C CE & TECHNOLOGY:

LOCAL AREA NETWORKING

NBS Special Publication 500-31

U.S. DEPARTMENT OF COMMERCE
National Bureau of Standards

COMPUTER SCIENCE & TECHNOLOGY:

Local Area Networking

Report of a 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

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 [001.6'44'04] 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).

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

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 COMPARATIVE EVALUATION OF THE PERFORMANCE OF ALTERNATIVE COMMUNICATIONS TECHNOLOGIES

A. K. Agrawala
Department of Computer Science
University of Maryland

Several communications technologies lend themselves well to use in a local area computer network. Techniques such as rings, random access cablebus, message switching, circuit switching and shared memory have been proposed, reported in the literature and/or are at various stages of development. In selecting a technology for a particular system, the designer is faced with the problem of evaluating the relative performance of these approaches. Performance results available in the literature are usually based on assumptions which are different for each approach, making a comparative evaluation extremely difficult.

In an ongoing study, being conducted at the University of Maryland, we are evaluating the performance of several communications technologies under the same loading conditions. To date, we have evaluated the performance of the Common Data Buffer, Ethernet, ring, and message switching technologies for a set of 8, 16, and 32 interconnected hosts. The communication load consists of a sequence of messages of random length generated by a Poisson process at each host. The host-to-host communication performance for each alternative is then studied using this workload model.

These models have been constructed using SIMSCRIPT II.5 and were studied from light loading through saturation conditions. The results of these studies will be presented.

REAL-TIME NETWORK FOR THE CONTROL OF A VERY LARGE MACHINE

J. Altaber
European Organization for Nuclear Research
Geneva, Switzerland

In 1971, the European Organization for Nuclear Research (CERN) undertook the construction of a Super Proton Synchrotron (SPS) which would accelerate protons to an energy of 400 GeV. The control requirements which are distributed all around the machine involve several tens of thousands of bits of status information and many thousands of analogue measurements and digital controls. The overall control of the machine is made from a Main Control Room by a small operating team. The geographical layout of the machine, which consists of 2.2 km diameter ring, and the necessity to reduce the number of control cables and their length, have led to a decentralized structure for the control system which has been finalized into a star network of 24 computers.

The design constraints of the control network were the following:

- facilities of embedding an urgent message (such as an alarm) into the transfer of any other message. This implies that the centre of the star is controlled by a computer known as the Message Handling Computer (MHC) which works in a store and forward mode;

- to obtain a satisfactory response time for operator actions on remotely accessed equipment, a 32 word-length message has to be transferred from source to destination with 5 ms elapsed time;

- anticipated traffic flow shows that the size of the messages can reach 32 Kwords with a traffic load of 30 Kwords/s.

A serial asynchronous data transmission system was designed to link satellite computers to the MHC. The transmission speed is equivalent to 30 us per 16-bit word. Each inter-computer link consists of four separate uni-directional channels with one in each direction for data and one in each direction for traffic control informataion. Messages are transferred by DMA for data channel line and by PIO for transfer control word (TCW) on the control channel line -- the TCW controlling and handshaking the data channel.

A point-to-point protocol minimizing the number of messages exchanged has been developed. This protocol, associated with a simple transport station procedure, provides asynchronous access to all resources in the network, messages being sent from source to destination without any special prior agreement.

The system has been in operation since the middle of 1975 and has been used with complete success for the commissioning and operation of the SPS machine. It has been adopted for controlling other machines or experiments on the CERN site and the current development is to connect all the individual control networks together with the possibility of inter-network transactions.

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

A local network is being investigated for use at the National Bureau of Standards' Gaithersburg site. The user community has been identified to include two major computers, about 60 smaller computers, and about 300 to 500 user terminals. Seldom, if ever, would most of the attached devices be simultaneously active. Full connectability is required. The two parties to a connection must not be required to operate at the same baud rate, and flow control must be provided. A broadcast packet network has been chosen for a number of reasons: incremental implementation, reliability, use of existing rf cables, and suitability for bursty users.

This system differs from other broadcast networks such as ETHERNET in that most of the user terminals will be "dumb," and an inexpensive Terminal Interface Equipment (TIE) will have to be provided for each user. Initially asynchronous serial RS-232 has been chosen as a lowest common denominator for all users. Most of the minicomputers at NBS have a spare port of this type. User data rates up to 19,200 bits per second are expected. The standard TIE, four of which exist in breadboard form, is based on an 8-bit NMOS microprocessor, with hardwired logic handling the transfer of data between the 1 megabaud network and multiport buffers.

A maximum packet length of 128 bytes has been chosen. Each packet contains both destination and source addresses; a command field containing control, acknowledgement, and sequence information; a byte count field; an optional data field; and a 16-bit CRC. Both addresses are 12 bits in length to accommodate the large number of potential users. The addressing scheme must be able to cover both our Gaithersburg and Boulder laboratories, should systems at the two sites ever be connected.

Plans for future improvements include TIEs suitable for two to four low-speed user terminals, TIEs to multiplex multiple connections into large hosts, and possibly wider-band TIEs for file transfers to minicomputers.

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.

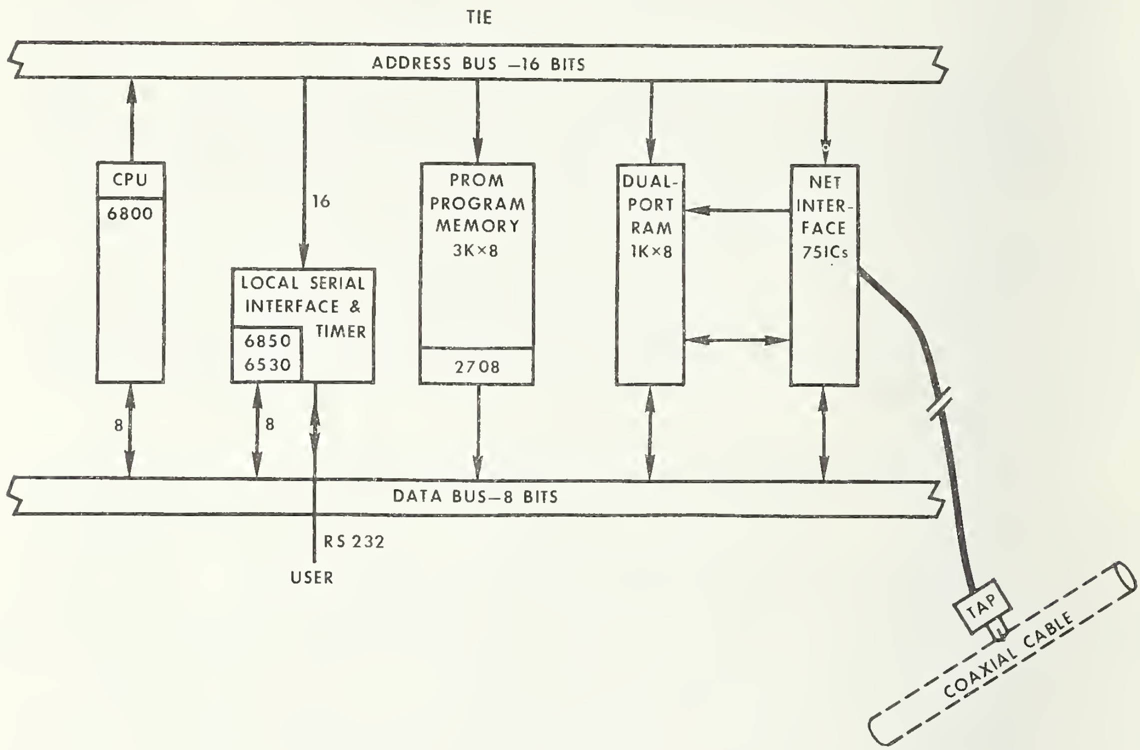

TYPICAL PACKET

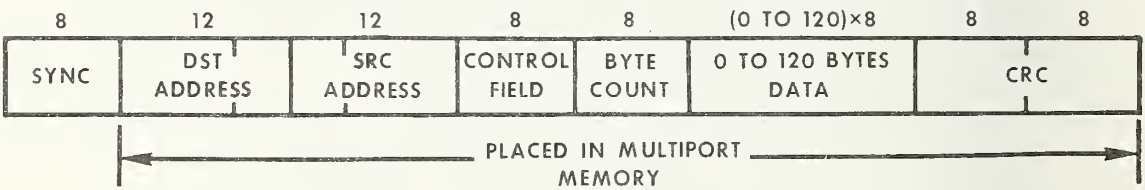

MINIMUM SIZE 8 BYTES, 64 BITS

MAXIMUM SIZE 128 BYTES, 1024 BITS

Estimated
1986 DATA COMMUNICATIONS REQUIREMENTS

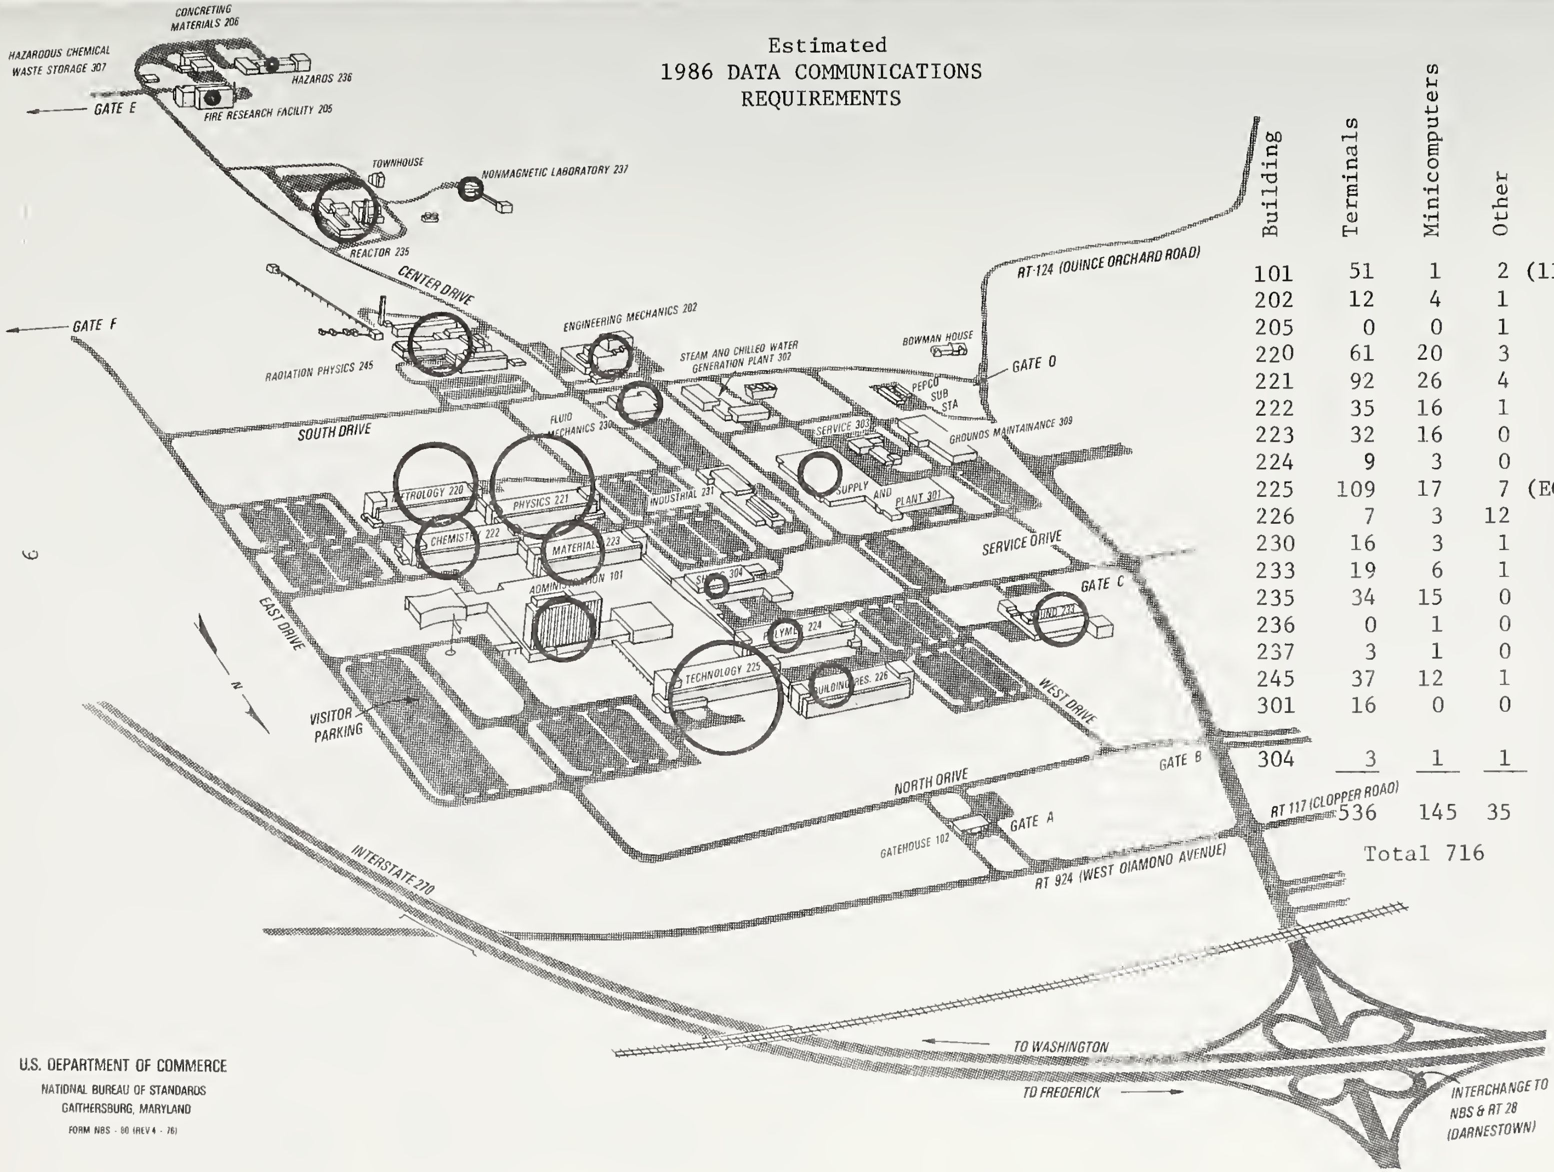

Building	Terminals	Minicomputers	Other	
101	51	1	2	(1108)
202	12	4	1	
205	0	0	1	
220	61	20	3	
221	92	26	4	
222	35	16	1	
223	32	16	0	
224	9	3	0	
225	109	17	7	(ECF)
226	7	3	12	
230	16	3	1	
233	19	6	1	
235	34	15	0	
236	0	1	0	
237	3	1	0	
245	37	12	1	
301	16	0	0	
304	3	1	1	
	536	145	35	

Total 716

APPLICATIONS OF HYPERCHANNEL (TM)

Gary S. Christensen
Network Systems Corporation
Brooklyn Center, Minnesota

Presents examples of actual and planned applications of the HYPERCHANNEL local computer network. These applictions include:

- linking of main frames and frontend processors,
- operation of peripheral subsystems on the network.

Also discussed will be the purpose of the example applications--vendor independence, application of new technologies, improved availability, common data bases, and common access networks.

MITRENET -- INTRODUCTION AND OVERVIEW

James P. Hanks
The Mitre Corporation
Bedford, Massachusetts

For the past two years, MITRE's Bedford Division has been designing and building an internal computer-to-computer network based on digital communications over coaxial cable. The design is intended to accommodate heterogenous computers and provide access to their resources through a network command language user interface. Network control is fully distributed, each host node having an associated copy of a Network Operating System. The prototype system, including IBM 370 Host and Data General NOVA 800 Host, was demonstrated in June, and is currently being retrofitted to use a different cable communication technique which will provide better use of available cable bandwidth. Plans for future expansion and operational use are in progress.

In our work to develop an operational local computer network, many of the issues common to large networking activities have been faced and solved with varying degrees of success.

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 communications path between computers and other devices in the Computer Laboratory. These devices are connected through the ring on an individual basis and as yet there are no global high level protocols to provide automatic call establishment. The primary uses of the ring are for equipment sharing and file dumping.

Ring Organization

The original design was based on the register insertion principle where the packet to be transmitted is placed in a shift register which is inserted into the ring at the appropriate moment in time. As the delay in inserting a register and thus transmitting is at most one packet time, hogging does not occur and bandwidth is distributed to all nodes symetrically.

In due course it was realised that a more attractive system would be one based on the empty slot principle. In its simple form the empty slot system suffers from hogging. This defect can be overcome if each packet makes a complete revolution of the ring and is not marked empty until it has passed the original source. With this scheme the interaction with the ring at each node is minimised and reliability is improved. As performance characteristics of the two systems are very similar the empty slot system was adopted as the basis for the Cambridge ring.

The structure of the ring is shown in Figure 1. Repeaters are used to regenerate the signal at each node and can operate autonomously from the stations which perform the logic functions for transmission and reception of packets. Each station is interfaced to its host via a specially built access box. The access box tends to be sophisticated for a simple device such as a line printer and fairly simple for an intelligent device such as a computer which can perform most of the required logic functions internally. There is a unique station called the monitor station which is used for setting up the slot structure during turn on, for monitoring the ring and clearing lost packets, and for accumulating error statistics.

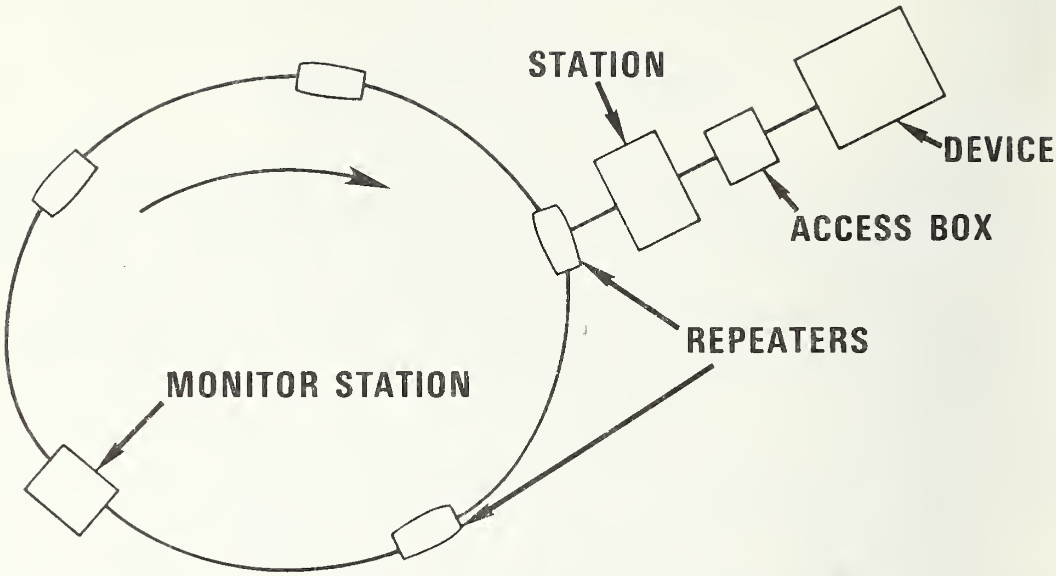

FIGURE 1. RING STRUCTURE

The packet structure is shown in Figure 2 and is chosen to allow the maximum timing tolerance and minimum delay at the transmitter and receiver. The leading bit is always a one and is followed by a bit to indicate whether the slot is full or empty. Now follows a control bit used by the monitor station to mark as empty packets which are circulating indefinitely due to an error in the full/empty bit. This is followed by 4 eight bit bytes the first two of which are used for destination and source addresses and the last two for data. Finally there are two control bits used for acknowledgement purposes.

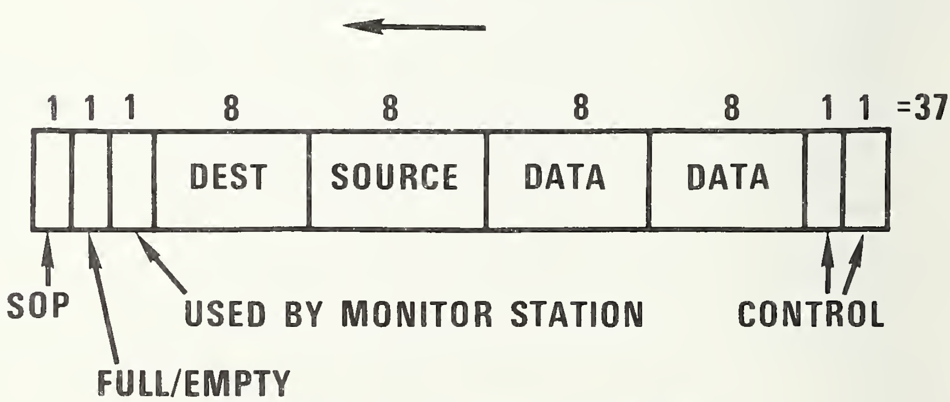

FIGURE 2. PACKET FORMAT

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 a 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 node.

The transmitted packet makes its way to the destination where the control bits are set on the fly to indicate accepted, busy, or rejected. It now returns to the source where the slot is marked empty. If the packet returns with the control bits unchanged it was not recognised by any destination. Each station automatically computes the total number of slots in the ring and can thus clear the full/empty bit immediately.

It can thus be seen that on transmission the packet is delayed until an empty slot is found but then the transmission is rapid. This is in contrast to the register insertion scheme where the delay round the ring can be large but the initial delay before the register is inserted is small. As the destination does not explicitly signal when it is ready to receive the next packet the ring can easily become clogged with the packets returning marked busy when devices with varying speed characteristics are being interconnected. In order to overcome this the following algorithm is incorporated in the hardware to reduce the number of busies. If a source transmits the same packet twice and both times it returns marked busy then it is not allowed to retransmit it again until some time later. This additional delay is dependent on ring loading and is the time to acquire the next empty slot. If any further retransmissions are attempted the extra delay is increased to about 16 x ring delay x traffic density. Thus the number of busies is decreased and performance is improved.

Each station possesses a station select register which is initialised by the host. This register can be set to accept or to reject all packets, or to receive from one source only. When combined with a time out mechanism it can be used to allocate resources on the ring.

Error Recovery

There are no CRC or parity checks on the transmitted packet; however, a copy of the information is retained at the source and is compared with the returning packet. This provides a powerful error detection facility but does not indicate that the packet was correctly copied at the 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 source. An error in the response bits might have a more serious effect as it will not be detected by the transmitter, which might repeat the packet or assume it was received correctly when this was not the case. Under some circumstances such errors can propagate but generally they are detected by the source or monitor station within one ring delay. Unnoticed errors are rare.

On turn on and if the number of slots in the ring changes each station determines the slot count by using the gap digits. The gap digits are set to zero and at least one such digit must be present in the system.

Additional error detection facilities are provided by the monitor station which can issue test packets, store erroneous ones and provoke a response from any station. This response is independent of the returning data so that such a test can be carried out when the ring is broken.

Hardware

The ring is built using TTL technology and operates at 10 MHz with a maximum distance of 200 meters between repeaters. Higher rates would be readily attainable with faster logic. The signals are transmitted along twisted pairs of the type normally used for the duplex operation of teletypes. Transformers are used throughout for isolation and common mode rejection. As the repeaters have to operate reliably whether they are connected to a station or not they are powered directly from the ring. This power is injected into the system at the monitor station.

Each station is full duplex so that it can transmit, and receive, concurrently and independently. The number of bits delay at a station is a fraction of a bit and the minimum ring delay is about 5 microseconds.

A number of modulation techniques were considered and some of these are shown in Figure 3. The four wire scheme was chosen as it is suitable for a pair of twisted wires and 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:

- it is d.c. balanced
- there is very little encoding or decoding delay
- both pairs are treated identically thus minimizing skew
- a choice of clock techniques is available
- it is easy to provide a control signal for a PLL
- it is easy to detect some errors (e.g., no change on either pair)
- it cannot move 1/2 digit out of phase

A number of repeaters and stations have been built and these are currently undergoing trials.

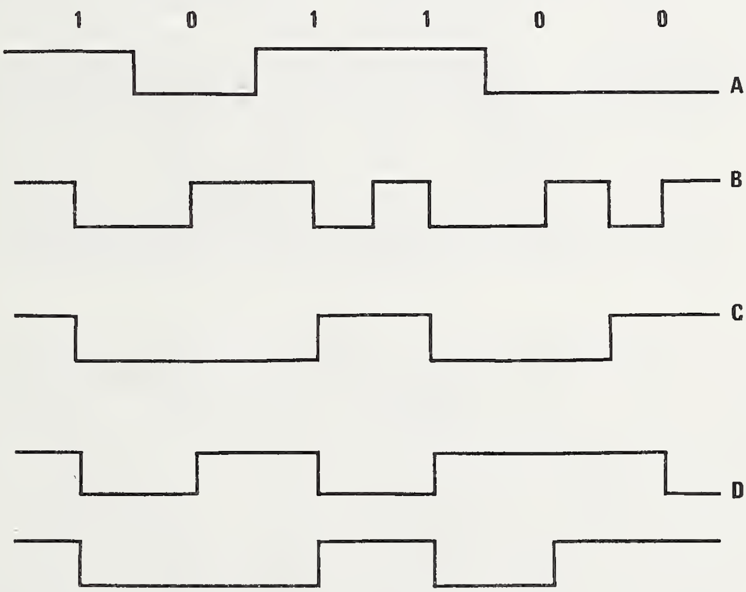

A. PULSE FOR 1 NO PULSE FOR 0
B. PHASE MODULATION
C. P.M./2
D. FOUR WIRE SYSTEM

FIGURE 3. MODULATION TECHNIQUES

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 an inexpensive as possible and it was thus kept simple. Nevertheless, many options are left and some of these are discussed below.

In a simple scheme each source transmits to only one destination at a time. This can be extended to allow the multiplexing of packets to different destinations. If this is done an algorithm has to be developed for matching the speed of transmission and reception and to ensure no sources are continually blocked. For devices with similar speed characteristics this can easily be done by employing a round robin scheme. Where the communicating devices operate at different data rates a speed number could be associated with each destination. This number is updated when the delay does not match the estimate. Other algorithms have been developed which achieve this in different ways.

Under some circumstances the services of a particular node might be required by a number of stations at the same time. In the Cambridge system such requests are arbitrated on a random basis and thus some sources experience additional delay before successfully transmitting. Where the application demands more precise performance characteristics such requests should be queued.

Acknowledgement

The ring was designed with Professor M. V. Wilkes and Professor D. J. Wheeler whom the author had help in preparing this document.

LOCAL MISSION-ORIENTED NETWORK

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

During the past two years, we have been investigating 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 components, as well as work on system architecture and configuration trade-offs.

As part of this effort, research has been sponsored at the University of Maryland, specifically in the area of communications technology for local computer networks. The principal methodology employed has been simulation, although analytic formulations have been used where appropriate. Specific communication alternatives have been identified (e.g., cable-bus, ring, shared memory, irregular topologies, etc.) and have been analyzed for both standard traffic loading assumptions (Poisson message arrival, etc.) and specific traffic loads representative of actual mission operations. Efforts have been focussed on producing results which enable comparison of the available technologies, and, hence, are usable in configuration trade-off studies.

INTERCONNECTION OF LOCAL NETWORKS USING SATELLITE BROADCAST TECHNOLOGY

David L. Mills
Communications Satellite Corporation
Washington, D. C.

The broadcast nature of the communication satellite is uniquely well suited for the interconnection of a number of geographically dispersed local networks. Using this medium, burst rates comparable to digitized voice (64 kilobaud) can be achieved with complete connectivity using packet protocols similar to those now used by typical local-network transceivers. In spite of the relatively long propagation times (up to 1/2 second round-trip delay) it is possible to achieve a performance equivalent to a floppy disk drive, as viewed by a microprocessor attached to the local network. This short report will discuss some of these issues and suggest some possible applications.

ETHERNET: DISTRIBUTED PACKET SWITCHING FOR LOCAL COMPUTER NETWORKS

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

Ethernet is a branching broadcast communication system for 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 computer networks or loosely coupled multiprocessors. An Ethernet's shared communication facility, its Ether, is a passive broadcast medium with no central control. Coordination of access to the Ether for packet broadcast is distributed among the contending transmitting stations using controlled statistical arbitration. Switching of packets to their destinations on the Ether is distributed among the receiving stations using packet address recognition. Design principles and implementation are described, based on experience with an operating Ethernet of 100 nodes along a kilometer of coaxial cable. A model for estimating performance under heavy loads and a packet protocol for error controlled communication are included for completeness.

THE ARPA LOCAL NETWORK INTERFACE

Paul V. Mockapetris
Michael R. Lyle
University of California
Irvine, California

David J. Farber*
University of Delaware
Department of Electrical Engineering
Newark, Delaware

The LNI, local network interface, project had, as its goal, the design and development of a low cost, fail-safe, and flexible local communications system centered about the design of a single chip LSI transmission controller. The background for this effort included the early (1972) development of a prototype ring system supporting a fully distributed operating system (DCS). This early system has

----
* Principal Investigator

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 addresing 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 shoud 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 that 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 El 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 [Metcalfe 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 LSI-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, <u>Communications of the ACM</u>, 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 environement 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 Messaage 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 similiar 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 comunication 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 charateristics 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 ether 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 funtions 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.

Examples were given of using the center conductor of a coaxial cable for transmitting power for powering electronic devices along the cable, such as repeaters.

The suggestion was offered that carriers can be imposed on power lines within a facility and that signals in the range of 100 kHz to 1 MHz can be used for this purpose. The best use of this was felt to be for narrow band applications within a single building.

The MITRE two-cable system was cited as having very low loss, less than 2 dB, with 200 drops, within a two-mile range. This system uses one-way transmission and amplifiers on each cable with a head-end amplifier bridging the cables. The relative merits of baseband versus RF techniques were discussed and it was pointed out that cable equalization can be used to compensate for frequency-dependent characteristics. However, this is mainly limited to one-way fixed configurations, and is less applicable where path lengths are arbitrary. The use of repeaters was offered as a means of overcoming cable losses, though it was felt that they could have an adverse effect on reliability. The comment was made that if a sufficient investment is made in modems, satisfactory operation cn be achieved in spite of enormous losses, in excess of 100 dB.

Tapping of coax was discussed, with a description of the use of pressure taps. These were felt to be an improvement over the use of tee's. An example was given in which over 150 taps were employed with low loss and minimal reflections. In this case, the stubs were short and the electronics attached to the stub were placed very close. Other examples were given of the use of considerably longer stubs and higher losses.

The emergence of optical fibers was discussed and it was agreed that these held potential as a transmission medium for local area networks. At present, they are the subject of intensive development. Taps appear to be a problem, though there are couplers becoming available. Modems are presently in the $1,000 range, though it was stated that the parts cost is perhaps only $60. With volume production and use, this technology should become quite inexpensive.

Radio systems were discussed briefly, but it was noted that there are severe limitations on available bandwidth and that FCC regulations were a further obstacle. The Aloha modems were cited at a cost of about $1,000.

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 decesion. An example was given of using CRC hardware for receiving, because of the large volume of received data, but doing the CRC in software for transmission, since much less data orginated at the station.

The use of "floating" names was discussed, in lieu of fixed hardware addresses. An example was given of process-to-process addressing. A hardware implementation of a Name Table was described; there was a divergence of opinion as to whether the Name Table should be in hardware or software.

The following diagram was provided by a participant to illustrate the hierarchical distribution of functions:

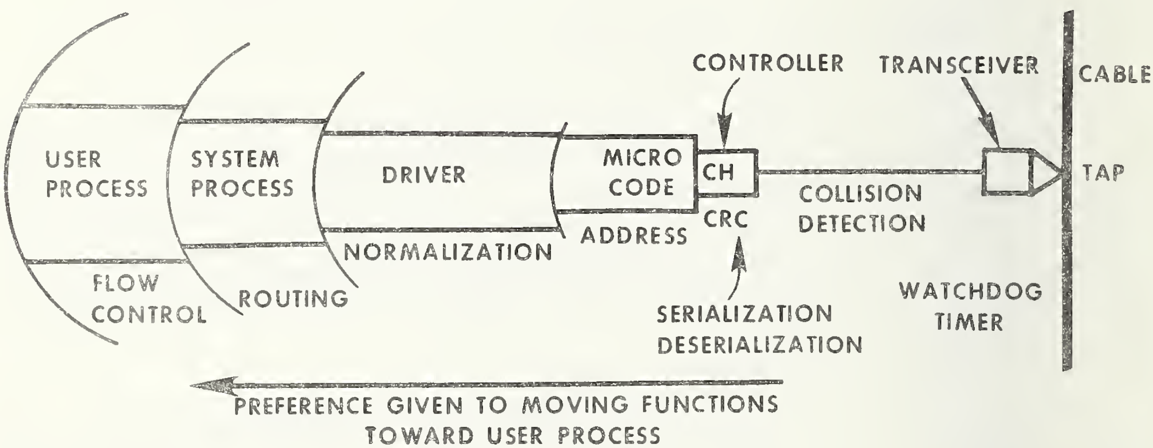

The design objective was stated as being to move functions toward the user process in order to make maximum use of the computer (in contrast to other design philosophies).

Another participant's approach is shown below:

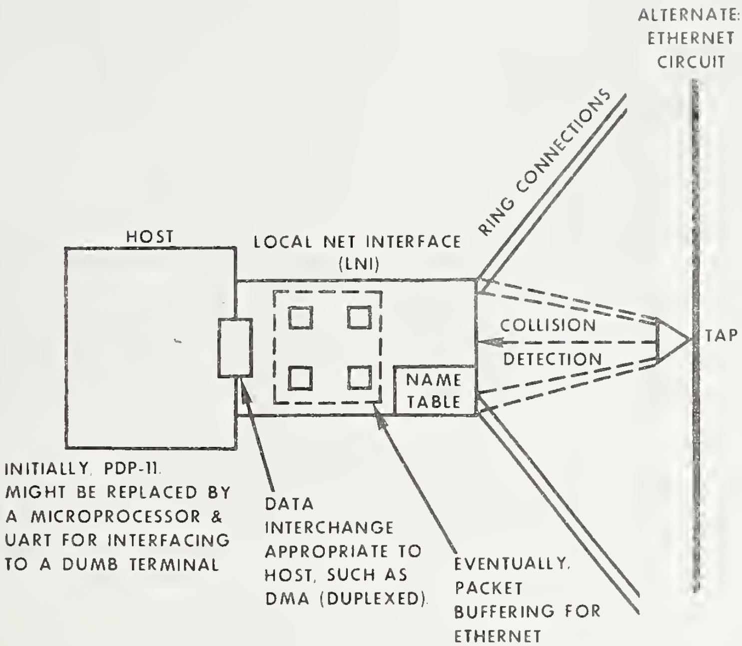

The comment was offered that slotted systems are very expensive compared to unslotted systems.

A system was described in which collisions could still be detected even with losses of up to 25 dB through the use of a sophisticated encoding scheme which has good correlation properties.

The issue of virtual circuits versus datagrams was debated. The observation was made that people concerned with end-to-end performance preferred datagrams, while Post, Telephone and Telegraph (PTT) people preferred virtual circuits. It was agreed that the choice should be an engineering decision based on the requirements of the application. Some users may prefer datagrams because they can get a little higher throughput, provided they have ways of dealing with anomalies such as packets out or order. The view was expressed that it is easier to build a virtual circuit system on top of a datagram system than vice versa.

A question was raised as to why there should be so much concern for this issue. It was answered by noting that the carriers are saying they must settle upon one system or the other and proceed with its implementation. The concept embodied in X.25 is that of virtual circuits.

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 in the Packet Radio Network was cited as a means of determining circuit loading.

Comments from the Chairman

The session was a lively one and exchanges of views flowed freely and frankly. Many of the participants had been or were about to be prime movers in their respective implementations, and that insured a competent level of discussion. The chart on the next two pages summarizes the technical characteristics of several systems represented at the workshop.

Some of the most important issues raised, it seems to me, are the following:

1. There is a need for some good backroom work in the engineering tradeoffs in cable systems, taps and signal design. I sense that the performance of some of the systems could be improved through a critical examination of the signal design (correlation properties, d.c. shift, efficiency, etc.) and cable plant (taps, hybrids, amplifiers, etc.). The CATV industry has developed a good deal of technology for this, which, if exploited, could allow expansion of the local-area system via common carrier, CATV and satellite video channels.

2. The use of carrier-current technology, which has been lying around college radio stations for years, is an inviting and interesting medium for relatively low bandwidth ALOHA systems within a building or dense cluster of buildings. Very cheap narrow-band FM transceivers can be made for data transmission at rates to 9.6 kilobaud or so, which is the rate used in the ALOHA system. A little shoebox using the ALOHA technology at a 180 kHz carrier frequency, for example, could provide truly portable communication sans telephone line within a fair size cluster of buildings, depending upon the nature of their power distribution system.

3. It was pointed out that contention-bus and ring systems have more in common than not. This is a very useful point of view, allowing meaningful comparisons in analytic and experimental models. Hopefully, the rapidly escalating interest in local-area networks will spur research in queuing models of these systems, which should lead to better flow-control and error- control algorithms. Stability

control in ALOHA/Ethernet systems is a prime target for these activities.

4. I sense two not necessarily mutually-exclusive views of local-area network utility. One group sees them as the ultimate in channel-to-channel adapters, while the other sees them in a terminal-to-host management role. The terminal-to-host viewpoint obviously is encouraged by a bullish view of the economics of these things and requires a smart cheap shoebox, probably with its own microprocessor and packet buffers. The channel-to- channel viewpoint sees the interface design much more tightly coupled to a host processor, using its microprogrammed capability to implement sophisticated protocols and its DMA facility to access packet buffers in the host memory. I tend more to the bullish view and embrace designs allowing operation disjoint from any smart host, with host or terminal interfacing on either a parallel or serial basis, depending on the data rates involved. Whether the interface connects to a programmed or autonomous bus would, it seems to me, be a matter of host system requirements and capabilities.

5. The issue of functionality in the hardware/software/firmware is a matter of some practical interest. It is rather disingeneous to point out that one man's fancy hardware controller is another man's firmware program, but the issues are the same. My particular horse in this derby is flexibility in methodology-implementation in hardware, firmware or software is not in itself important, but only the selection of one or another of them as determined by a comprehensive engineering study of the system as a whole.

	Mitrenet	Ethernet	MIT-LCS	Ford
Collision Detection	interface hardware	interface hardware	interface hardware	interface hardware
Flow Control	end-end	μprog(9)	prog	prog
Error Control	end-end block	CRC/ARQ end-end	CRC prog(1)	prog
Buffering	internal + host	internal + host	host DMA(2)	internal 2K RAM
External Interface	host/term.	host	host	serial 19.2K + parallel
Structure	head-end	bus	ring or bus	bus
Protocol	slot/ALOHA	contention	contention	contention
Rate	.3-7 Mb	3 Mb	1-8 Mb	1.3 Mb
Code	RF	Manchester	baseband	ternary
Primary Use	host + terminal	host/host	host/host	host/proc. control
Construction	μp	Nova 800, Alto μp	PDP-11 bus periph.(4)	multi- μp (SMS + PM11)
Cost ($)	1k±	.5K	2K(3)	1K
Status	private	'74	Sept. '77	private
Taps	passive + repeater	passive T	active repeater	passive (lossy)

1. ether version only
2. packet buffers to be included in ether version
3. subsequent versions lower cost
4. subsequent versions M6800, PDP-10, etc.
5. includes cost of μp system
6. may convert to Z80
7. may construct UNIBUS system
8. 4 nodes now
9. end-end flow control; distributed congestion control
10. G.P. O.S. available also

NBS	CNET	ENET	DCS	Cambridge
interface hardware	μp software	interface hardware	no collisions	no collisions
prog	μp software	prog	hardware name table	hardware
CRC/ARQ μprog	μp software	CRC prog. chip	CRC + soft. seq.	packet compare
internal 2R x 4T	internal 4K RAM	4K RAM + host	host DMA	internal DMA
9.6 K	serial 19.2K + parallel	serial 19.2K +parallel+DMA	host UNIBUS	host
bus	bus	bus	ring	ring
contention	contention	contention	marker	marker
1-2 Mb	<.1 Mb	3-5 Mb	2.3 Mb	10 Mb
Manchester	baseband	phase mod.	baseband	baseband
terminal/ mini/host	host/host I/O	host/host	process/ process	host/host I/O
M6800(6)	M6800(6)	M6800(6)	hardware FSM	TTL custom
.5K-2.5K	<.2K(5)	.5K(5)	1K	<.3K
'76(8)	Oct. '77	Dec. '77	'73(10)	'77
passive T	passive T	passive T	active repeater	active repeater

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

Chairman: Stuart Wecker, Digital Equipment Corporation

Recorder: Robert Rosenthal, National Bureau of Standards

Participants:

Jacques Altaber	CERN
Gary Christensen	Network Systems Corp.
David Nelson	Prime Computer, Inc.
David Reed	MIT/LCS
Edward Rowe	National Security Agency
Richard Sherman	Ford Motor Company
David Wilner	Lawrence Berkeley Laboratory
John Wood	Lawrence Berkeley Laboratory
Jeff Yeh	Lawrence Livermore Laboratory

PROTOCOL ISSUES IN LOCAL AREA NETWORKS

An agenda consisting of four major topics was proposed as a basis for discussion of protocol issues in local area networks. The agenda consisted of the following:

I. Layered Structures and Protocol Functions

II. Protocol Issues at Each Layer

III. Internetworking
- local vs. local
- local vs. global

IV. Standards

A brief dichotomy of protocol layering provided a "common ground" and framework on which to proceed. Consensus was easily reached that the following layers of protocol are representative of layered structures and protocol functions in most networks today:

1. Hardware Level (port)
   - modem control

2. Link Level Protocol (data link control)
   - framing
   - bit error detection/correction
   - link management and control

3. Network Level Protocol
- routing
- congestion control
- virtual circuit creation and management
- end-to-end acknowledgement
- flow control
- sequencing

4a. Application Level Protocol
- I/O device control/access
- file access

4b. User Level Protocol

User level protocol issues, thought to be extremely important, were only briefly discussed; some members expressed hopes that additional time could be devoted to broader coverage -- another session devoted to this issue was suggested.

With this layered structure and these protocol functions in mind, the group proceeded to tackle the protocol issues peculiar to local area networking. Discussion centered around functional requirements of the protocol; the following issues were addressed as they relate to local area networks:

- error freeness of data links and end-to-end recovery
- message sequencing
- flow control
- security
- response requirements and user interface
  - single message-at-a-time service
  - no pipe-lining
  - connection establishment procedures
  - data and message priorities
- guaranteed delivery
- connection vs. telegram vs. broadcast techniques
- end-to-end addressing
- allocation of resources
- interrupt facilities
- failure and recovery techniques

After this discussion of requirements, the group began to question the attributes of a local area network. Hopefully, the previous discussions of protocol requirements and protocol layering together with a discussion of local network attributes would shed light on our understanding of protocol issues for local networks. The following attributes were discussed and agreement was reached that the following were attributes of local area networks:

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

2. low delay (messages do not remain in the network very long)

3. relatively error free (compared with phone type links)

4. may be fully connected (at least richer then global networks)

5. cost to interface is low (communications system may be tailored to host interface)

6. heavy data movement is expected (networking is a primary activity)

7. network is optimally designed for particular application

8. packet life is short

9. very little buffering in the network (usually none)

10. transmission medium itself is [illegible] a bottleneck

11. central management of the system by a single person or local group

On return from a coffee break, the group proceeded to discuss how these attributes might affect the protocol designs of local networks. Aside from obvious conclusions that local networks have different topological characteristics than global networks and that protocol designs will be greatly influenced by the applications intended for the network, several generalizing comments were made:

> Local networks respond quickly with acks and/or data; and, therefore, pipelining won't buy much and can be ignored.
>
> We still need to identify packets with sequence numbers to handle lost packets; for some applications, a 1-bit number might do.
>
> Local networks with high bandwidth give you room to play. You usually start with simple less efficient protocols because you have "bandwidth to burn". This

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 -- physically easier to handle but the problem is very much applications dependent. At the media level we may have little problem. Especially since the network itself might be physically contained in a small geographic area under one management. Obviously, physically larger local nets may have more severe security problems.

The broadcast capabilities of many local networks are usually used in "one-of-a-kind" applications that seldom require acknowledgement. Examples are: resetting all time of day clocks, down loading all microprocessors, informing nodes that the network is going down, etc. Links are reliable enough so that broadcasting works most of the time.

Addressing is more complicated in local area networks because the user is more intimate -- he requires a more dynamic addressing capability then global networks seem to provide. This is due, in part, to a greater dependance and higher use of networking capabilities. This seems to be a Network Operating Systems problem. Generally, it is felt that in a local area network you might be willing to "burn some bandwidth" searching through address spaces or global catalogs that you wouldn't consider doing in a global network.

Homogeneity -- Local networks seem to be built so that heterogeneous devices can interact with each other, or so that the performance characteristics of homogeneous devices are increased.

The group concluded that indeed the characteristics of local area networks are different enough from global networks to warrant serious consideration in the design and structure of the protocols and interfaces used in such systems.

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

Chairman: Philip Stein, National Bureau of Standards

Recorder: Gary Donnelly, National Security Agency

Attendees:

Phil Stein	NBS
Gary Donnelly	NSA
George McClure	Ford Motor Co.
Jim Hanks	MITRE/Bedford
Louis Pouzin	Institut de Recherche d'Informatique et d'Automatique
Ira Cotton	NBS

Summary

There are in general two situations that exist in the analysis of computer network specifications. The first is where the user specifies what a net is to do, and then builds the net (or has it built) to those specifications. This involves the generation of applications-driven network specifications. The second is where a user looks at an existing or proposed net to see if it does what the application requires. For the second case, it is necessary that the network specifications be given in terms that are meaningful to the user.

A general feeling prevailed that computer networks are being designed and built with little or no user involvement. Furthermore, network specifications are given in terms that often are meaningless to the user of that network. While it is recognized that these specifications are necessary in defining the characteristics of the net, it is also necessary to give the user the specifications in terms that are of interest to him.

An attempt was made to define several parameters that a user would like to know, in addition to the standard specifications typically given with networks. It was felt that the measures usually given are not of interest in the applications area, and could sometimes mislead the user of the network. For example, data rates are usually specified via Poisson distributions while the user might need to know an absolute upper bound that can be delivered by the network. Response times are also given statistically. However, an example of an industrial workstation requiring service in two seconds or less was cited. In this example,

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 necessary for an application user of the network to know. While the definitions are not complete, it was felt that they represent a minimal set.

1. T1: Time from a request for service until completion of service by the network. This should be specified as a dead time plus a burst rate. (This is a one way communication .)

2. T2: Time from a request for service until receipt of service, excluding processing at server end. (This is a two-way communication.)

3. CL: An overall confidence level of the above parameters. This is also the probability of success of the above responses.

   The above parameters assume that the network is working at the time of the request. The following definitions specify what the user needs to know about the networks not working.

4. Availability and schedule of network and terminal services. Is the system available 24 hours a day? Can the user get to terminals during non-prime time? Will the net be down during some period on a recurring basis?

5. Response during degradation of services (Sometimes described as: "What is the network not doing while it is doing things that is isn't supposed to be doing?"). Will the user get meaningful error diagnostics? Will any damage occur to his files? Is there any priority of services during a degraded mode?

6. Response during restoration of services. Is the network available to all users, but with limited services? Does the net require a period of time during restoration where it is unavailable to all users?

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 consensus that design specifications are not always generated with the end user in mind, 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.

If the network is being implemented on top of an existing computer environment, such as connecting a group of individual computers , then the network should appear to be totally transparent to the user. There should be no control characters for the net, and perhaps an out-of-band signalling scheme should be used by the net.

It was also decided that a local network should be completely deterministic in its operation. That is, it should give predictable responses for given actions. This is not the case in general with large computer systems or networks. The use of a deterministic network would allow the user always to see predictable interactions from the net and its resources, despite the origin of the resources - including behavior during degradation. This would also allow the use of "cookbooks" for naive users, and would also allow automata to be used at terminal nodes. This was important in an industrial application, where a microprocessor was acting as the user of the network instead of a human. Such is the case in several networks already in use in industry.

Finally, there was some discussion as to the differences users saw between a global net and a local net. It was generally felt that there were only slight differences. Mostly they had to deal with the way in which a network failed. A global network tends not to fail completely. For example, the ARPAnet might have individual hosts that are down, or links between computers not functioning, but the entire net won't be inoperative. This is not generally true of a local net. In an industrial environment, it is imperative that the network designers consider the case where the net could be down due to a power failure of an entire area. This requires careful thought to things such as graceful degradation, and restart procedures.

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 transfering 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, Pogran 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.

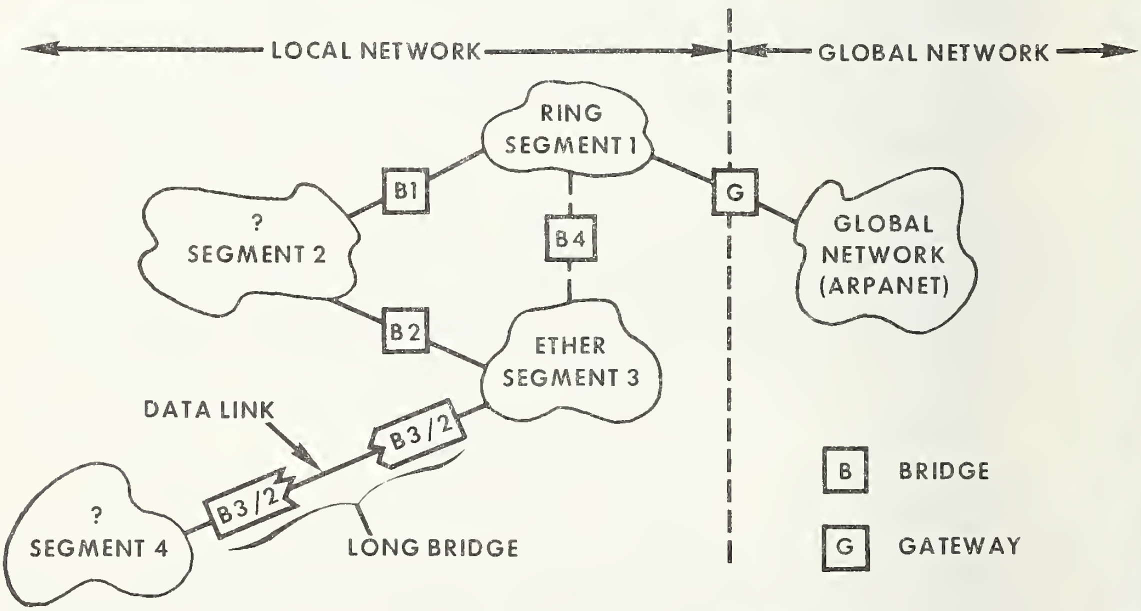

FIGURE 1. MULTISEGMENT LOCAL NETWORK EMPLOYING BRIDGES

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 desination 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.

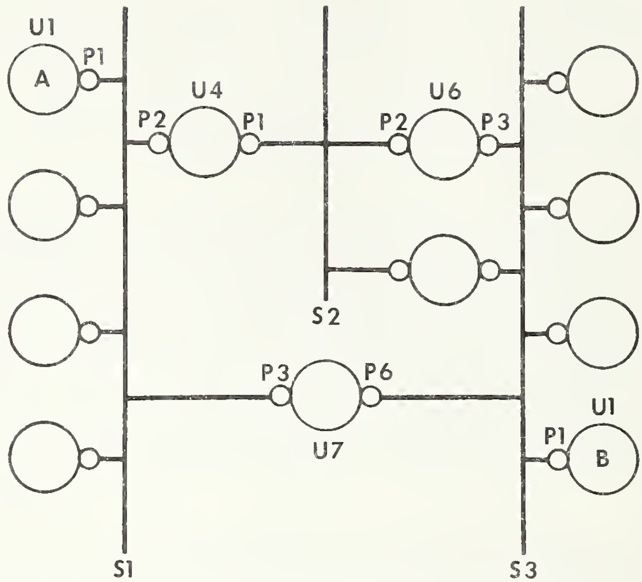

PATH AB IS VIA U7...P3, P1, S3, U1; P6, P1, S1, U1; CONTROL; DATA
OR VIA U4, U6...P2, P2, P1, S3, U1; P3, P1, P1, S1, U1; CONTROL; DATA

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.

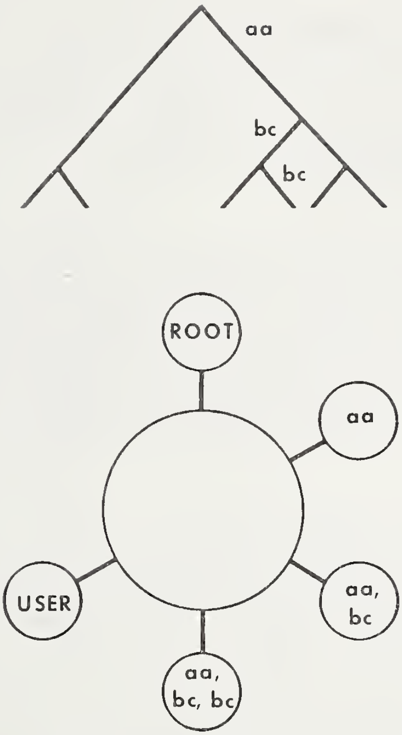

FIGURE 3. LOGICAL NAME STRUCTURE

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 stategies 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."

<u>Inter-process Communication (IPC) Levels</u>

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.

<u>Dynamic Network Reconfiguration</u>

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 excuted 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 langauge 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. Mockapetris, U. C. Irvine
Mel Gable, Ford Motor/Science Lab
Jeff Yeh, L. 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).

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 moving the majority of the tasks to the interface unit and away from the host software.

c) Instrumentation of the Network

The discussion centered not so much on what instrumentation should be provided, but rather on what instrumentation is provided on extant systems, and whether or not it was included in the original design.

System	Instrumentation	Included in Original Design
DCS IRVINE	Software Probes	No
Packet Radio Bay Area	Software & hardware support including special packets	Yes
MITRENET	Essentially None	?
Ford Motor Co.	None	No
Xerox	1. Collision Counts 2. Special packets and test routing (promiscuity mode) 3. Oscilloscope	Yes

Little additional commentary on instrumentation was possible except:

1. it should be considered as part of the original design
2. it should reflect the data needs of extant or proposed models
3. distributions (histograms) are more important than means and variances.

Finally it was pointed out that existing instrumentation in the ETHERNET has indicated traffic patterns force individual interfaces to be non-Poisson.

d) Performance Measures

The following list of performance measures was developed.

1. Throughput - rate of successful completion of transmission of message/packet.

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. Queueing delays - including those associated with host-host and communication subnet delays.
4. Utilization - fraction of bandwidth actually used.
5. Reliability
6. Buffering requirement - in host and/or interface unit (if any).
7. Processing requirement - on host, (see discussion in b).
8. Availability - several standard definitions are possible, including MTBF. It was suggested that the best measure may be probability (of finishing what I start), and is obviously related to the individual user's service requirements.
9. Recovery time - following failure and degree of recovery from user's view.
10. Fairness - in delivering service to users.
11. Cost
12. Error rate

An attempt was made to order the measures by importance. After some deliberation it was decided that measure importance is application dependent and general ordering is, therefore, impossible, although there was some consensus that cost and availability may predominate most application dependent orderings. Two additional comments are necessary. First, instrumentation must ensure data to quantify the measures, and second, it is not clear whether the measures should be applied to the communications subset or to host-host (end-to-end) transmissions. Some studies suggest that host operating system code and buffering can be the overwhelming influence on subnet behavior. Additional evidence is required.

e) Selection of modelling tools and model features

It was pointed out that modelling can serve one or more of at least six objectives. Namely modelling can serve:

1. To provide a predictive tool for gaining insight into system performance.
2. To aid in design feature selection.
3. To aid design validation.
4. To aid selling the network or network service.
5. To aid "tuning" of the system, i.e., performance improvement.

It was subsequently decided that the traffic model 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 amalgam 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 via instrumentation of an existing network.
7. Both finite and infinite population models are useful.

The tools available to the modeller include queueing theory (qt), graph theory (gt), network analysis (na), and simulation employing special purpose (s.p) or general purpose (g.p) languages. The following somewhat incomplete table was the result of an attempt to summarize the situation.

Technology	Traffic Models			Prototype in Existence	Reliability Analysis Done	Known Simulation Models & Traffic Types	
	qt models	na models	gt			sp	gp
Random Access	1A,1B			x		1B,1A	1,4,B
Ring	1B			x		1B	
Polling	1B			x			
TDMA	1A			x			1,4,B
Direct Connect/ Store& Forward				x	x	1,5,B	
Circuit Switching	x			x		1,5,B	
Shared Memory	x			x			

Where A => infinite population model
B => finite population model

and

The numbers 1-6 refer to the traffic model types given above.

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. Limited or influenced by technology.
2. Inversely proportional to confidence that a user will place in the resulting system.
3. Limited by available manpower.

On efficiency it was pointed out that packets generally contain small amounts of data. As regards protocol efficiency no determinations were made, but it was decided that efficiency is dictated (in part) by the issue of static (node resident) versus flowing data (contained in packets) to describe path (routing). The issue is application dependent.

AB70 Abramson, Norman. "Another Alternative for Computer Communications." Fall Joint Computer Conference, 1970, pp. 695-702.

AB73 Abramson, Norman. "Packet Switching With Satellites." National Computer Conference, June 1973, pp. 695-702.

AC76 Acree, John and Arthur Lynch. "Ring Architecture Supports Distributed Processing." Data Communications, March/April 1976, pp. 51-55.

AG77 Agrawala, A. K., R. M. Bryant and J. Agre. "Analysis of an Ethernet-like Protocol." Technical Report TR-522, University of Maryland, Department of Computer Science, April 1977.

AS75 Ashenhurst, R. L. and R. H. Vonderohe. "A Hierarchical Network." Datamation, Vol. 21, No. 2, February 1975, pp. 40-44.

BA76 Ball, J. Eugene, Jerome Feldman, James R. Low, Richard Rashid and Paul Rovner. "RIG, Rochester's Intelligent Gateway: System Overview." IEEE Transactions on Software Engineering, Vol. SE-2, No. 4, December 1976, pp. 321-328.

BI75 Binder, R., N. Abramson, F. Kuo, A. Okinaka and D. Wax. "ALOHA Packet Broadcasting -- A Retrospect." National Computer Conference, May 1975, pp. 215.

BU75 Burchfield, J., R. Tomlinson and M. Beeler. "Functions and Structure of a Packet Radio Station." National Computer Conference, May 1975, pp. 245-251.

DE76 DeMarines, Victor A. and Lawrence W. Hill. "The Cable Bus in Data Communications." Datamation, Vol. 22, No. 8, August 1976, pp. 89-92.

FA72a Farber, David J. and Kenneth C. Larson. "The System Architecture of the Distributed Computing System -- The Communications System." Proceedings of the Symposium on Computer-Communications and Teletraffic, New York: Polytechnic Press, 1972; pp. 21-27.

FA72b Farber, David J. and Kenneth C. Larson. "The Structure of a Distributed Computing System -- Software." Proceedings of the Symposium on

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

FA75 Farber, David J. "A Ring Network." Datamation, Vol. 21, No. 2, February 1975, pp. 44-46.

FA73 Farber, David J., et. al. "The Distributed Computing System." Compcon 73, February 1973, pp. 31-34.

FRAL75a Fralick, Stanley C. and James C. Farrett. "Technological Considerations for Packet Radio Networks." National Computer Conference, May 1975, pp. 233-244.

FRAL75b Fralick, Stanley C., David H. Brandin, Franklin F. Kuo and Christopher Harrison. "Digital Terminals for Packet Broadcasting." National Computer Conference, May 1975, pp. 253-261.

FO77 Forsdick, Harry C., Richard E. Schantz and Robert H. Thomas. "Operating Systems for Computer Networks." Bolt Beranek and Newman Inc., Report No. 3614, July 1977.

FRAN75 Frank, Howard, Israel Gitman and Richard Van Slyke. "Packet Radio System -- Network Considerations." National Computer Conference, May 1975, pp. 217-231.

FRAN76 Frank, Howard, Israel Gitman and Richard Van Slyke. "Recent Advances in Ground Packet Radio Systems." Network Analysis Corporation, Sixth Semiannual Technical Report for the Project: Local, Regional and Large Scale Integrated Networks, Volume 2, February 1976.

FRAN76 Frank, Howard, Israel Gitman and Richard Van Slyke. "Local and Regional Communication Networks -- Technologies and Architectures." Network Analysis Corporation, Sixth Semiannual Technical Report for the Project: Local, Regional and Large Scale Integrated Networks, Volume 4, February 1976.

FRAS75 Fraser, A. G. "A Virtual Channel Network." Datamation, Vol. 21, No. 2, February 1975, pp. 51-53.

HAS73 Hassing, Thomas E., Raymond M. Hampton, Gerald W. Bailey and Robert S. Gardella. "A Loop Network for General Purpose Data Communications in a Heterogeneous World." Third Data Communications Symposium, November 1973, pp. 88-96.

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 Analysis of a Ring Switched Data Transmission System." The Bell System Technical Journal. November 1971, pp. 2947-2978.

HU76 Huynh, D., H. Kobayashi and F. F. Kuo. "Design Issues for Mixed Media Packet Switching Networks." National Computer Conference, June 1976, pp. 541-549.

IN75 Innes, D. R. and J. L. Alty. "An Intra University Network." Fourth Data Communications Symposium, October 1975, pp. 1-8 - 1-13.

KA75 Kahn, Robert E. "The Organization of Computer Resources into a Packet Radio Network." National Computer Conference, 1975, pp. 177-186.

KL75 Kleinrock, Leonard and Fouad Tobagi. "Random Access Techniques for Data Transmission Over Packet Switched Radio Channels." National Computer Conference, May 1975, pp. 187-202.

KU74 Kuo, Frank F. and D. R. Binder. "Computer-Communications by Radio and Satellite: The ALOHA System." In Computer Communication Networks, Grimsdale and Kuo, editors, NATO Advanced Studies Institute Series, E-4, Noordhoff Int., Leyden, Netherlands, 1975, pp. 397-408.

KU73 Kuo, Franklin F. and Norman Abramson. "Some Advances in Radio Communications for Computers." Compcon 73, February 1973, pp. 57-60.

LA73 Labonte', Robert C. "A General Purpose Digital Communications System for Operation on a Conventional CATV System." Compcom 73, February 1973, pp. 85-88.

LI76 Lidinsky, William P. "The Argonne Intra-Laboratory Network." Proceedings, Berkeley Workshop on Distributed Data Management and Computer Networks, Lawrence Berkeley Laboratory, Berkeley, California, May 25-26, 1976, pp. 263-275.

ME73 Metcalfe, Robert M. "Packet Communication." Massachusetts Institute of Technology Project MAC Technical Report TR-114, December 1973. (Ph. D.

Thesis, Harvard University).

ME76 Metcalfe, Robert M. and David R. Boggs. "Ethernet: Distributed Packet Switching for Local Computer Networks." Communications of the ACM, Vol. 19, No. 7, July 1976, pp. 395-404.

MO77 Mockapetris, Paul V., Michael R. Lyle and David J. Farber. "On the Design of Local Network Interfaces." IFIP Congress, August 1977, pp. 427-430.

PI72 Pierce, John R. "How Far Can Data Loops Go?" IEEE Transactions on Communications, Vol. COM-20, No.3, June 1972, pp. 527-530.

RE76 Reams, Cecil C. and Ming T. Liu. "Design and Simulation of the Distributed Loop Computer Network (DLCN)." Third Annual Symposium on Computer Architecture, January 1976, pp. 124-129.

RE75 Reams, Cecil C. and Ming T. Liu. "A Loop Network for Simultaneous Transmission of Variable-Length Messages." Second Annual Symposium on Computer Architecture, January 1975, pp. 7-12.

ROB73 Roberts, Lawrence G. "Dynamic Allocation of Satellite Capacity Through Packet Reservation." National Computer Conference, June 1973, pp. 711-716.

ROD77 Rodgers, John C. "Computer Networking With a Data Bus." Sixteenth Annual Technical Symposium, Association for Computing Machinery (Washington, D.C. chapter) and the National Bureau of Standards, June 1977, pp. 45-50.

ROS73 Rosen, Saul and John M. Steele. "A Local Computer Network." Compcon 73, February 1973, pp. 129-132.

SP71 Spragins, John D. "Analysis of Loop Transmission Systems." Second Symposium on Problems in the Optimization of Data Communications Systems, October 1971, pp. 175-182.

SW72 Switzer, I. "The Cable Television System as a Computer-Communications Network." Proceedings of the Symposium on Computer-Communications and Teletraffic, New York: Polytechnic Press, 1972; pp. 339-346.

TO76 Tobagi, Fouad A., Stanley E. Lieberson and Leonard Kleinrock. "On Measurement Facilities in Packet Radio Systems." National Computer Conference, June 1976, pp.

589-596.

WEa77 West, Anthony R. "A Broadcast Packet-Switched Computer Communications Network -- Design Progress Report." Queen Mary College (University of London), Department of Computer Science and Statistics, unpublished report, February, 1977.

WE172 West, Lynn P. "Loop Transmission Control Structures." IEEE Transactions on Communications, Vol. COM-20, No. 3, June 1972, pp. 531-539.

WILK75 Wilkes, M. V. "Communication Using a Digital Ring." Proceedings of the Pacific Area Computer Communication Network System Symposium, August 1975, pp. 47-56.

WILL74 Willard, David G. "A Time Division Multiple Access System for Digital Communication." Computer Design, June 1974.

YU76 Yuill, Stuart J. "A Survey of Multiple Access Satellite Packet Communication." Fifteenth Annual Technical Symposium, Association for Computng Machinery (Washington, D.C. chapter) and the National Bureau of Standards, June 1976, pp. 9-16.

NBS Workshop on Local Area Computer Networking

August 22 and 23, 1977

Attendees

Abrams, Marshall D. National Bureau of Standards Bldg. 225, Rm. B222 Washington, D. C. 20234	301/921-2601
Agrawala, Ashok K. University of Maryland Computer Science Department College Park, MD 20742	301/454-4251
Altaber, Jacques CERN Prevessin F-01631 CERN Cedex, France	833570
Carpenter, Robert National Bureau of Standards Bldg. 225, Rm. A219 Washington, D.C. 20234	301/921-3427
Christensen, Gary S. Network Systems Corp. 6820 Shingle Creek Pkwy. Brooklyn Center, MN 55430	612/566-5050
Cotton, Ira W. National Bureau of Standards Bldg. 225, Rm. B212 Washington, D.C. 20234	301/921-2601
Donnelly, Gary National Security Agency 9800 Savage Road Ft. Meade, MD 20755	301/688-8147
Elischer, Victor Lawrence Berkeley Lab. Real Time Systems Group #1 Cyclotron Rd. Berkeley, CA 94720	415/843-2740, x6378 FTS - 451-6378

Forsdick, Harry C. 617/491-1850, x638
Bolt, Beranek & Newman, Inc.
50 Moulton Street
Cambridge, MA 02138

Fralick, Stanley 408/733-1300
Comtech. Adv. Syst. Inc.
754 N. Pastoria
Sunnyvale, CA 94086

Franta, W. R. 612/373-9751
University of Minnesota
143 Space Science Bldg.
100 Union St., S.E.
Minneapolis, MN 55455

Gable, Mel 313/594-0904
Ford Motor Co.
P.O. Box 2053
Science Lab., Rm. S2037
Dearborn, MI 48121

Gordon, Karen Delouche 301/454-4251
University of Maryland
Computer Science Dept.
College Park, MD 20740

Hart, Edwin F. 301/953-7100, x2318
Applied Physics Lab.
Johns Hopkins Rd.
Laurel, MD 20810

Hopkins, Gregory T. 617/271-2038
The MITRE Corp.
P.O. Box 208
Bedford, MA 01730

Hanks, James P. 617/271-2595
The MITRE Corp.
P.O. Box 208
Bedford, MA 01730

Hopper, Andrew 02-235-2435
Computer Laboratory
University of Cambridge
Corn Exchange St.
Cambridge CB2 3QG, England

Kimbleton, Stephen R. 301/921-2601
National Bureau of Standards
Bldg. 225, Rm. B214
Washington, D. C. 20234

Larsen, Ronald L. 301/982-2777
NASA - Goddard Space Flight Center
Code 570.1
Greenbelt, MD 20771

Lyle, Michael
University of California - Irvine
Computer Science Department
Irvine, California 92717

McClure, George 313/323-4760
Ford Motor Co.
Scientific Research Lab.
Rm. S-2037
P.O. Box 2053
Dearborn, MI 48121

Meissner, Paul 301/921-3427
National Bureau of Standards
Tech. Bldg., Rm. A219
Washington, D.C. 20234

Metcalfe, Robert M. 415/494-4274
Xerox SDD
3408 Hillview Ave.
Palo Alto, CA 94304

Mills, David L. 202/554-6246
COMSAT, Advanced Systems
950 L'Enfant Plaza
Washington, D.C. 20024

Mockapetris, Paul 714/833-5233
University of California
ICS Dept.
Irvine, CA 92717

Nelson, David L. 617/879-2960, x493
Prime Computer, Inc.
145 Pennsylvania Ave.
Framingham, MA 01701

Pilipchuk, Andrew — 301/454-4526
University of Maryland
Computer Science Center
College Park, MD 20742

Pogran, Kenneth T. — 617/253-6012
MIT Laboratory for Computer Science
545 Technology Sq.
Cambridge, MA 02139

Pouzin, Louis — 954.90.20
Institut de Recherche d'Informatique
et d'Automatique
Domaine de Voluceau
Rocquencourt
France

Reed, David P. — 617/253-6004
MIT
545 Technology Sq., Rm. 510
Cambridge, MA 02139

Rogers, Sig. G. — 415/843-2740, x5354
Lawrence Berkeley Lab.
1 Cyclotron Rd.
Bldg. 50A-6102
Berkeley, CA 94720

Rosenthal, Robert — 301/921-2601
National Bureau of Standards
Bldg. 225, Rm. B212
Washington, D.C. 20234

Rowe, Edward, Jr. — 301/688-8147
National Security Agency
9800 Savage Rd.
Ft. Meade, MD 20755

Sherman, Richard H. — 313/594-0744
Ford Motor Co.
20000 Rotunda Dr.
Science Research Lab., S2037
Dearborn, MI 48170

Stein, Phil — 301/921-2805
National Bureau of Standards
Bldg. 221, Rm. A345
Washington, D.C. 20234

Wecker, Stuart Digital Equipment Corp. ML3-2/E41 146 Main St. Maynard, MA 01754	617/897-5111, x4366
West, Anthony R. Queen Mary College Computer Systems Lab. Mile End Rd. London E1-4NS, England	01 980-4811, x751
Wilner, David Lawrence Berkeley Lab. Real Time Systems Group #1 Cyclotron Rd. Berkeley, CA 94720	415/843-2740, x6378 FTS - 451-6378
Wood, John Lawrence Berkeley Lab. 1 Cyclotron Rd. Bldg. 50A, Rm. 1144 Berkeley, CA 94720	415/843-2740, x5972
Yeh, Jeffry Lawrence Livermore Lab. Computation Dept., L-316 P.O. Box 808 Livermore, CA 94550	415/447-1100, x5100
Young, Carol National Bureau of Standards A-223 Administration Building Washington, D.C. 20234	301/921-3384
Zobrist, Dale Eldec Corp. 16700 13th Ave. West P.O. Box 100 Lynnwood, WA 98036	206/743-1313

NBS-114A (REV. 7-73)

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. Report of a Workshop Held at the National Bureau of Standards, Gaithersburg, Maryland, August 22-23, 1977			5. Publication Date April 1978
			6. Performing Organization Code
7. AUTHOR(S) Ira W. Cotton, Chairman & Editor			8. Performing Organ. 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.
			11. Contract/Grant No.
12. Sponsoring Organization Name and Complete Address *(Street, City, State, ZIP)* Same as item 9			13. Type of Report & Period Covered
			14. Sponsoring Agency Code

15. SUPPLEMENTARY NOTES

Library of Congress Catalog Card Number: 78-606029

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:

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.

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

18. AVAILABILITY	19. SECURITY CLASS (THIS REPORT)	21. NO. OF PAGES
[x] Unlimited [ ] For Official Distribution. Do Not Release to NTIS	UNCLASSIFIED	82
[x] Order From Sup. of Doc., U.S. Government Printing Office Washington, D.C. 20402, SD Stock No. SN003-003— [ ] Order From National Technical Information Service (NTIS) Springfield, Virginia 22151	20. SECURITY CLASS (THIS PAGE) UNCLASSIFIED	22. Price $2.40

ANNOUNCEMENT OF NEW PUBLICATIONS ON COMPUTER SCIENCE & TECHNOLOGY

Superintendent of Documents,
Government Printing Office,
Washington, D. C. 20402

Dear Sir:

Please add my name to the announcement list of new publications to be issued in the series: National Bureau of Standards Special Publication 500-.

Name ______________________________

Company ______________________________

Address ______________________________

City ______________ State ______________ Zip Code ______________

(Notification key N-503)

☆U.S. GOVERNMENT PRINTING OFFICE:1978 261-238/83 1-3

NBS TECHNICAL PUBLICATIONS

PERIODICALS

JOURNAL OF RESEARCH—The Journal of Research of the National Bureau of Standards reports NBS research and development in those disciplines of the physical and engineering sciences in which the Bureau is active. These include physics, chemistry, engineering, mathematics, and computer sciences. Papers cover a broad range of subjects, with major emphasis on measurement methodology, and the basic technology underlying standardization. Also included from time to time are survey articles on topics closely related to the Bureau's technical and scientific programs. As a special service to subscribers each issue contains complete citations to all recent NBS publications in NBS and non-NBS media. Issued six times a year. Annual subscription: domestic $17.00; foreign $21.25. Single copy, $3.00 domestic; $3.75 foreign.

Note: The Journal was formerly published in two sections: Section A "Physics and Chemistry" and Section B "Mathematical Sciences."

DIMENSIONS/NBS

This monthly magazine is published to inform scientists, engineers, businessmen, industry, teachers, students, and consumers of the latest advances in science and technology, with primary emphasis on the work at NBS. The magazine highlights and reviews such issues as energy research, fire protection, building technology, metric conversion, pollution abatement, health and safety, and consumer product performance. In addition, it reports the results of Bureau programs in measurement standards and techniques, properties of matter and materials, engineering standards and services, instrumentation, and automatic data processing.

Annual subscription: Domestic, $12.50; Foreign $15.65.

NONPERIODICALS

Monographs—Major contributions to the technical literature on various subjects related to the Bureau's scientific and technical activities.

Handbooks—Recommended codes of engineering and industrial practice (including safety codes) developed in cooperation with interested industries, professional organizations, and regulatory bodies.

Special Publications—Include proceedings of conferences sponsored by NBS, NBS annual reports, and other special publications appropriate to this grouping such as wall charts, pocket cards, and bibliographies.

Applied Mathematics Series—Mathematical tables, manuals, and studies of special interest to physicists, engineers, chemists, biologists, mathematicians, computer programmers, and others engaged in scientific and technical work.

National Standard Reference Data Series—Provides quantitative data on the physical and chemical properties of materials, compiled from the world's literature and critically evaluated. Developed under a world-wide program coordinated by NBS. Program under authority of National Standard Data Act (Public Law 90-396).

NOTE: At present the principal publication outlet for these data is the Journal of Physical and Chemical Reference Data (JPCRD) published quarterly for NBS by the American Chemical Society (ACS) and the American Institute of Physics (AIP). Subscriptions, reprints, and supplements available from ACS, 1155 Sixteenth St. N.W., Wash., D.C. 20056.

Building Science Series—Disseminates technical information developed at the Bureau on building materials, components, systems, and whole structures. The series presents research results, test methods, and performance criteria related to the structural and environmental functions and the durability and safety characteristics of building elements and systems.

Technical Notes—Studies or reports which are complete in themselves but restrictive in their treatment of a subject. Analogous to monographs but not so comprehensive in scope or definitive in treatment of the subject area. Often serve as a vehicle for final reports of work performed at NBS under the sponsorship of other government agencies.

Voluntary Product Standards—Developed under procedures published by the Department of Commerce in Part 10, Title 15, of the Code of Federal Regulations. The purpose of the standards is to establish nationally recognized requirements for products, and to provide all concerned interests with a basis for common understanding of the characteristics of the products. NBS administers this program as a supplement to the activities of the private sector standardizing organizations.

Consumer Information Series—Practical information, based on NBS research and experience, covering areas of interest to the consumer. Easily understandable language and illustrations provide useful background knowledge for shopping in today's technological marketplace.

Order **above** *NBS publications from: Superintendent of Documents, Government Printing Office, Washington, D.C. 20402.*

Order **following** *NBS publications—NBSIR's and FIPS from the National Technical Information Services, Springfield, Va. 22161.*

Federal Information Processing Standards Publications (FIPS PUB)—Publications in this series collectively constitute the Federal Information Processing Standards Register. Register serves as the official source of information in the Federal Government regarding standards issued by NBS pursuant to the Federal Property and Administrative Services Act of 1949 as amended, Public Law 89-306 (79 Stat. 1127), and as implemented by Executive Order 11717 (38 FR 12315, dated May 11, 1973) and Part 6 of Title 15 CFR (Code of Federal Regulations).

NBS Interagency Reports (NBSIR)—A special series of interim or final reports on work performed by NBS for outside sponsors (both government and non-government). In general, initial distribution is handled by the sponsor; public distribution is by the National Technical Information Services (Springfield, Va. 22161) in paper copy or microfiche form.

BIBLIOGRAPHIC SUBSCRIPTION SERVICES

The following current-awareness and literature-survey bibliographies are issued periodically by the Bureau:

Cryogenic Data Center Current Awareness Service. A literature survey issued biweekly. Annual subscription: Domestic, $25.00; Foreign, $30.00.

Liquified Natural Gas. A literature survey issued quarterly. Annual subscription: $20.00.

Superconducting Devices and Materials. A literature survey issued quarterly. Annual subscription: $30.00. Send subscription orders and remittances for the preceding bibliographic services to National Bureau of Standards, Cryogenic Data Center (275.02) Boulder, Colorado 80302.