the ranges of addresses reached through each of its ports. An alternate solution would be to adaptively form a routing table based on the source field of each received packet.

Some Things Bridges Don't Do

Bridges do not originate packets.

Bridges do not add protocol.

~~Bridges do not contain routing information beyond that required to direct output to their correct port.~~

Bridges do not confirm correct delivery of packets, this must be an end-to-end function.

The Long Bridge

A special version of the bridge can be used where a segment of a local network is some distance away. In this case the bridge is split down the middle with a high-speed data link between the two halves. To conform with the definitions of the bridge, this data link must have sufficient bandwidth that it does not form a bottleneck to data flow. Since each bridge is bidirectional, and loss of contention efficiency is generally associated with the placing of packets on a segment, the packets would presumably be buffered at the end of the link nearest their destination. An interesting consequence of the idea of the long bridge is that a network consisting of two similar segments connected by a long bridge spanning many kilometers (a time delay of several packets), would still meet the definition of a local network.

Gateways

When two networks, or segments, differ substantially in protocol or packet design; or if sophisticated routing strategies are to be employed, the interconnecting device must be of greater complexity. The concept of this Gateway is fairly well understood. It was accepted that the complexity of connection between dissimilar local networks was essentially equal to that of connection between a local network and a global network. A gateway joins two (or more) networks and may perform translations at all levels, from electrical, link-level protocol, end-to-end protocol, teletype protocol or character conversions, and user-to-user conversions such as with error messages.

The performance of a gateway will be poor if it is asked to connect between networks which differ greatly, since similar functions may not exist in the two networks. The example of ARPANET vs. Tymnet was given to contrast a terminal-oriented network and a host-oriented network.

Gateways may do complex routing and may initiate communication with other gateways and hosts.

Reed mentioned a few points concerning gateways that all designers of local networks should remember. The likelihood of eventually desiring to connect a local network to a global network is so great that it is foolish to design a local network without considering this interconnection in the design. Be prepared to use X.25 (or some equally widespread) standard for connection outside the local network. This should minimize the translation requirements in the gateway.

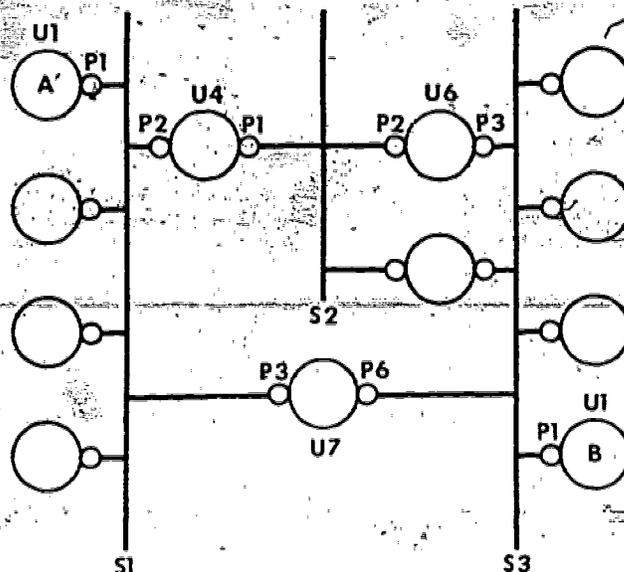
Be prepared to add security within your local network should it be connected to a global network.

NETWORK ARCHITECTURE

Distributed Computing

There was an extended discussion of process/user naming in distributed systems. Sherman described an approach in which, after initial connect, communication between processes was by means of packets in which the whole path name was concatenated at the start of the packet. As the packet progressed toward the receiver, each level in the path removed the first item in the destination field (its own name) and prepended it to the source address string. Thus when a packet reached the intended recipient the destination field would contain only the name of the destination, but the source field would contain the entire path from the sender, in correct order to be used as a destination field for a return packet. See Fig. 2.

The full pathname between the sender and receiver would be obtained by a broadcast enquiry of a directory. Duplicate directories might exist but would (hopefully) be identical. The connection between process name and pathname would be provided by the directory, and would allow access control to be maintained by the directory to offer some security.



PATH AB IS VIA U7...P3, P1, S3, U1; P6, P1, S1, U1; CONTROL; DATA
 OR VIA U4, U6...P2, P2, P1, S3, U1; P3, P1, P1, S1, U1; CONTROL; DATA

FIGURE 2. EXAMPLE-INTERCONNECTED NETWORK

Reed described a similar dynamic naming scheme, similar to a Multics pathname structure, Fig. 3. In this case each node has a process which knows the names of connected branches (nodes). These processes (directories) must be followed in sequence to find the intended destination. For example the process of name aa.bc.bc would be located by determining from process aa the location aa.bc. Each level strips off its level name and appends route information. Process aa.bc would then be required to indicate the location of process aa.bc.bc. This kind of approach can be followed to any depth. Once determined, the routing information can be used directly for further packets.

Two approaches were presented to maintain up-to-date directories.

Sherman: Grow the tree information by broadcasting from a node when it comes alive, and occasionally thereafter.

Prune the tree by discarding directory information if a node has not been heard from in longer than the broadcasting period.

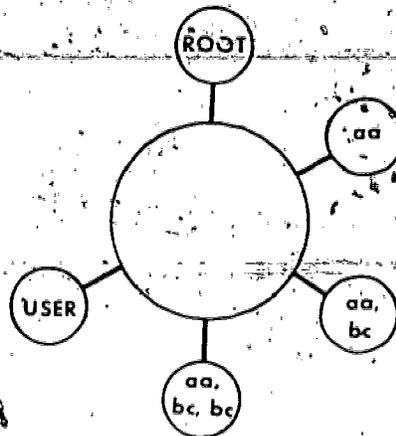
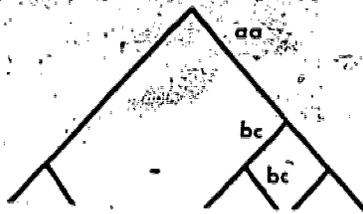


FIGURE 3. LOGICAL NAME STRUCTURE

Reed: Grow directory by each node telling its parent(s) that it is present and will be for a specified timeout period.

Prune information by deleting from directory if the timeout expires.

There was some surprise that such a distributed approach would be of interest to a person involved in industrial manufacturing automation. Sherman pointed out the manufacturing consequences of failures in inflexible systems and the desire to obtain continued operation in the event of failure of some servers.

Monitoring

Rosenthal felt that it was often important for the user, or at least the system control personnel, to know the current status of a network. This is a high-level protocol issue not generally faced. He also enquired if any participants had been able to find an important use for the "all-points" broadcast feature (with ACKs) built into many

systems. There was little response.

Types of monitoring which might be done are:

1. Statistics on busy-ness of network (network load).
2. May listen only to header information to gather traffic information.

The information obtained through monitoring may be used for operational purposes such as making high-level decisions for the user about strategies to follow for network traffic optimization. It may also be used for accounting purposes. It is clearly of interest in diagnosis of network malfunctions. It is also important in evaluation of network "tuning".

If detailed monitoring, such as full header information, percentage of packets damaged, etc., the available speed of monitoring equipment may be the limiting factor on network data rate. In the case of less detailed information gathering, the addition of monitoring will generally not require a reduction in network data rate.

Language-driven Approaches

There was a short discussion of the locality for action on commands. It was noted that there is some pressure to move this nearer the user. The concept of a Network Access Machine (NAM) which can translate between a common set of commands and those required by various servers was presented. The user-level interaction can be tailored to even the individual user, including correction of his habitual errors.

It was suggested that the incompatibility of system commands may be transitory problem, at least if the National Software Works is successful with a common network operating system.

Rogers wondered whether local networks will be run by more cooperative people than global networks. No one was very optimistic.

Zobrist emphasized that many users wanted a distributed computing system in which the user could ask for a type of service rather than for a specific machine. If in fact, the assigned machine proved inadequate, the process should automatically migrate to a suitable machine without user request or detection.

Session: NETWORK OPERATING SYSTEMS
August 21, 1977
0900 - 1230

Chairman: Stephen R. Kimbleton, National Bureau of Standards

Recorder: James P. Hanks, Mitre

Participants:

Jacques Altaber, CERN
Gary Christensen, NSC
Harry Forsdick, BBN
Jim Hanks, MITRE
Steve Kimbleton, NBS
Ron Larsen, NASA-GSFC
George McClure, FORD
David L. Mills, COMSAT
Dave Nelson, PRIME
Louis Pouzin, IRIA
Edward Rowe, NSA
Stu Wecker, DEC
Anthony West, Queen Mary College
David Wilner, UCLBL
John Wood, LBC

Dr. Kimbleton opened the proceedings by outlining a possible set of functions and objectives of Network Operating Systems as follows:

The underlying assumptions of the discussion are:

1. Heterogeneous host computers
2. Bursty transmission characteristics
3. A host operating system is inviolate

The network operating system should meet the following objectives which collectively will provide a uniform user viewpoint of the network resources:

1. Terminal support
2. Network Job Execution (somewhat related to remote job entry)
3. Network data support
4. Control

These four requirements break out into:

a. User-System interface:

1. Command language
2. File management (resources)
3. Network Job Execution

b. System-System Interface:

o Inter-process communication (IPC)

- o Remote record access (access to files or data sets at the record level thus avoiding the need to actually transfer files).

There are a set of data format mapping problems associated with preserving the logical structure of data as well as the data types. This problem is generally referred to as "data translation."

Inter-process Communication (IPC) Levels

- a. end-to-end - similar to capabilities provided by a Job Control Language
- b. call/return based - implied wait for return
- c. message based - send a message, continue until a response returns later
- d. problems of synchronization and mutual exclusion must be resolved

One mechanism looked upon favorably is a version of the UNIX "PIPE" generalized to be more independent - i.e., not just between siblings, and also to include a mechanism for mutual exclusion from resources.

Dynamic Network Reconfiguration

Another aspect of the restorability issue raised below is the need for a higher level means of configuring a network at initialization time. For example, if you are running a 10 to 12 computer network, odds are higher that failed components will exist than if you are running a 6 computer network. The issue is, how do you cope with these outages? What is desirable is a descriptive language interface that allows a network operator to define the

mapping of processes into processors and how the processes "pipe line" together in a network functional configuration.

To achieve this, a set of "schema" much like a data base mapping language might employ are required. These schema would be executed by some system "manager." The ability to dynamically "bind" and "re-bind" processes is an important reliability factor to people using networks in a real time environment.

Errors

Error control is a problem somewhat aggravated by the complexities of a network. There are several approaches to the management of error conditions.

In some cases, operational requirements place the emphasis on the prevention of error conditions. For example, in the manufacturing industry, a network failure in a parts assembly line could cause the entire production line to stop. Similarly, when instrumenting an expensive laboratory experiment, it is a disaster if all the data is not captured. In these situations there is strong motivation to minimize the potential for error. In a university time-sharing environment on the other hand, there is correspondingly less motivation for flawless operation because the users can tolerate and recover from outages.

Another aspect of the issue is how recovery from failure is effected. There are some unanswered questions, e.g., do you reinitialize everything? Do you leave failed components off until the following day in order to preserve uninterrupted although degraded service? Again, particular strategies depend on the nature of the network.

The conclusions were that errors should be dealt with in terms of system reliability and availability when designing a real time system. Reliability connotes a low failure rate for component parts of a network; but availability, a perhaps more important criterion, connotes the expectancy that all of the parts needed by a user are capable of doing the job when they are wanted. Restorability is in the more traditional sense of "mean time to restore" but the restoration strategy must be developed to meet operational requirements of a particular network.

Network Language Issues

One important need is for a language (or hierarchy of languages) suitable for use in a network environment. In a network of heterogeneous hosts, not only are data represented in different ways but differing implementations of "standard" languages produce inconsistent results.

At CERN, the use of interpreters to deal with source statements has been the mechanism for achieving some semblance of compatibility. An interpretive approach to representation of data has also been employed.

It was pointed out that a Command Language Interpreter is just that - an interpreter - and that inefficiencies are common with the interpretive approach. Even so, a command procedure language, based on assemblages of lower level functions, would be of benefit to naive users. For example, the statement "EXECUTE = Group 2 in Computer b" is easy enough to deal with (where group 2 is the function being invoked).

Although there is need for a network oriented language or family of languages, there is always the problem of getting a new language accepted by users. People are reluctant to learn a new one when they can do what they want to do with languages they already know.

Compilers/Data Structures

Current research in compilers and how they treat data structures is important to networking technology. It is necessary to deal with data in the system in independent ways for inter-host inter-operability. Traditionally, compilers have taken advantage of local conventions and not retained information describing the data structures to be dealt with by the compiled code.

Access Control - Versus Capability Based Control

Security can be provided by access control mechanisms based on access control lists or "tickets." In distributed systems, either approach presents a problem. However, local networks with high bandwidth provide a reasonable environment for a ticket-based mechanism because control information can be exchanged rapidly. Another alternative for access control is to provide an authority mechanism for naming "pipes." If an object to be operated on is viewed as a capability, a "pipe" can be invoked by a user as long as it is within his name space.

Session: ANALYSIS AND PERFORMANCE EVALUATION

Chairman: A. K. Agrawala, University of Maryland

**Recorder: W. R. Franta, University of Minnesota and
Honeywell, Systems Research Center**

Participants:

P. Mockapetras, U. C. Irvine
Mel Gable, Ford Motor/Science Lab
Jeff Yeh, E. L. L.
Ira Cotton, N. B. S.
Mike Lyle, U. C. Irvine
Bob Metcalfe, Xerox
Andy Hopper, Cambridge University
Greg Hopkins, Mitre
Stan Fralich, Comtech
Karen Gordon, University of Maryland
Gary Donnelly, N. S. A.
Victor Euscher, Lawrence Berkeley Laboratory

At the outset the attendees decided to center workshop discussion around issues germane to, first, the user's viewpoint of the network, and subsequently to the designer's viewpoint of the network.

I. User's Viewpoint

Network availability, message throughput, and message response time (including a measure of guaranteed service) were identified as the three network attributes most significant to the user. The meaning of availability and throughput are discussed in the Section II. Guaranteed service is meant to imply a guarantee that each message will be transmitted before a specified number of time units has elapsed since its presentation to the communications subnetwork.

II. Designer's Viewpoint

The areas identified as being of major concern to the network designer included:

- a) technology selection (including access protocol),
- b) transmission path medium selection,
- c) instrumentation of the communication network to measure performance,
- d) selection of performance measures, and
- e) the selection of modelling tools and model features.

a) Technology Selection

The following list of technology dependent options was constructed. Namely it was concluded that the designer must select from among:

1. random access broadcast
2. rings
3. polling
4. TDMA
5. Direct connection
6. Store and forward
7. Circuit switching
8. Shared memory

It was also observed that systems representing realizations of each alternative either exist or are under design.

b) Path Medium Selection

It was observed that most links are realized by either:

1. twisted pairs
2. coaxial TV cable, but that
3. fiber optic links are being investigated by a number of researchers (e.g., Xerox, Honeywell).

The transmission rates expected are in the 1-5 megabit per second range, but some are lower, and some (Network Systems Corporation) are rated as high as 50 megabits per second.

It was observed that as a "rule of thumb" a processor produces 1 bit of I/O per instruction executed. On the basis of this rule a machine's ability to develop (require) a trunk in the 100 Mb/s range would dictate that it operate at the 500 Mi/s level. That is fast. This observation suggests that from an efficiency point of view fiber optics is not an efficient medium, in the sense defined in the ETHERNET paper (CACM, 19, 7, July, 1976).

6. Following is a side discussion that developed concerning the placement of protocol (e.g., access, collision detection, realization of retransmission policy) related matters: specifically, what and how much of this machinery should be cast in the hardware, and how much should be relegated to software in the connected host. One approach (Xerox) has been to place as much as possible in software to minimize hardware cost, while another (Network Systems Corporation) is to remove most functions from host software and handle them in Bus Interface Unit firmware. No consensus of opinion evolved, although the attendees seemed to favor

2. Response time - elapsed time from generation of a message/packet until its successful transmission. We are interested at least in the mean, variance, minimum and maximum.

3. Queuing delays - including those associated with host-host and communication subnetwork delays.

4. Utilization - fraction of bandwidth actually used

5. Reliability

6. Buffering requirements at intermediate unit (if any)

Processing requirements at intermediate unit (in b).

7. Availability - the system should be available as possible. In finding M.T.B.F. it was suggested that best measure may be probability of finishing what I start) and is obviously related to the individual user's service requirement.

Recovery time following failure and degree of recovery from user's view.

Fairness - in allocating

Cost

8. Error

the mean

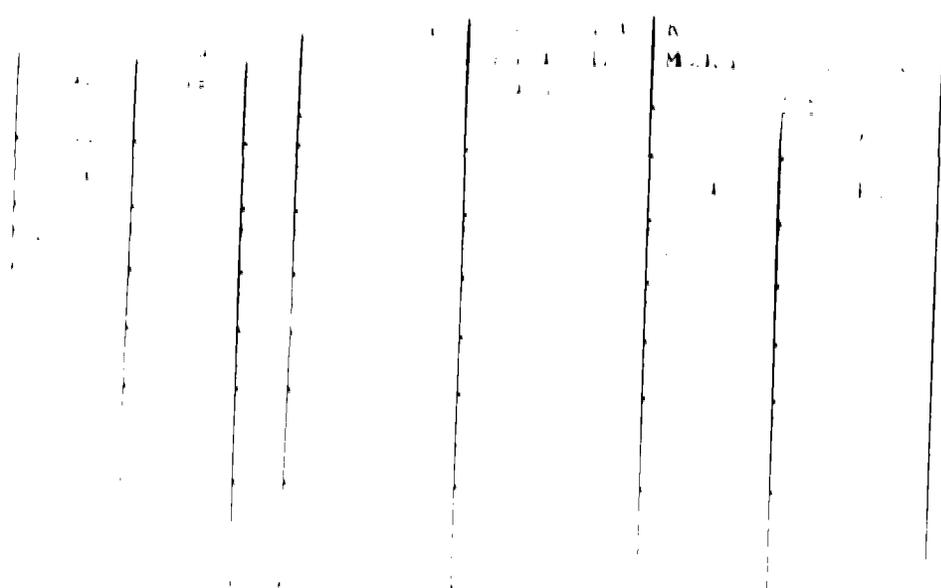
It was de

dependent on average
although there as
way performance
two additional
ensure that
not clear whether
communication
Some
and
Additional

represented a crucial element of the model, and then that the following traffic models should be considered:

- 1 Independent Poisson process traffic
- 2 Deterministic (periodic) traffic
- 3 Bi-modal traffic, representing the analysis of long and short traffic requests.
- 4 Dominant user traffic, perhaps including the distribution of message destinations for messages emanating from the dominant user.
- 5 Application dependent traffic (e.g. voice traffic is a mix of Poisson and periodic).
- 6 Real traffic resulting from data collected from instrumentation of an existing network.
- 7 Both finite and infinite population models.

A. ...
B. ...
C. ...
D. ...
E. ...



The session concluded with several discussions of the (at least in some instances) designer's quest for novelty and the effect of protocol on efficiency. It was noted that the degree of novelty in a design is:

1. Inversely proportional to the confidence in the place in the resulting system
2. Limited by available manpower

It was also noted that the degree of novelty is inversely proportional to the confidence in the place in the resulting system (that efficiency is dictated by the static (not resilient) load flow, rate (in pockets) to describe the routing) in a situation dependent

BIBLIOGRAPHY ON LOCAL AREA COMPUTER NETWORKING

Communications " 11 Int. Computer Soc. J. 1970,
pp. 695-702

No.
11

Mar. 1970

...

...

SOEWA
76 pp. 32-328.

Computer-Communications and Teletraffic, New York:
Polytechnic Press, 1972; pp. 539-545.

PAID Parler, David J. "A Ring Network." Datamation, Vol
21, No. 2, February 1975 pp. 44-40.

... distributed computer
system Comp on 3, 68, 1973, pp. 31-34

... State ... James
"Technological Consideration for P. ...
Networks" National Computer Conference May 1975 pp
233-244.

... and ...
Broadcasting National Computer Conference
May 1975 pp. 253-260.

...
1977

... 1973

...
1976

...

...

HAY71a Hayes, J. F. and D. N. Sherman "Traffic and Delay in a Circular Data Network." Second Symposium on Problems in the Optimization of Data Communications Systems, October 1971 pp 102-107

HAY71b Hayes, J. F. and D. N. Sherman "Traffic and Delay in a Circular Data Network." Analysis of Data Communications Systems, The Bell System Technical Journal, November 1971, pp 2947-2970

100
National

100
100
100
100
100
100
100

Thesis, Harvard University).

ME 70 Distributed Packet Switching for Local Computer Networks. Communications of the ACM, Vol 19, No 7, July 1967, pp. 35-40.

IP.P.C. The Design of a Packet Switching System. August 1971, p. 27.

June 1972, p. 750. 3,

April 1973, p. 4.

Record of the Proceedings of the ACM National Conference, January 1975, p. 12.

589 596.

Report on the design of a computer communication network
Department of Computer Science (University of London)
unpublished report February 1977

3.7.7 19 2 1 539

19
11

NBS Workshop on Local Area Computer Networking

August 24 and 25, 1977

Billings
Wash. D.C.

Computer Science
College Park, Md. 20742

ERN
P.O. Box
01031
Neden

Billings
Wash. D.C.

Forsdick, Harry C.
Bolt, Beranek & Newman, Inc.
50 Moulton Street
Cambridge, MA 02130

617/491-1850, x638

Comtech Adv Sys
754 N P St
San Jose, CA 94108

Univer. of D.
143 Sp...
100 Union St
Minneapolis, MN 55454

PG...
PG...
PG...
PG... MA

Univer. of D.
PG...
PG... MA

PG...
John Hopkins
PG... MD 20184

PG...
PG...
PG...

PG...
PG...

PG...
PG...
PG...



Kimbleton, Stephen R
National Bureau of Standards
Bldg. 225, Rm. B214
Washington, D. C. 20234

301 921-2001

NASA
NASA
Code 5701
Greenbelt, Md. 20771

18

University of
Computer Science
Irvine, Calif.

State University
P.O. Box 2
Irvine, Calif.

Dr. Albert M.
State University
Irvine, Calif.



Pilipchuk, Andrew
University of Maryland
Computer Science Center
College Park, MD 20742

301/454-4526

MIT Laboratory 10.
545 Technology Sq
Cambridge, MA 213

Institut de Recherche
et d'Automatique
Jean Diez de Voltaire
Rocquencourt
France

411
545
211

100
100
100
100

100
100
100

100
100
100

100
100
100

100
100
100

Wecker, Stuart
Digital Equipm
ML3-2/E-1
146 Main St
Maynard, MA

017 537 0111 x4300

Queen Mary Col
Computer Systems
Mile and Rd
London E 4

London, UK
Real Time Sys
#10 Chichester
Berkeley, CA 947

London, UK
1 Cyclot
Bldg 50
Berkeley, CA 947

Lawrence
Computations
P O Box 80
Livermore

Lawrence
A 123
MSB01

Lawrence
E11
1670
P O Box 1
Livermore