


MTP-3.57

SURVEY OF COMPUTER NETWORKS

JACK J. PETERSON
SANDRA A. VEIT

SEPTEMBER 1971



THE
MITRE
 CORPORATION
 WASHINGTON OPERATIONS

SURVEY OF COMPUTER NETWORKS

JACK J. PETERSON

SANDRA A. VEIT

The work reported here was sponsored by the
Defense Communications Agency under
contract F19628-71-C-0002

SEPTEMBER 1971

THE
MITRE
CORPORATION

ABSTRACT

This paper presents the results of a survey of state-of-the-art computer networks. It identifies ten major networks: ARPA, COINS, CYBERNET, the Distributed Computer System, DLS, MERIT, Network/440, Octopus, TSS, and TUCC and outlines their capabilities and design. A tabular presentation of the most significant network features and a brief discussion of networks that were examined but rejected for the survey are also included.

ACKNOWLEDGMENTS

The authors of this survey thank the organizations mentioned herein for their assistance in providing much of the basic information from which this survey was compiled. We wish to extend special thanks to the individuals named below, who gave a good deal of their time for site interviews, telephone conversations, and correspondence with us:

Jack Byrd, Jim Caldwell, and Jim Chidester of Control Data Corporation; Doug McKay and Al Weis of IBM; Don Braff, John Fletcher, Mel Harrison, and Sam Mendicino of the Lawrence Radiation Laboratory; Eric Aupperle and Bertram Herzog of MERIT; Peggy Karp and David Wood of MITRE; Dan Cica, Wayne Hathaway, Gene Itean, Marge Jereb, and Roger Schulte of NASA; Doug Engelbart and Jim Norton of the Stanford Research Institute; Leland Williams of TUCC; and David Farber of the University of California at Irvine.

FOREWORD

Data for this survey was gathered primarily from interviewing individuals at the various network sites. A questionnaire was used as a checklist during the interviews, but not as a tool for comparative evaluation of the networks because of the wide range of questions and because of the vast differences among the networks. In many cases additional information was obtained from literature provided by the interviewees or their installation.

Most of the information furnished by this survey was gathered between January and April 1971; however, in this rapidly expanding area most networks are in the process of changing. This document gives a picture of these networks as they were at a given point in time; where possible, proposed or impending changes have been indicated. Each section of the survey has been reviewed by the cognizant organization to ensure greater accuracy, although errors are inevitable in an undertaking of this magnitude.

TABLE OF CONTENTS

SECTION		PAGE
	LIST OF FIGURES	vii
I	INTRODUCTION	1
II	NETWORKS SURVEYED	3
	The ARPA Computer Network	3
	The COINS Network	11
	The CYBERNET Network	14
	The Distributed Computer System	18
	Data Link Support (DLS)	23
	The MERIT Computer Network	25
	Network/440	32
	The Octopus Network	35
	The TSS Network	48
	The TUCC Network	53
III	MATRIX OF NETWORK FEATURES	59
	Configuration	60
	Communications	64
	Network Usage	66
IV	EXCLUDED NETWORKS	69
V	SUMMARY	71
	GLOSSARY	73
	APPENDIX	83
	BIBLIOGRAPHY	85

LIST OF FIGURES

FIGURE		PAGE
1	ARPA Network Topology, February 1971	4
2	Inventory of Nodes and Host Hardware in the ARPA Network	5
3	The Interface Message Processor	6
4	COINS Configuration	12
5	The CYBERNET Network	15

FIGURE		PAGE
6	Typical CYBERNET Configurations	16
7	The Distributed Computer System Topology	19
8	Inventory of Planned Hardware	20
9	Communications Interface	21
10	DLS Configuration	23
11	Overview of the MERIT Network	26
12	Inventory of MERIT Host Hardware	27
13	MERIT Communications Segment	28
14	Communication Computer System	29
15	Logical Structure of Network/440	33
16	Nodes in Network/440	33
17	The Octopus Network	37
18	Octopus Hardware	38
19	Television Monitor Display System (TMDS)	39
20	6600/PDP-10 File Transport Channel	41
21	7600/PDP-10 File Transport Channel	42
22	Octopus Teletype Subnet	43
23	Remote Job Entry Terminal (RJET) System and Network Connections	45
24	An Overview of the TSS Network	49
25	TSS Network Hardware	50
26	Usage of the TSS Network	52
27	An Overview of the TUCC Network	55
28	Configuration of the 360/75 at TUCC	56

SECTION I INTRODUCTION

As defined in this paper, a computer network, is an interconnected group of independent computer systems which communicate with one another and share resources such as programs, data, hardware, and software. This paper presents the results of a survey of state-of-the-art computer networks by MITRE under the sponsorship of the Defense Communications Agency. It identifies the major networks according to the working definition given above and includes a discussion of their purpose, configuration, usage, communications and management.

The bulk of the paper consists of a discussion of the selected networks and a matrix presentation of some of the more predominant characteristics of each. Section II presents much of the information gathered in the course of the study; it is divided into ten subsections, one for each of the networks surveyed. Each of the subsections (networks) is further divided into five topic areas: Introduction, Configuration, Communications, Usage, and Management. A comparative matrix in Section III gives an overview of the characteristics of the networks. Section IV briefly examines networks that were not included in the survey. Section V presents a summary of the survey. The Glossary provides definitions of terms and acronyms which may be unfamiliar to the reader.

SECTION II NETWORKS SURVEYED

Each subsection in Section II presents the findings pertaining to one network. All network discussions are organized in the same manner and deal with five basic topics: *Introduction* gives background information such as the sponsor, purpose and present status of the network. *Configuration* provides an inventory of network hardware, generally accompanied by a topological diagram of the network, and information on network software. *Communications* relates the relevant factors in the communications subsystem of the network. *Usage* discusses the present or intended use of the network. *Management* presents a view of the network management structure.

THE ARPA COMPUTER NETWORK

Introduction

The Advanced Research Projects Agency (ARPA) network is a nationwide system which interconnects many ARPA-supported research centers. The primary goal of this project is to achieve an effective pooling of all of the network's computer resources, making them available to the network community at large. In this way, programs and users at a particular center will be allowed to access data and programs resident at a remote facility.

At the present time, network activity is concentrated in three major areas. The first is the installation of the network interface hardware, and the development and testing of its associated software modules. Secondly, network experimentation is being carried out at several operational sites. These experiments are designed to develop techniques for measuring system performance, for distributing data files and their directories, and for disseminating network documentation. Finally, expansion and refinement of the original system design are being investigated, with considerations being paid to both long-range and immediate goals.

Configuration

The ARPA Network is a distributed network of heterogeneous host computers and operating systems. ARPA's store-and-forward communication system consists of modified Honeywell DDP-516 computers located close to the hosts and connected to each other by 50 kilobit-per-second leased telephone lines. The 516 is called an Interface Message Processor, or IMP.

The Network Control Program (NCP) is generally part of the host executive; it enables processes within one host to communicate with processes on another or the same host. The main functions of the NCP are to establish connections, terminate connections, and control traffic flow.

Figure 1 is a topological diagram of the ARPA Network. Figure 2 lists the network nodes along with a brief description of the hardware and software at each. Although this compilation is approximate at the time of this writing, it provides a general idea of the resources available at various nodes in the ARPA Network.

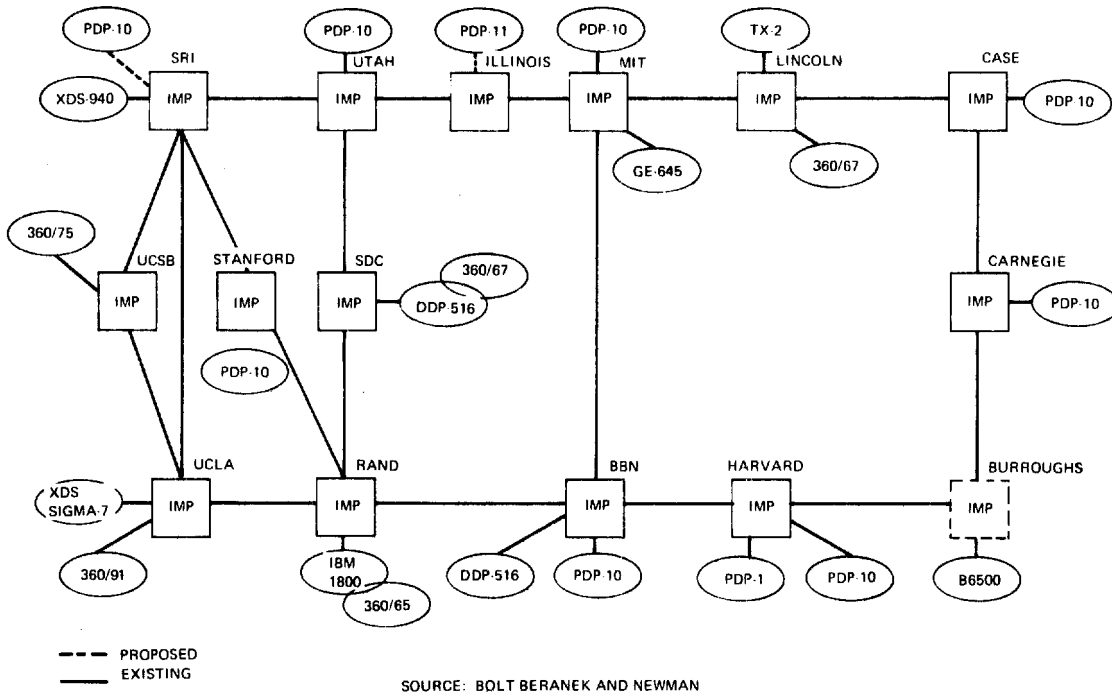


Figure 1 ARPA Network Topology, February 1971

Communications

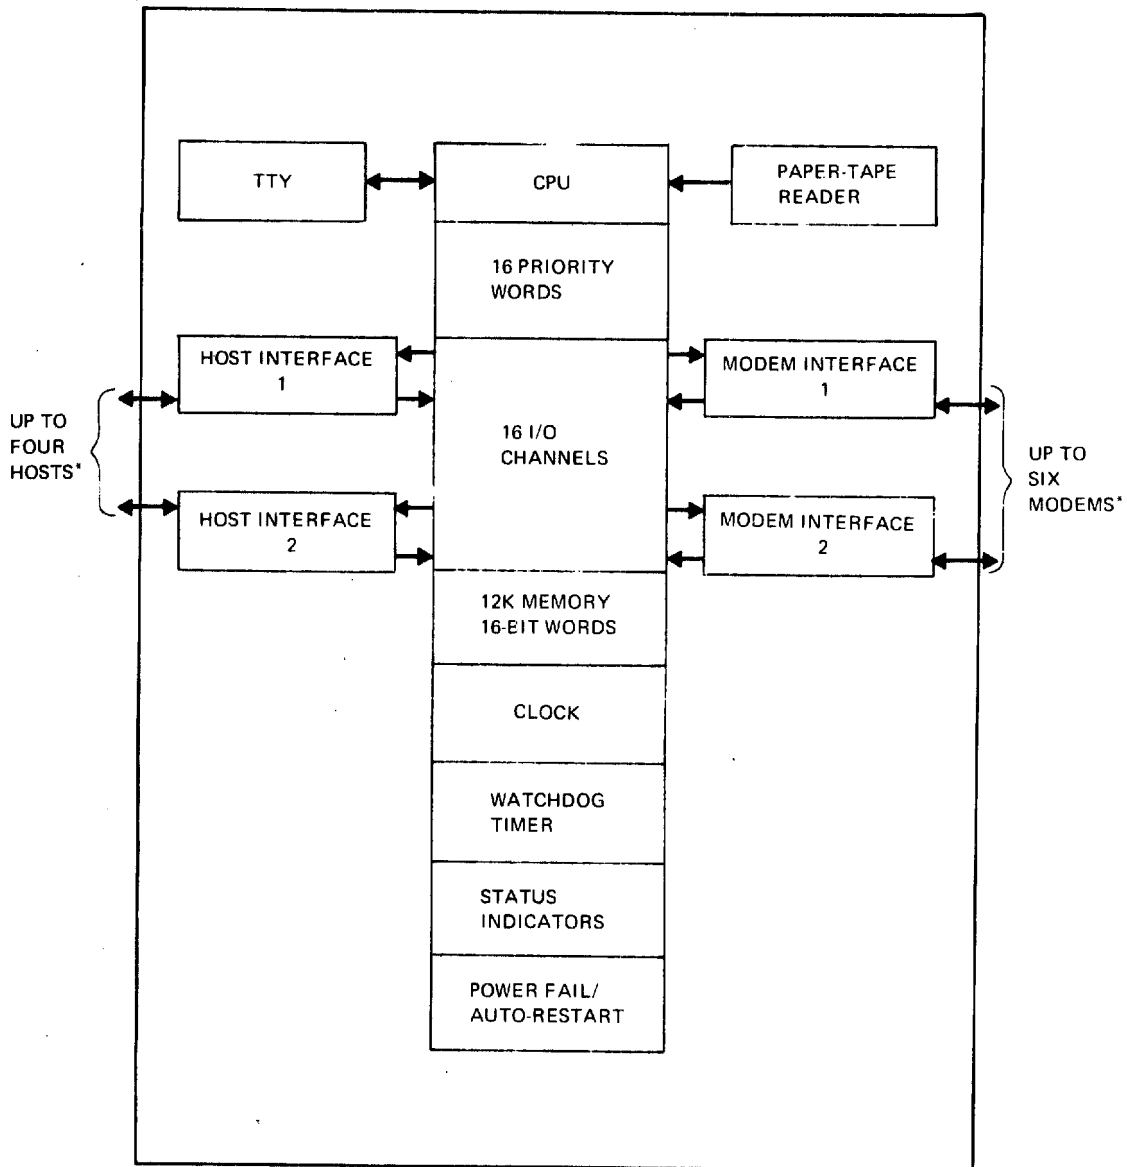
Communications in the ARPA network are achieved using a system of leased lines, operated in a synchronous, full-duplex mode at 50,000 bps. The interconnection of the host computers to the telephonic network is the primary function of a specially developed communications computer system, the Interface Message Processor (IMP).¹ Each IMP, as shown in Figure 3, is an augmented, ruggedized version of the Honeywell DDP-516, and includes 12K 16-bit words of core memory, 16 multiplexed channels, 16 levels of priority interrupt, and logic supporting host computers and high-speed modems. Special

¹ A second device, the Terminal Interface Processor (TIP), is under development for use on the ARPA network. It not only performs the same function as an IMP, but can also directly support user terminals, eliminating the need for a host. The first TIP is scheduled to go into operation in August 1971 at NASA Ames.

NODE	PROCESSOR	SPECIAL NODE FUNCTIONS OR SOFTWARE
BOLT BERNAEK AND NEWMAN (CAMBRIDGE MASSACHUSETTS)	PDP 10 (TENEX)	LISP VERSION ALLOWING 250K WORDS OV VIRTUAL MEMORY NATURAL LANGUAGE PROCESSORS NETWORK CONTROL CENTER
BURROUGHS (PAOLI, ILLINOIS) ¹	B8500	
CARNEGIE MELLON UNIVERSITY (PITTSBURGH, PENNSYLVANIA)	PDP 10 50	LC ² AN INTERACTIVE ALGOL LANGUAGE
CASE WESTERN RESERVE (CLEVELAND, OHIO)	PDP 10	LOGOS (DATA MANAGEMENT)
HARVARD UNIVERSITY (CAMBRIDGE, MASSACHUSETTS)	PDP 10 PDP 1	GRAPHICS
LINCOLN LABORATORY (LEXINGTON, MASSACHUSETTS)	TX 2	LEAP - A GRAPHIC LANGUAGE LIL - LOCAL INTERACTION LANGUAGE
	360/67 TSP	
MASSACHUSETTS INSTITUTE OF TECHNOLOGY (CAMBRIDGE, MASSACHUSETTS)	GE 645 (MULTICS)	
	PDP 10	ARTIFICIAL INTELLIGENCE MATHLAB
	PDP 6	DYNAMIC MODELING
RAND (SANTA MONICA, CALIFORNIA)	360 65 VIA IBM 1800	ADAPTIVE COMMUNICATION PROJECT NETWORK SERVICES PROGRAM CONVERSATIONAL PROGRAMMING SYSTEM
STANFORD RESEARCH INSTITUTE (MENLO PARK, CALIFORNIA)	PDP 10 (TENEX)	NETWORK INFORMATION CENTER TODAS NLS
STANFORD UNIVERSITY (STANFORD, CALIFORNIA)	PDP 10 PDP 6	ARTIFICIAL INTELLIGENCE PROJECTS
SYSTEM DEVELOPMENT CORPORATION (SANTA MONICA, CALIFORNIA)	360 67 VIA ODP 516 (ADEPT)	CONVERSE DISPLAY 70 ORBIT
UNIVERSITY OF CALIFORNIA AT SANTA BARBARA	360/75	UCSB COLLIER-FRIED ON-LINE SYSTEM
UNIVERSITY OF CALIFORNIA AT LOS ANGELES	XDS SIGMA-7 360/91 (OS MVT)	NETWORK MEASUREMENT CENTER REMOTE JOB SERVICE
UNIVERSITY OF ILLINOIS (URBANA, ILLINOIS)	PDP 11	
UNIVERSITY OF UTAH	PDP 10/50	SPECIALIZED GRAPHICS SERVICES
AIR WEATHER SERVICE ¹ (OFFUTT AIR FORCE BASE, OMAHA, NEBRASKA)	UNIVAC 1108	
ETAC ¹ (WASHINGTON, D.C.)	USER ONLY	
MITRE/WASHINGTON ¹ (MCLEAN, VIRGINIA)	USER ONLY	
NASA AMES RESEARCH CENTER ¹ (MOUNTAINVIEW, CALIFORNIA)	360/67 ILLIAC IV VIA B6500	
	PDP 10 DATA MACHINE	
NATIONAL BUREAU OF STANDARDS ¹ (GAITHERSBURG, MARYLAND)	PDP 11	
NATIONAL CENTER FOR ATMOSPHERIC RESEARCH ¹ (BOULDER, COLORADO)	CDC 6600 CDC 7600	SPECIALIZED FOR DATA MANAGEMENT
RAOC ¹ (ROME, NEW YORK)	GE-645	
LONDON UNIVERSITY ² (LONDON, ENGLAND)	PDP-9 CDC 8800	
OCAMA ² (TINKER AIR FORCE BASE, OKLAHOMA CITY, OKLAHOMA)	UNIVAC 418	
SAAC ² (ALEXANDRIA, VIRGINIA)		
SAAMA ² (MCLELLAN AIR FORCE BASE, SACRAMENTO, CALIFORNIA)	UNIVAC 418	
UNIVERSITY OF SOUTHERN CALIFORNIA ² (LOS ANGELES, CALIFORNIA)	360/44	

¹SCHEDULED NODES.
²PROPOSED NODES.

Figure 2 Inventory of Nodes and Host Hardware in the ARPA Network



*THE NUMBER OF HOSTS PLUS THE NUMBER OF MODEMS MAY NOT EXCEED SEVEN.

SOURCE: HEART, F. E., et al., "THE INTERFACE MESSAGE PROCESSOR FOR THE ARPA COMPUTER NETWORK," PROCEEDINGS OF THE SPRING JOINT COMPUTER CONFERENCE, MAY, 1970, P. 558.

Figure 3 The Interface Message Processor

hardware is provided to detect certain internal failures, and to either correct them or to gracefully power down if correction is not possible. Each IMP is capable of supporting up to four hosts, with the restriction that the number of hosts plus the number of transmission lines may not exceed seven. Software support is derived from a specially developed operating system which requires approximately 6K words of core memory; the remaining 6K words are used for message and queue storage. The operating system is identical for all IMP's except for a protected 512-word block which contains programs and data unique to each. This allows an IMP which has detected a software failure to request a reload of the program from a neighboring IMP.

The IMP hardware is activated by a host computer whenever a message is ready for transmission. Such messages are variable length blocks with a maximum size of 8095 bits. The host interface portion of the IMP, which is its only host-dependent component, operates in a bit-serial, full-duplex fashion in transferring the message between the host and IMP memories. A data-demand protocol is used in the interface to match the transfer rates of the two processors.

Messages received by the IMP are segmented into variable length "packets," each having a maximum size of approximately 1000 bits. Packets serve as the basic unit record of information interchange between IMP's. Their smaller size places a reduced demand on intermediate message-switch storage, and increases the likelihood of an error-free transmission. Parity check digits, which provide an undetected error rate of about 10^{-12} , are appended to the packets. The packets are then queued for transmission on a first-in, first-out basis.

The selection of the particular link over which a packet is to travel is determined by the IMP's estimation of the delay in reaching its destination over each of its available lines. These estimates, which are recomputed at approximately 500-millisecond intervals, are based on the exchange of estimates and past performance records between neighboring IMP's. As a consequence of this estimation capability, transmission paths which maximize effective throughput are selected. In addition, since these estimates are dynamic, the several packets which comprise a message need not use the same physical path through the network to their destination.

IMP activity is also initiated upon receipt of a packet from another IMP. A packet error check is performed first. If the packet is error-free, it is stored and a positive acknowledgment is returned to the sending IMP, allowing it to release the packet from its storage area. If the packet contains errors, or if the receiving IMP is too busy or has insufficient storage to accept it, the packet is ignored. The transmitting IMP waits a pre-determined amount of time for a positive acknowledgment; if none is detected, the packet is assumed lost and retransmitted, perhaps along a different route.

Once a positive acknowledgment has been generated, the receiving IMP must determine, by an examination of the destination field in the packet header, whether the packet is to be delivered to a local host or forwarded. In the latter case, the packet is queued for transmission in a fashion similar to that used for locally initiated messages. Otherwise, the IMP must determine whether all the packets comprising the message have arrived. If so, a reassembly task is invoked to arrange the packets in proper order and to transfer the message to the host memory.

In addition to its message handling functions, the IMP provides special capabilities for the detection of communication failures and the gathering of performance statistics. In the absence of normal message traffic, each IMP transmits idling packets over the unused lines at half-second intervals. Since these packets must be acknowledged in the usual manner, the lack of any packet or acknowledgment traffic over a particular line for a sustained period (about 2.5 seconds) indicates a dead line. Local routing tables may be up-dated to reflect the unavailability of such a line. The resumption of line operation is indicated by the return of idling packet traffic.

The IMP is capable of gathering statistics on its own performance. These statistics, which are automatically transmitted to a specified host for analysis, may include summaries, tabulation of packet arrival times, and detailed information describing the current status of the packet queues. All network IMP's can provide these statistics on a synchronized, periodic basis, allowing the receiving host to formulate a dynamic picture of overall network status.

An additional capability supporting performance evaluation is tracing. Any host-generated message may have a trace bit set. Whenever a packet from such a message is processed, each IMP records the packet arrival time, the queues on which the packet resided, the duration of the queue waits, the packet departure time, etc. These statistical records, which describe the message-switch operation at a detailed level, are automatically transmitted to a specified host for assembly and analysis.

Usage

The use of the ARPA Network has been broken into two phases related to the network implementation plans:

- initial research and experimental use; and
- external research community use.

The first phase involves the connection of approximately 14 sites engaged principally in computer research into areas such as computer systems architecture, information system

design, information handling, computer augmented problem solving, intelligent systems and computer networking. The second phase extends the number of sites to about 20.

During the final phase one the network usage consists primarily of sharing software resources and gaining experience with the wide variety of systems. This enables the user community to share software, data, and hardware, eliminating duplication of effort.

The second phase activities will consist of adding new nodes to take advantage of other research in such areas as behavioral science, climate dynamics and seismology. Data distribution, data sharing and the use of the ILLIAC IV in climate dynamics and seismology modeling are areas of special interest. One of the uses of the network will be to share data between data management systems or data retrieval systems; this is regarded as an important phase because of its implications for many government applications.

A network node for data management is being designed by Computer Corporation of America (CCA); it will consist of a PDP-10, one trillion bits of on-line laser memory, interfaced with the B6500/ ILLIAC IV processing complex. CCA plans to implement a special data language to talk to the "data machine," having disk storage and a slower, trillion bit direct-access store that will provide an alternative to storage at network sites.

The network is also used to access the Network Information Center (NIC) at SRI; the NIC serves as a repository of information about all systems in the network that can be dynamically updated and accessed by users.

Another use of the network is measurement and experimentation; because of the nature of the network, much effort has been expended developing appropriate tools for collecting usage statistics and evaluating network performance. Bolt, Beranek and Newman (BBN), the Network Control Center, gathers information such as:

- the up/down status of the hosts and telephone lines;
- the number of messages failing to arrive over each telephone line;
- the number of packets successfully transmitted over each telephone line; and
- the number of messages transmitted by each host into its IMP.

Additional information is being gathered by UCLA, the Network Measurement Center.

Management

Although the several nodes of the ARPA network are at least partially supported by ARPA, each is an independent research facility engaged in many activities beyond its participation in the network. One of the primary considerations of the network design philosophy and of its management is the preservation of this autonomy.

As a consequence, administrative control of the computer systems has remained with the individual facilities, while the responsibility for intercomputer communications has been assumed by network management.

The management of the network is functionally distributed between two organizations. Fiscal policy, particularly the disbursement of funds, is determined by the Information Processing Office of ARPA. The technical pursuit of the network is the responsibility of the Chairman of the Network Working Group (NWG), who is appointed by ARPA.

The NWG itself is composed of at least one member from each participating site. It meets every three months and operates in a somewhat informal fashion. Its main purpose is to propose and evaluate ideas for the enhancement of the network. To this end, several subcommittees have been formed within the NWG, each involved with a single major aspect of network operation. Their respective areas of inquiry include the following:

- data transformation languages;
- graphics protocol;
- host-host protocol;
- special software protocol; and
- accounting.

The critical need for the timely dissemination of technical information throughout the network community is satisfied by means of a three-level documentation scheme. The most formal papers are called "Documents," and are issued by the Chairman of the NWG as a statement of network technical policy. A "Request for Comments" (RFC) is issued by any member of the NWG as a means of proposing technical standards. RFC's are therefore technical opinions and serve to promote the exchange of ideas among the NWG. An RFC Guide which indexes and defines the status of all RFC's is published periodically by The MITRE Corporation. Finally, RFC's, Documents, substantive memoranda, telephone conversations, site documents, and other appropriate material are cataloged by the NIC at the Stanford Research Institute (SRI), which periodically publishes a comprehensive index to these materials. SRI has also developed two sophisticated software systems to enable a network user to effectively utilize the information in the catalog files. The first of these is the Typewriter Oriented Documentation Access System (TODAS). This system, as its name implies, is intended to provide the teletype terminal user with appropriate capabilities for manipulating the library catalogs. These facilities include text editing, record management, keyword searching, and display of formatted results. The second system, which is similar to TODAS but far more powerful, employs graphic display devices with specially developed keyboards in place of the teletype.

THE COINS NETWORK ¹

Introduction

The Community On-Line Intelligence System (COINS) was proposed in 1965 as an experimental program. Its primary purpose is to assist in determining methods of improving information handling among the major intelligence agencies.

The COINS network is currently operational as an experimental system. The research that has been carried out to date has been concerned almost exclusively with the means of sharing pertinent data among the network users. This is a particularly complex problem in the intelligence community because of the variety of hardware, software, and standards that are used. Studies are also underway to demonstrate the applicability of a common network control language and a common data management system to be implemented at all sites.

Configuration

COINS is a geographically distributed network of heterogeneous computers and operating systems working through a central switch, an IBM 360/30. Linked to the switching computer are a GE 635, and two Univac 494 installations (one of which is a triple processor). The configuration is illustrated in Figure 4. Some agencies participate in the network via terminal connection to one of the participating computer systems.

Communications

Communications are achieved in the COINS network by a centralized message switch and conditioned, leased voice-grade lines. The lines, which connect each host computer to the central switch, are operated in a full-duplex, synchronous mode at 2400 bps using modems. The transmission system is completely secure, using cryptographic equipment throughout the network.

A host computer may transmit a message of up to 15,000 characters to another host; however, a message must be subdivided into segments of no more than 150 characters prior to transmission. All characters transmitted use the 7-bit ASCII code with additional bit for parity. Each segment of a message must be sent and acknowledged.

¹COINS is no longer operational; it is included here as a matter of historical record.

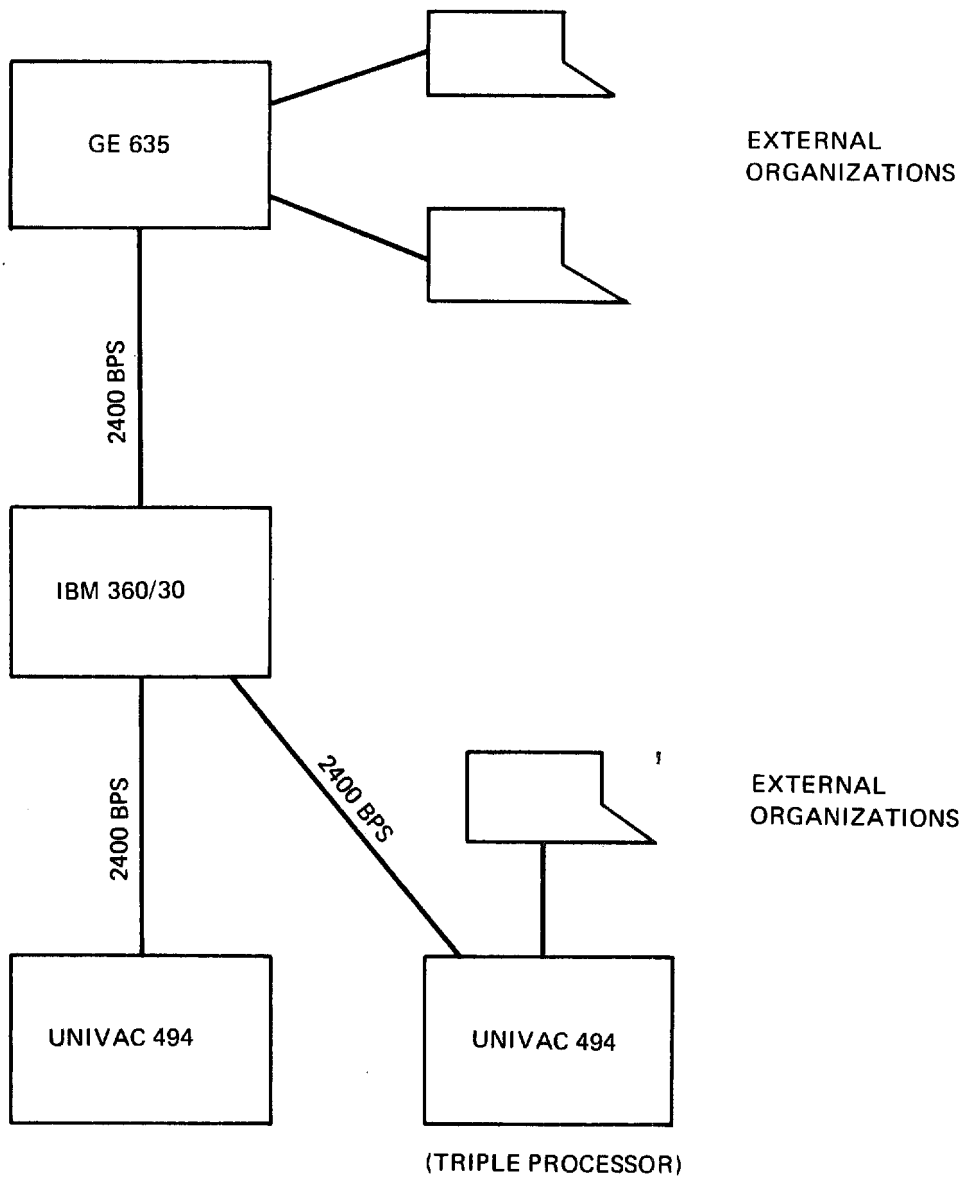


Figure 4 COINS Configuration

Usage

COINS is being used experimentally to enable various intelligence agencies to share their data bases with each other. These data bases are constantly changing, and the responsibility for building, maintaining and updating a data base rests solely with its sponsor. Users at terminals cannot change the data bases; they can only query them. A response time of less than 15 minutes is the goal, but in practice it ranges from five minutes to two hours. The response time achieved is dictated to a great extent by the workload of the file processors responding to interrogations.

Management

The management of the COINS network is vested in the Project Manager who is responsible for the design and operation of the network. He is assisted by a Subsystem Manager from each of the participating agencies who represents the interests of his agency.

One of the more critical problems faced by the Project Manager is the establishment of acceptable procedures governing the inclusion of files. Currently, through a formalized nomination procedure, a network user may request that a file maintained by one of the participating agencies be made available for network access. The Project Manager coordinates such requests by determining whether other users also require the files or by establishing the necessary justifications. Subsequently, the request is forwarded to the particular agency, which maintains the exclusive right to accept or deny the request.

A forum for the presentation and discussion of interagency problems is provided by four panels, each consisting of one or more individuals from each agency. Although the panels can make specific recommendations, final decisions rest exclusively with the Project Manager and Subsystem Managers. The User Support Panel is responsible for conducting training seminars in network usage and for distributing network documentation among the users. The Security Panel is tasked with investigating procedures for ensuring adequate security on the network computers. The gathering and evaluation of network performance statistics is the responsibility of the Test and Analysis Panel. Finally, the Computer and Communications Interface Panel is concerned with the system software and network communications, and the protocol used in network operation.

THE CYBERNET¹ NETWORK

Introduction

The CYBERNET network is a nationwide commercial network offering computing services to the general public. CYBERNET is operated as a division of the Control Data Corporation, and represents a consolidation of their former Data Center operation. By interconnecting the individual service centers, CDC feels that the user is offered several unique advantages which include the following:

- better reliability, by offering local users a means for accessing a remote computer in the event of local system failure;
- greater throughput, by allowing local machine operators to transfer parts of an extra heavy workload to a less busy remote facility;
- improved personnel utilization, by allowing the disperse elements of a corporation to more readily access one another's programs and data bases; and
- enhanced computer utilization, by allowing the user to select a configuration which provides the proper resources required for the task.

Configuration

CYBERNET is a distributed network composed of heterogeneous computers, mainly CDC 6600's and CDC 3300's linked by wideband lines. Figure 5 gives a geographic picture of CYBERNET and a partial inventory of its hardware.


The 6600's are considered the primary computing element of the network and are referred to as "centroids" where many jobs are received and processed. Future centroids will include a 7600 and other CDC machines to be announced. The 3300's serve as front ends" and concentrators for the 6600; they are referred to as "nodes." In addition, small satellite computers can be used as terminals to the CYBERNET network; they are distinguished by the fact that they have remote off-line processing capabilities and are able to do non-terminal work while acting as a terminal. These satellites include CDC 3150's, CDC 1700's, and lower scale IBM 360's. Figure 6 gives some typical system configurations.

CYBERNET supports essentially four types of terminals:

- Interactive/conversational (MARC² I);
- Low-, medium-, and high-speed peripheral processors (MARC II, III, IV);
- Small- to medium-scale satellite computers (MARC V); and
- Large- to super-scale computers with terminal facilities (MARC VI).

¹CYBERNET is a registered trademark of the Control Data Corporation.

²Multiple Access Remote Computer.



SOURCE: THE CONTROL DATA CORPORATION

Figure 5 The CYBERNET Network

PROCESSOR	CORE AND SECONDARY MEMORY	OTHER COMPONENTS
TYPICAL 6600 CONFIGURATION	131K MAIN CORE 2 167-MILLION WORD DISKS	SCOPE OPERATING SYSTEM DEPENDENT SUBSYSTEMS - FORTRAN, COBOL, ETC. EXPORT 8231 - WIDEBAND REMOTE TERMINAL SUBSYSTEMS SYSTEM 2000 - A BATCH OR INTERACTIVELY ACCESSIBLE LIST STRUCTURED DATA BASE MANAGEMENT SYSTEM
CDC 6400	131K MAIN CORE 6638 DISK FILE 848.4 MULTIPLE SPINDLE DISK DRIVE	KRONOS - TIME SHARING OPERATING SYSTEM DEPENDENT SUBSYSTEMS - FORTRAN, COBOL, ETC. TELEX - TELETYPEWRITER COMMUNICATIONS MODULE AND SWAPPING EXECUTIVE EXPORT 200 - CDC 200 BATCH COMMUNICATIONS MODULE EXPORT 8231 - WIDEBAND TERMINAL COMMUNICATIONS MODULE IMPORT 6600 - WIDEBAND LINK TO 6600 MODULE WITH INPUT CONCENTRATION AND OUTPUT DISPERSION FACILITY.
TYPICAL 3300 CONFIGURATION	131K MAIN CORE 854 DISK DRIVES	MASTER MULTIPROGRAMMING OPERATING SYSTEM DEPENDENT SUBSYSTEMS - FORTRAN, COBOL, ETC. SHADOW - COMMUNICATIONS AND MESSAGE SWITCHING SUBSYSTEM SHADE - RECORD MANAGING SUBSYSTEM

Figure 6 Typical CYBERNET Configurations

Terminals are categorized in the above manner to indicate their hardware and/or software characteristics. For example, the CDC 200 User Terminal is the CYBERNET standard for low- and medium-speed devices; other devices which have been programmed to resemble this terminal include the CDC 8090, the CDC 160A, the IBM 1130, the Univac 9200, UCC's COPE series, the IBM 360/30 and higher, and the Honeywell 200.

Software available through CYBERNET includes FORTRAN, COBOL, COMPASS assembly language, ALGOL, SIMSCRIPT, GPSS, SIMULA, JOVIAL, BASIC, the SYSTEM 2000 Data Management System, the EASE structural analysis package, the STARDYNE dynamic structural analysis package, and a large statistical library. Linear programming systems include OPHELIE II, PDQ/LP, OPTIMA and NETFLOW (transportation).

Communications

The communications facilities of the CYBERNET network consist of two primary elements: the transmission system and the nodes. The transmission system itself includes four major components: lines, modems, multiplexers, and switches.

CYBERNET employs a variety of lines connecting terminals with computers and computers with one another. For the most part the lines are either switched or leased lines, but private lines are occasionally used, and at least one satellite communications link is in use. Switched lines are operated at low speeds, and include both local and Direct Distance Dial (DDD) facilities. Leased lines include Foreign Exchange (FX)

facilities and point-to-point connections. Measured and full period inward WATS lines are also provided for operation at moderate speeds. Finally, wide-band full period lines are used between computer complexes.

A corresponding complement of modems is used throughout CYBERNET. Typewriter-like terminals are supported by Western Electric 103A modems, operating at rates of up to 300 bps. Medium speed terminals use Western Electric 201A and 201B modems, operating at 2000 bps and 2400 bps respectively on switched and leased lines. High speed terminals use Milgo and ADS modems operating at up to 4800 bps on leased lines. Western Electric 303 modems are used on the wideband lines, operating at 40800 bps.

Multiplexers are used to increase the transmission efficiency of voice-grade lines supporting low-speed terminals. The principal multiplexing configurations are designed to drive the leased-lines at their full capacity of 2400 or 4800 bps by operating as many as 52 low-speed devices simultaneously on the same line. Cost savings are realized by having low-speed terminal users dial in to the local multiplexers rather than directly to a remote computer.

Western Electric line switches have been used throughout CYBERNET to establish terminal-to-computer and computer-to-computer connections. The switches are operated similarly to a telephone exchange system. The switches are not dependent on any of the computer systems, providing a highly reliable mode of operation.

CYBERNET supports two types of nodes: remote job entry and conversational. Each type of node can concentrate message traffic, perform message switching, and provide a user processing capability. The remote job entry node is a CDC 3300 operating with the MASTER multiprogramming operating system. A special subsystem called SHADOW has been developed for this configuration to provide the necessary support for the communications and message-switching functions of the nodes. The SHADOW software is capable of supporting remote-job entry from typewriter-like and CDC 200-Series terminals. Communication from the 3300 to either another 3300 or to a 6600 is also supported by SHADOW.

The conversational nodes of CYBERNET are CDC 6400's operating under an extended version of the KRONOS time-sharing operating system. At the present time, the 6400 is capable of supporting teletypes in a conversational mode, and remote batch 200-Series terminals. Planned additions to the system include communications capabilities for 3300 and 6600 support and complete message-switching facilities.

Usage

The CYBERNET Network is intended to make the computer utility concept available to all of its commercial users by offering the following services:

- "super-computer" processing;
- remote access;
- a multi-center network;
- file management; and
- an applications library and support.

Load sharing, data sharing, program sharing and remote service are possible over the network. The CDC 3300 nodes are used for remote job entry, and the CDC 6400 is used for time sharing. Both nodes can also serve as front ends or concentrators, can relay messages, and can process jobs. CYBERNET's nodes are intended to provide the following facilities:

- generalized store-and-forward message switching;
- the ability to send work to a system that is not loaded;
- the ability to send work to a system which is not inoperative;
- the ability to utilize a unique application at a particular location;
- the ability to access a data base at another location; and
- the ability to utilize a specific computer system.

Management

The management of the CYBERNET network is centralized, vested in the Data Services Division of CDC. All activities including hardware/software development, resource accountability, and documentation development and dissemination are controlled through this central office.

THE DISTRIBUTED COMPUTER SYSTEM

Introduction

The Distributed Computer System is an experimental computer network being developed by the Information and Computer Sciences Department at the University of California at Irvine. The immediate goals of the project are to design, construct, and evaluate the intercomputer communications network.

The Distributed Computer System is currently in the planning stage. Upon completion of the overall design, the communications interfaces are to be constructed, followed by an experimentation program using small computer systems.

Configuration

When the Distributed Computer System at Irvine becomes operational, it will consist primarily of a store-and-forward communications system with a uni-directional, ring-structure topology. Messages will be forwarded around the ring (which is to be composed of two megabit coaxial cables) until the appropriate destination is reached. Figure 7 illustrates this topology.




Figure 7 The Distributed Computer System Topology

The initial Irvine network will consist of heterogeneous mini-computers located at several nodes on the Irvine Campus. A simple device such as a teletype can be considered a host computer on this network. Figure 8 gives the planned inventory of hardware.

FORTRAN and BASIC will be provided through the network; plans call for other capabilities to be added later.

NODES	CORE SIZE	SECONDARY MEMORY
VARIAN 620/i VARIAN 620/a MICRO 800 3 TELETYPES	8K (16 BIT) 8K (16 BIT) 8K (16 BIT)	IBM 2314 (ONE SPINDLE)

Figure 8 Inventory of Planned Hardware

Communications

Two principal elements comprise the communication subsystem: the transmission lines and the communications interface. The transmission lines actually form three distinct subnets, as Figure 7 shows. The primary subnet forms a closed ring connecting all of the nodes. This is the path which is normally used for all message traffic. The other two subnets, one connecting the even nodes, the other the odd, are included solely for reliability. In the event a particular node should fail, the two adjoining nodes could communicate directly over the backup link.

All of the transmission paths will be coaxial cable carrying digital transmissions using pulse-code modulation (PCM). The links are operated using a simplex protocol with all message traffic traveling in one direction around the ring. Data rates of two million bps are expected to be used in the initial configuration; this rate may be increased to as high as six million bps if conditions warrant.

The communications interface is functionally illustrated in Figure 9. Its primary components and their functions are as follows:

- an input line selector switch which automatically switches to the backup input line whenever the primary line drops out for a predetermined period of time;
- a pair of passive PCM repeaters which autonomously forward messages through the interface;
- a repeater selector which activates the backup PCM repeater in the event that the primary unit fails;
- a shift register which provides intermediate storage for messages originating from and delivered to the host computer; and
- logic modules which operate the previously mentioned components, and determine whether a message is to be delivered or forwarded.

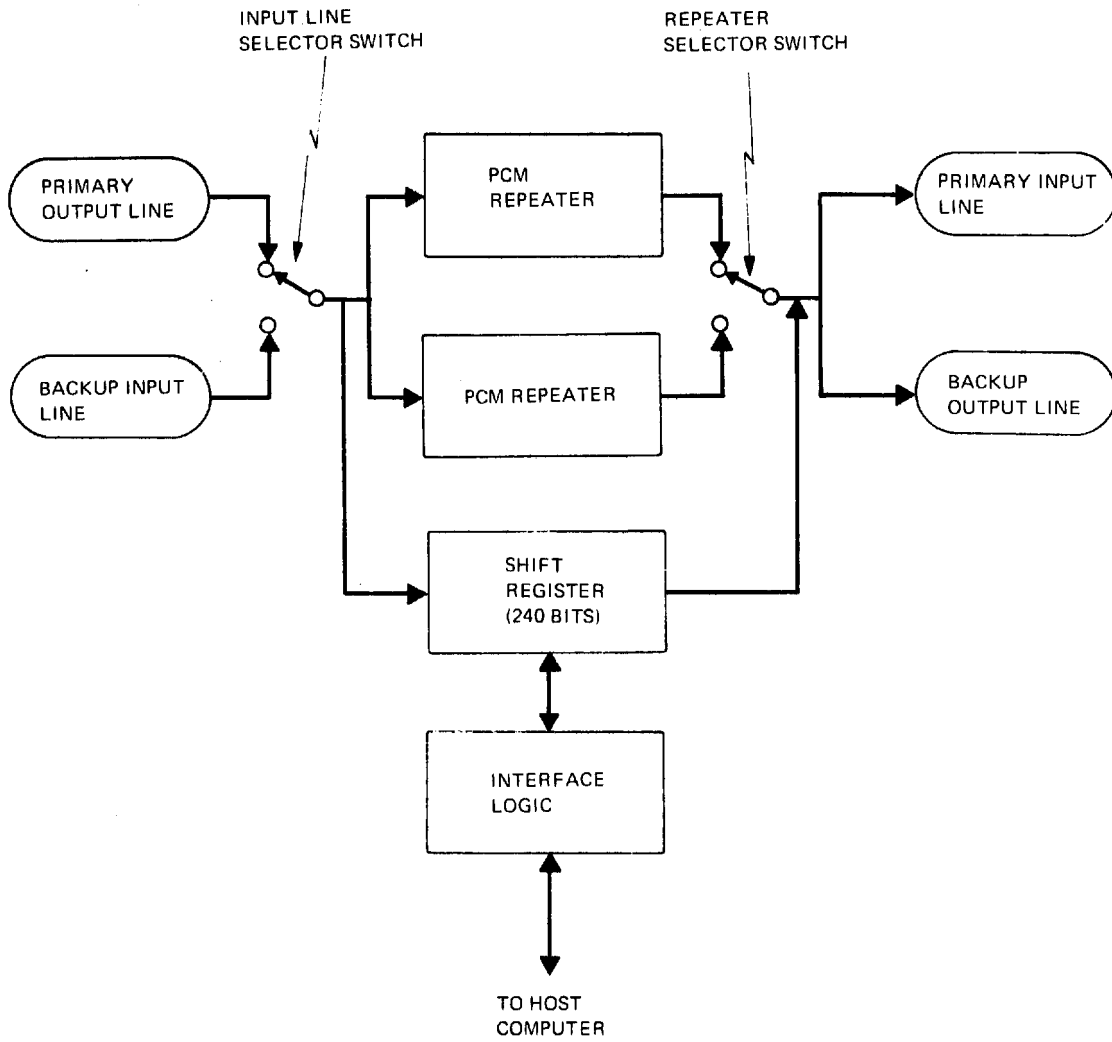


Figure 9 Communications Interface

The communications interface can operate in one of two modes: idle or busy. In the idle mode, the interface can accept messages from either the transmission line or its host. In the former case, the message header is examined to determine whether the destination is the local host; if not, the message is ignored and the PCM repeaters forward it to the next node. If the message is to be delivered, it is removed from the line, placed in the shift register, and checked for errors. If none are detected, a positive acknowledgment is generated and sent to the originating host, and the message is passed to the destination host. If errors are detected, a retransmission request is sent to the originating host.

Upon receipt of a message from its host, the communications interface places the message in the shift register and on the output lines and goes into the busy mode. In this condition, the interface routinely forwards all messages received over the lines, checking only for acknowledgments or retransmission requests. The receipt of a retransmission request indicates that the previously transmitted message was received incorrectly by the destination node. The interface subsequently places the message on the output lines again. A positive acknowledgment, indicating receipt on an error-free message, is passed to the host, and the interface returns to the idle state.

There are two conditions in which a message may circulate in the ring for a protracted period, one of which is the non-existence of the destination node. The other occurs if a message arrives at the destination node when the node interface is in the busy state, unable to accept any messages. In most cases, if the message is allowed to circulate, it will eventually arrive at the destination node while the interface is idle. An interesting exception, however, is the case where two nodes independently and simultaneously send messages to each other. The two messages would circulate forever, since each destination is awaiting an acknowledgment which the other cannot generate. At the present time there is no facility for preventing infinite message loops, although such a capability will probably be added later.

Usage

Because this network is primarily an experimental communications system, very little has been done to provide software to assist users of the network. Load sharing, program sharing, data sharing and remote service are not anticipated in the near future. User software to provide these features and host/host protocol will be developed by the university's computer center once network viability has been demonstrated.

Management

At the present time, the Distributed Computer System is highly localized, involving mainly the resources of the Information and Computer Sciences Department. Consequently, there has been no need for a formalized management structure.

DATA LINK SUPPORT (DLS)

Introduction

The DLS system is a communication facility which connects the National Military Command System Support Center (NMCSSC) with the Alternate National Military Command Center (ANMCC). Its primary purpose is to provide an automated, high-speed capability allowing data bases to be exchanged between the sites and to facilitate computer load leveling by allowing remote program execution.

DLS was developed by the IBM Corporation for the Defense Communications Agency (DCA) during the period June 1969 to June 1970. Final testing was completed in September 1970. The system is currently undergoing further tests and evaluations.

Configuration

Data Link Support (DLS) transmits jobs and data over a 40,800 bps leased line between IBM 2701 Data Adapter Units connected to two IBM 360 computers (Model 50's or larger) operating in a point-to-point mode. Data is hardware encrypted prior to being transmitted and is decoded when received. DLS is currently being operationally tested using a 360/65 at the NMCSSC and a 360/50 at the ANMCC.

DLS is a software package that runs as a problem program in a single region under OS/MVT. Standard OS software is available when using DLS. The DLS configuration as currently used by the NMCSSC and the ANMCC is shown in Figure 10.

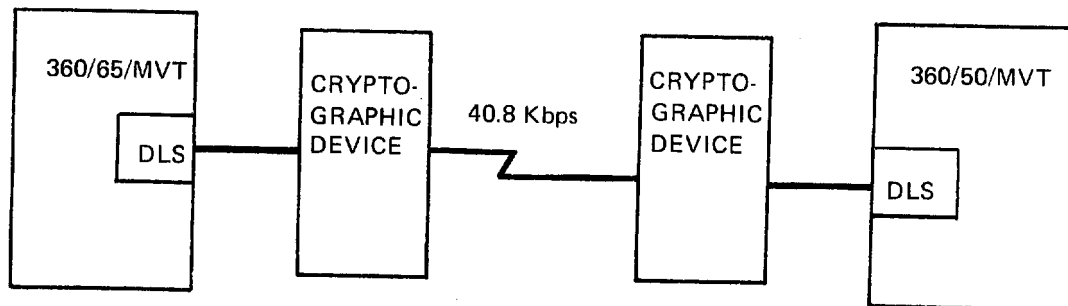


Figure 10 DLS Configuration

Communications

DLS operates between the Support Center and the Alternate Center using a secure, leased wide-band line operated at 40,800 bps. The transmission line is termi-

nated by an IBM 2701 Data Adapter at each end, connected to IBM 2870 Multiplexer Channels. The link uses standard binary-synchronous communication protocol.

Software support for the operation of the link is derived from the DLS program, a copy of which must be operating at each of the link end points. The portion of the DLS program which drives the communications link is called the Binary-synchronous Communication Controller (BCC).

The BCC consists of two primary modules, the Start/Restart Synchronizer (SRS) and the Continuous Communication Access Method (CCAM). SRS is responsible for recognizing the need to start DLS at the remote site, to read and spool job decks, to despool job decks, and dispose of data sets. SRS decodes operator requests, and invokes the necessary support routines to perform the desired function.

CCAM is responsible for maintaining an active channel program for the communication line. CCAM permits multitask usage of the communication link by supporting software multiplexing and demultiplexing functions. The module is also responsible for generating positive acknowledgments upon proper receipt of a message and for requesting retransmission for lost or garbled messages. Message compaction and decompression are also supported by CCAM, as is the gathering of statistics reflecting the performance of the communication subsystem.

Usage

The primary capabilities offered by DLS are data base transmission between remote locations and remote job processing. Thus far DLS has been used primarily to transmit data bases, rather than to achieve load leveling. DLS is used for program sharing, but not extensively because of the large data bases in the operational environment of the NMCS. DLS is designed for batch processing and has no on-line capability.

A job requiring no data base transmission can be transmitted under operator control to the remote site, executed, and the output returned without modification to the deck used when running locally. A job for remote execution which requires data sets located at the local site must include DLS control cards to transmit those data sets. The job is then placed in the reader destined for the appropriate site, or for whichever site is more desirable if the job can run at either location.

Management

The DLS system is controlled by the National Military Command System Technical Support (NMCSTS) Directorate of the Defense Communications Agency (DCA) and is being implemented by the NMCSSC.

THE MERIT COMPUTER NETWORK

Introduction

The Michigan Educational Research Information Triad, Inc. (MERIT) network is a cooperative venture among the three largest universities in Michigan: Michigan State University, Wayne State University, and the University of Michigan. The central purpose of this undertaking is the "development and implementation of a prototype educational computing network, whereby the educational and computing resources of each university may be shared as well as enhanced by each other."¹

The development of the MERIT network is proceeding in two stages. The first of these, funded by grants from the Michigan State Legislature and the National Science Foundation, calls for the development and installation of all network hardware and software modules by June 1971. Subsequently, network experimentation projects will begin, advancing research in information retrieval and computer-aided instruction systems.

Configuration

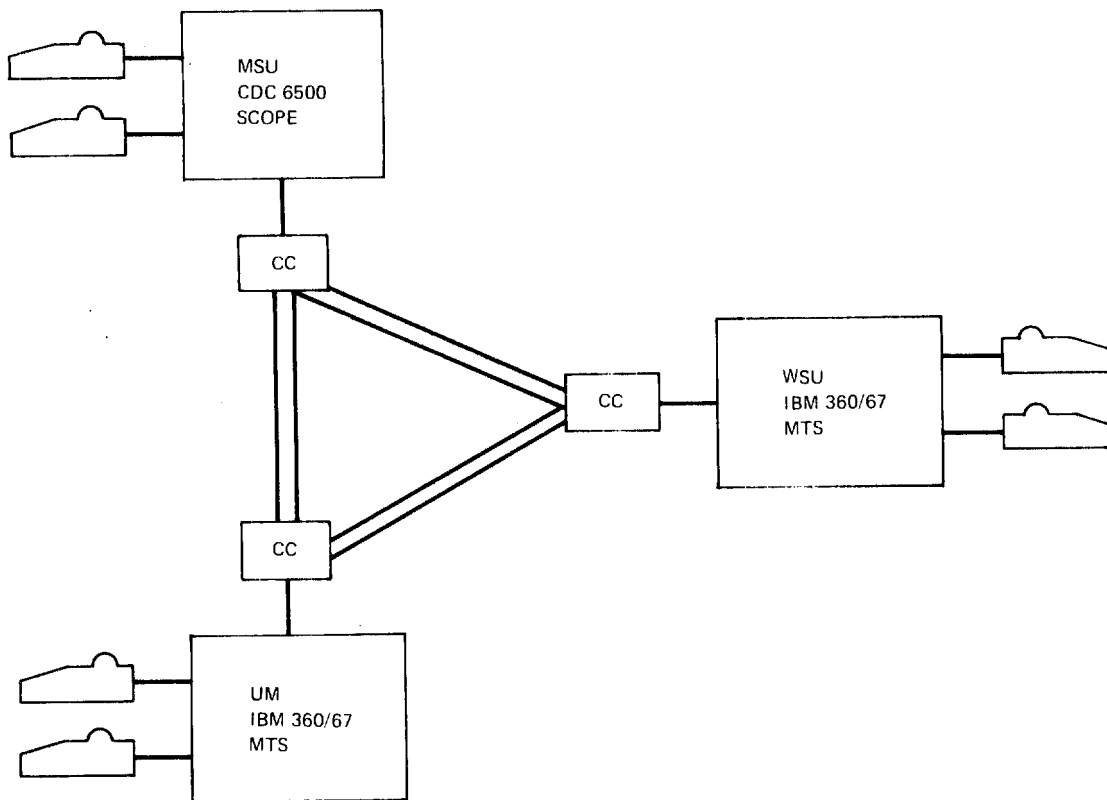
MERIT is a distributed network of heterogeneous computers with nodes at Michigan State University (MSU) in East Lansing, Wayne State University (WSU) in Detroit, and the University of Michigan (UM) in Ann Arbor. Each host is connected to the network through a communications computer (CC), a modified PDP 11/20 computer with a special purpose operating system for communications. Data is transmitted over 2000 bps, voice-grade lines with eight lines connected to each CC. Figure 11 presents an overview of the MERIT Computer Network.

UM runs Michigan Terminal System (MTS) on a duplex IBM 360/67; MTS can service over fifty time-sharing terminals and several batch jobs. MSU uses a CDC 6500 with the SCOPE operating system. WSU has an IBM 360/67 and runs the MTS operating system. Figure 12 presents an inventory of the host hardware at each of the three nodes.

Communications

The communications subsystem of the MERIT network consists of three functional units: the host interface, the communications computer, and the telephonic network. The interconnection of these modules along a typical communications segment is illustrated in Figure 13.

¹Bertram Herzog, "MERIT Proposal Summary," 2nd Revision, 28 February 1970.



CC: COMMUNICATIONS COMPUTER – DEC PDP-11

SOURCE: THE MERIT COMPUTER NETWORK

Figure 11 Overview of the MERIT Network

The host interface is a specially designed hardware module which interconnects the host computer with the communications computer (CC). This interace provides two primary capabilities. First, it is capable of independently transmitting a variable-length data record¹ to (from) the memory of the CC from (to) the host computer, performing whatever memory alignment operations are required by the different word configurations of the two processors. Secondly, it provides a multiple-address facility which permits the host to treat the CC as several peripheral devices. This greatly simplifies the host software, since a dedicated pseudo-device can be allocated to each user or task requesting use of the communications resources.

¹Record length is determined by a software parameter.

NODE	PROCESSOR	MAIN CORE	OTHER HARDWARE
MSU	CDC 6500	64K (60-BIT WORDS) 4K (12-BIT WORDS) FOR EACH OF THE TEN PERIPHERAL PROCESSORS.	1 CDC DISK SYSTEM (167) MILLION 6-BIT CHARACTERS) 3 854 DISK STORAGE DRIVES (8.2 MILLION 6-BIT CHAR- ACTERS PER DISK PACK). MODEL 33 TELETYPES 2 CDC 200 REMOTE BATCH STATIONS 1 217.2 REMOTE SINGLE STA- TION ENTRY/DISPLAY
UM	360/67 (DUPLEX)	6 UNITS (1.5 MILLION 8 BIT BYTES TOTAL) - VIRTUAL MEMORY MACHINE.	3 2314s (8 DRIVES EACH) 2 DATA CELLS, (800 MILLION BYTES TOTAL) 2 HIGH-SPEED DRUMS (APPROX- IMATELY 7.4 MILLION BYTES TOTAL) IBM 360/20 COMPUTER IBM 1130 COMPUTER IBM 2780 REMOTE JOB ENTRY TERMINAL TERMINALS (IBM 2741, DATEL 30, 33/35 TELETYPES).
WSU	360/67 (HALF DUPLEX) ¹	ONE MILLION BYTES ² - VIRTUAL MEMORY MACHINE	2 2314s ³ (8 DRIVES 8 EACH) 2 DRUMS TERMINALS (IBM 2741, TELE- TYPES, DATA 100)
¹ A DUPLEX SYSTEM BECOMES OPERATIONAL IN JUNE. ² AN ADDITIONAL 125K WILL BE ADDED WHEN THE DUPLEX SYSTEM BECOMES OPERATIONAL. ³ TWO MORE 2314s WILL BE ADDED IN JUNE.			

Figure 12 Inventory of MERIT Host Hardware

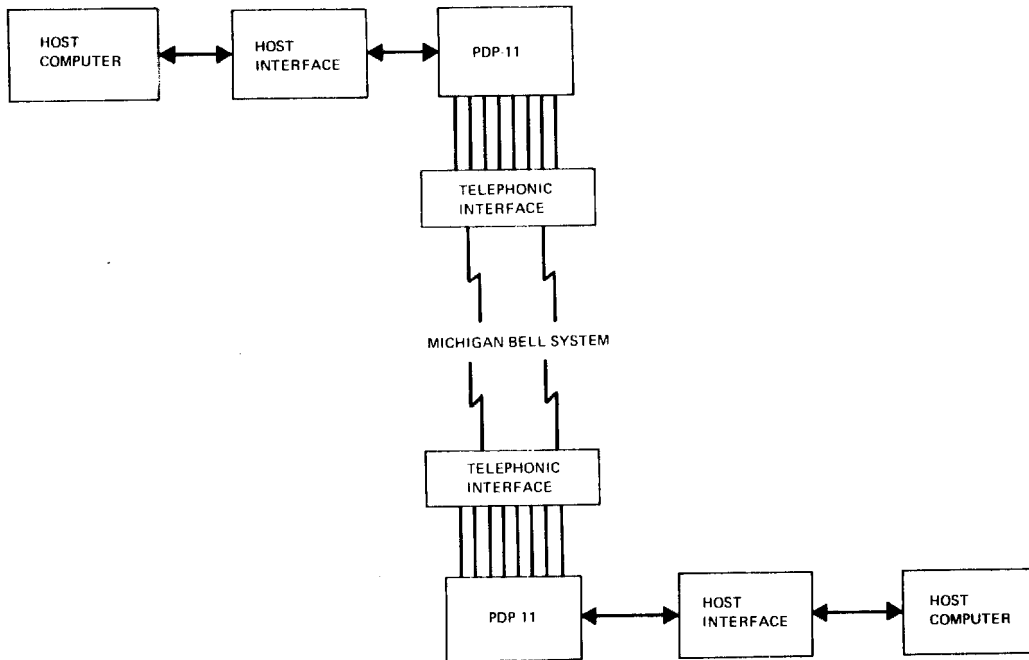


Figure 13 MERIT Communications Segment

The heterogeneous composition of the network has required the development of two philosophically similar but functionally different host interfaces, one for the IBM equipment, the other for the CDC system. The IBM interface attaches to a 2870 Multiplexer Channel and transmits data on eight parallel lines at rates of up to 480,000 bps. The CDC interface, on the other hand, couples the CC with the CDC 6500 Data Channel and its associated Peripheral Processor. Transmission is achieved over twelve parallel lines at an expected rate in excess of 3,000,000 bps.

The central element of the communications subsystem is the CC. As Figure 14 shows, the CC is a PDP-11/20 computer with 16K 16-bit words of core memory, augmented with interfaces that allow interconnection to the host computer and the telephonic network. The primary responsibility of the CC is to allocate its resources among messages to be transmitted, delivered, and forwarded. Software support is derived from the Communications Computer Operating System (CCOS), a specially developed multitasking monitor operating on the PDP-11. The present configuration of CCOS requires approximately 8K words of core memory; the remaining 8K words are used for message and message queue storage.

Upon receipt of a message from the host interface, the CC translates the local host character string into a standard ASCII code (unless the original message was in

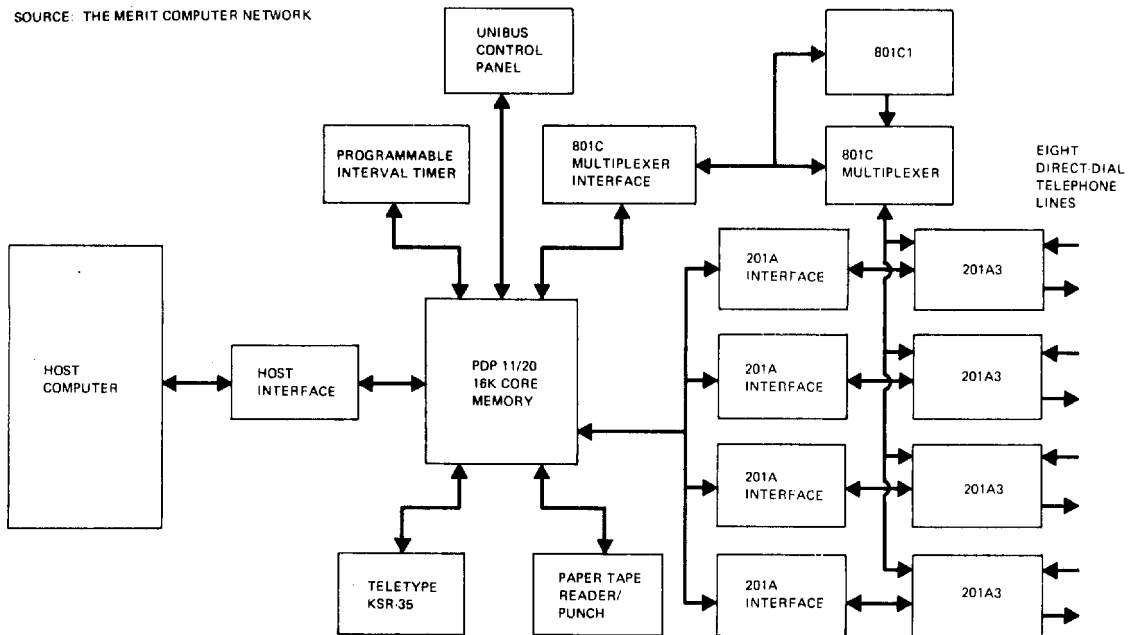


Figure 14 Communication Computer System

binary, eliminating the need for this operation). A message header is generated by CCOS and a 16-bit checksum is computed and checked by the CC hardware. The message is stored, and a transmission queue entry is generated. The order in which the queue is emptied and the physical link over which transmission takes place are subsequently determined by a CCOS task.

Each CC is capable of receiving messages from the others. In this event, a determination of whether the message was received free of errors is made, using the message checksum. If the message was error-free, an acknowledgment is returned to the transmitting CC, allowing it to release its record of the message. If errors were detected, a request for retransmission is returned in lieu of the acknowledgment. Upon receipt of an error-free message, the receiving CC determines whether the message is for its host or is to be forwarded. If the former, the message is queued for host interface activation and subsequent transfer to the host memory; otherwise, the message is queued for transmission toward its destination.

The telephonic network comprises the physical transmission medium and its termination equipment. The MERIT network employs voice-grade, dial-up lines exclusively. Some economy in line costs is achieved by sharing the existing tri-university Telpak lines on a competitive basis with normal voice traffic. Each site supports four Western Electric 201A modems, operating asynchronously in four-wire, full-duplex

mode at 2000 bits per second. Dial-up connections are made by a Western Electric 801C Data Auxiliary Set, which is multiplexed among the four modems. Because the modems operate in a four-wire configuration, the 801C is designed to allocate lines in pairs for each modem. Moreover, since the 801C is completely controlled by the PDP-11 software, it is possible to change the number of lines between two sites in accordance with the current traffic volume, achieving an optimum cost/performance tradeoff within the constraints of available bandwidth.

Usage

MERIT is seeking knowledge of the problems and solutions of operating a network in an established educational computing environment; through the development and implementation of a network, they expect to make contributions to computer and educational technology. MERIT management feel that a network linking the computers of the participating universities will have a synergistic effect on the combined resources of the separate computing facilities. Connecting machines with significantly different system characteristics enables the user to take advantage of the computer best suited for his work. For example, the University of Michigan's system was especially designed for time sharing; it will be available to those at other nodes needing a time-sharing capability. Because of its speed, the CDC 6500 at MSU is well suited for compute-bound jobs; once it is connected to the network, personnel at other universities will be able to take advantage of its facilities. Interconnecting computer systems can make possible a cooperative policy for acquiring some of the more unusual peripheral devices.

The MERIT Network is designed to provide a vehicle for a rapid exchange of information that would not be possible otherwise and to bring researchers in closer contact with those doing similar work at different locations, thereby eliminating much duplication of effort.

MERIT will provide remote service that will be transparent to the user; his job will look like a standard batch job except for the addition of a few network routing commands. MERIT feels that load sharing is infeasible on a per program basis.

Ultimately MERIT hopes to provide a service whereby real-time terminal users will be able to concurrently control programs on two or more host systems. This "dynamic communication" would enable the user to control this process, operating the programs simultaneously or sequentially and transferring data between them. Dynamic communication will facilitate "dynamic file access," the ability of a user at one node to access a file at another node without transferring the file to the user's node. MERIT feels that implementation of this capability will be difficult.

A standard data description language has been proposed by the Michigan Interuniversity Committee on Information Systems (MICIS) to facilitate transmission of data between computers, systems, and programs and to provide a convenient and complete format for the storage of data and its associated descriptor information. MICIS proposes a standard data set composed of two parts:

- a directory describing the logical and physical storage of the data and the properties of its variables; and
- the data matrix.

The directory is to be written in a standard character set, facilitating maximum transferability between various character codes. This restriction does not apply to the actual data described by the directory, however; data can be highly machine dependent, although its description is written in a standard character set.

The current plan is that data will be converted to ASCII prior to being transmitted over the network; upon receipt by the object node the data will be converted to a compatible form for processing on the target hosts. Programs and data must be transmitted in a form acceptable to the target host. MERIT feels that the network will eliminate the need for physical program transferability and that all users can share programs that exercise special features offered by a node but have not been written in a computer independent manner.

Management

The participants in the MERIT network are independent universities, each vying with the others for students, faculty, and research grants. One of the primary goals of the network is to maintain this autonomy at the maximum level consistent with effective network operation. Consequently, each of the universities is responsible for access authorization, resource accounting, systems development, and local hardware expansion. Communications facilities, intercomputer protocol, and similar aspects of the network are the proper concerns of the network management.

Network management is vested in the MERIT Central Staff, comprising a director, an associate director from each university, and a small technical staff. The Director is appointed by the Michigan Interuniversity Committee on Information systems (MICIS), the predecessor of MERIT, which is composed of representatives from each of the three participating institutions. Each associate director's position is filled by nominations from the university, selection by the Director, and approval by MICIS.

The Director is responsible for the technical development of the network and for the administration of its fiscal resources. He relies on his associate directors to insure that the implementation at each university is proceeding on schedule. The associate directors are also responsible for promoting and encouraging network activities at their respective institutions. Moreover, each associate director acts as a liaison between MERIT and his university, to ensure that the university's interests are equitably served with respect to the demands placed upon its resources.

The distribution of system documentation throughout the user community is the joint responsibility of MERIT and the individual universities. At the present time, MERIT disseminates information relevant to the design and operation of the communications subsystem and its interfaces. Each university is required to maintain and distribute its local facilities documentation and is responsible for issuing notices reflecting any significant changes.

The MERIT staff is developing procedures to closely monitor the performance of the network. Statistics gathered on message errors, traffic distribution, and overall throughput will significantly help in adapting the original network design to actual usage patterns. Moreover, a study of machine utilization should facilitate the development of an equitable interuniversity rate structure.

NETWORK/440¹

Introduction

Network/440 is an experimental network sponsored by the Watson Research Center of the IBM Corporation located at Yorktown Heights, New York. Its primary purpose is to facilitate the study of computer networks and to provide a vehicle for an experimentation program in network applications.

The network is currently operational using a 360/91 MVT region as a central switch. This architecture was chosen because of the ease with which performance statistics could be gathered. However, because of the inherent disadvantages of the centralized topology, the network is to become distributed.

Configuration

At the present time Network/440 is a centralized network of homogeneous computers as shown in Figure 15. The grid node of this network is a region of an

¹This network is no longer operational; it is included for historical purposes.




Figure 15 Logical Structure of Network/440

IBM 360/91¹ running under MVT. This node acts as a central switch for the store-and-forward communications, presently being carried out over 40,800 bps leased lines. The present and expected nodes in the network are listed in Figure 16.

LOCATION	MACHINES
IBM WATSON RESEARCH CENTER	360/916MVT (CURRENTLY IN THE NETWORK)
IBM WATSON RESEARCH CENTER	360/67/CP (CURRENTLY IN THE NETWORK)
IBM (BOULDER, COLORADO)	360/65/MVT
OTHER IBM INSTALLATIONS	360/MVT
NYU	CDC 6600
YALE	360/44
IBM (SAN JOSE, CALIFORNIA)	360/91/65/65

Figure 16 Nodes in Network/440

Standard OS/360 software is available to the user over the network.

¹This is the same 360/91 that is linked to the TSS Network.

Communications

Network/440 is a centralized network, comprising a grid node which performs all of the communications support functions, and a set of transmission links which extend radially outward from the grid node to the host computers. The transmission links are leased wide-band half-duplex lines operated at 40,800 bps using Western Electric 300-Series modems. Computer terminations are provided by IBM 2701 Data Adapters connected to 2870 Multiplexer Channels. The links are operated using the standard IBM Basic Telecommunications Access Method (BTAM).

Special communications capabilities are provided by a problem program operating in a single region of a 360/91 MVT system. The program comprises six primary segments, performing network control, operator interface, error recovery, line handling, message queue management, and transaction recording functions. The network control segment is responsible for handling user jobs and decoding appropriate network control messages. The operator interface handles messages going to and from the central machine operator. The error recovery segment is responsible for retransmitting messages which were lost or garbled, and for attempting to re-synchronize the lines after a line loss. The line handler provides the interface with the BTAM software for forwarding messages to the host computers. The message queue manager is responsible for queuing messages in core for forwarding if the target host is available, or on a disk if not. In this way, a host will always receive its messages whether or not it is operational at the time the message is sent. Finally, the transaction recording segment maintains an audit tape of all message traffic passing through the central switch.

Usage

Network/440 is a research project being used to gain a better understanding of computer networking; for this reason the centralized approach was taken in its design. The grid node monitors all messages passing through the network and makes network measurements more easily than would be the case on a distributed system.

Load sharing, data sharing and program sharing are possible over the network. The grid node provides a centralized catalog of all data sets available for network usage, but each node maintains control over its own data sets. One of the more important functions of the network is transferring data sets. This currently requires the user to spell out exactly what he is referring to when manipulating files. Current plans call for making these operations more user transparent. Network/440 is currently a batch-oriented network with plans to offer interactive facilities in the future.

Network/440 has developed several control languages, each providing the user with more capability and flexibility in a less machine-oriented form. Planned expansions of this control language include the following:

- grid node conversion of local job control language into the language required by the target computer or grid node mapping of one job control language into the target machine's job control language;
- grid node conversion of floating point numbers, integers, and character strings from one machine structure to any other; and
- automatic job scheduling to achieve load leveling among like machines or, by job control language conversions, among unlike machines.

IBM's concern about network usage of proprietary data has prompted the development of a grid node usage matrix that maintains a list of resources available to a specific user. Additionally, a node may disconnect itself from the network to process proprietary data; if this occurs, incoming messages are stored until the node is reconnected.

Management

Because of the nature of Network/440, no formal management structure exists. The network is administered as a research project of the IBM Corporation.

THE OCTOPUS NETWORK

Introduction

The Octopus network is a highly integrated system providing operational support for the research activities of the Lawrence Radiation Laboratory (LRL). The network was developed and is operated by LRL under the auspices of the United States Atomic Energy Commission. The computation requirements of LRL have necessitated the use of several large computer systems; the purpose of the Octopus network is to integrate these systems into a unified computer complex. In satisfying this responsibility, Octopus performs two primary functions:

- it provides the user with a single point of access to the several computers; and
- it allows each of the computers to access a large centralized data base.

Octopus was first conceived in the early 1960's and became operational in 1964. As the computation center has grown in size and complexity, Octopus has been expanded to meet changing needs. At the present time, the system services about 300 remote terminals, four main computers, and a trillion-bit data base.

Configuration

Octopus is a heterogeneous network of computers including two CDC 6600's, two CDC 7600's, and, in the future, a CDC STAR. Each of these worker (or host) computers is operated in a time-sharing mode and is linked to a central system providing two features:

- a centralized hierarchy of storage devices (a centralized data base) shared by the worker computers; and
- a provision for various forms of remote and local input and output, permitting the network to be viewed as a single computing resource.

Octopus uses a store-and-forward communications protocol. Communications lines between workers are 12 megabit, hardwire cables. Figure 17 gives a graphic description of the Octopus system.

Octopus can be more easily visualized as two independent, superimposed networks:

- the File Transport Subnet which is a centralized network consisting of worker computers, the Transport Control Computer (a duplex PDP-10 which serves as the grid node), and the central memory system (disk, data cell, Photo-store); and
- the Teletype Subnet (Figure 21) which is a distributed network consisting of worker computers, three PDP-8's (each servicing up to 128 teletypes) and the Transport Control Computer (the duplex PDP-10).

A third network, not yet installed, will comprise remote I/O terminals supported by duplex PDP-11's. While the networks can be considered logically independent, they are interconnected to provide alternate routes for data; for example, the PDP-10 provides an alternate path between the interactive terminals and the worker computers. Figure 18 shows some of the major hardware components in the Octopus system.

The Octopus network also supports a Television Monitor Display System (TMDS), shown in Figure 19. The purpose of TMDS is to provide a graphic display capability, with monitors distributed throughout the facility. Bit patterns representing characters and/or vectors are recorded on the fixed-head disk, which operates at a synchronous speed compatible with the standard television scan rate. Sufficient storage is available on the disk to store 16 rasters of 512 x 512 black or white picture points. The addition of a crossbar switching system will allow a particular raster to be displayed on several monitors simultaneously.

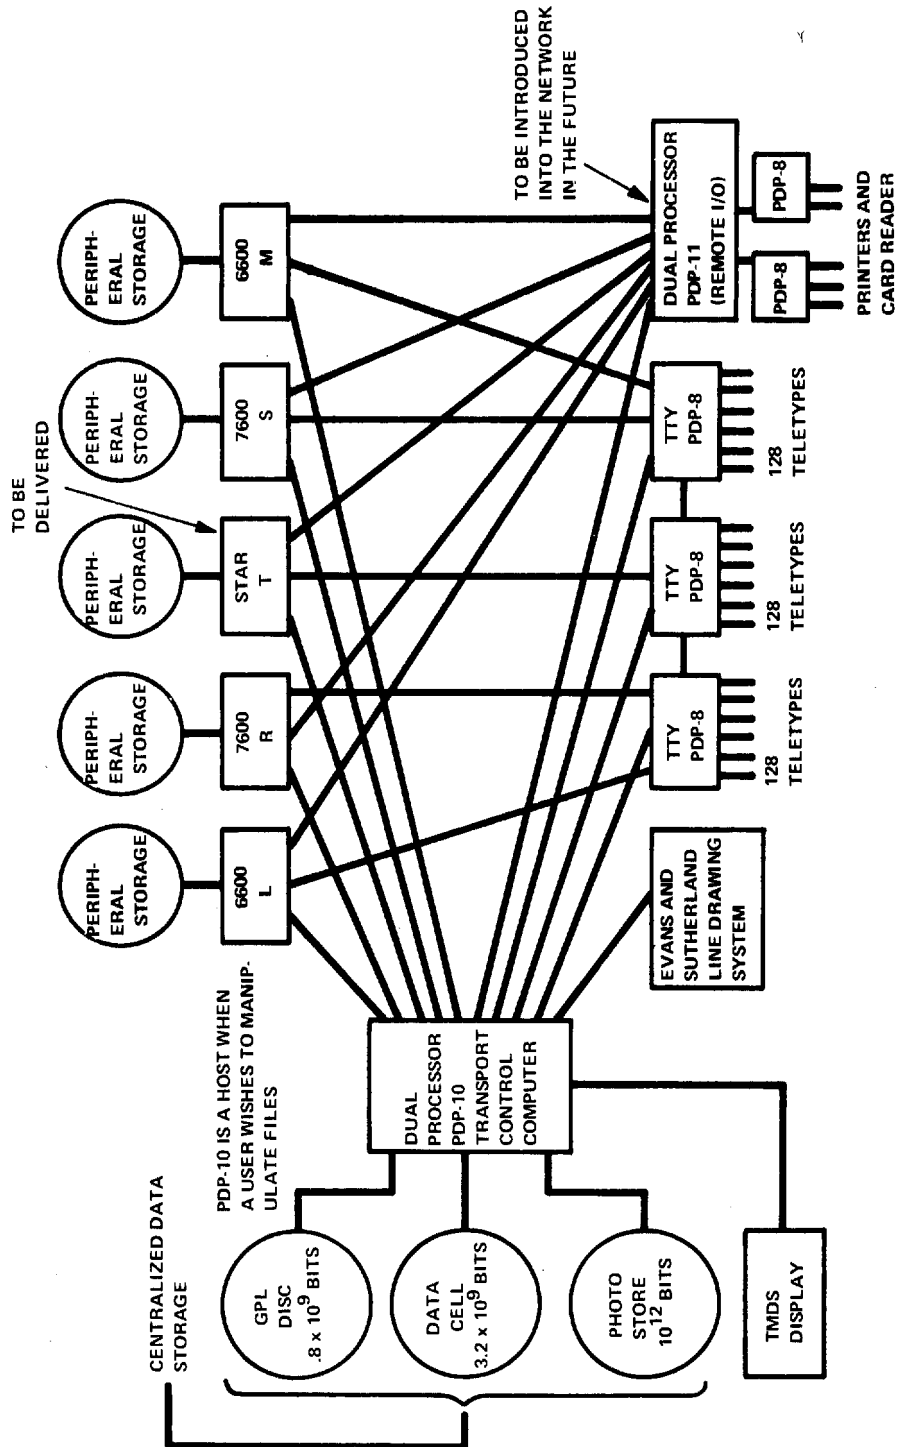
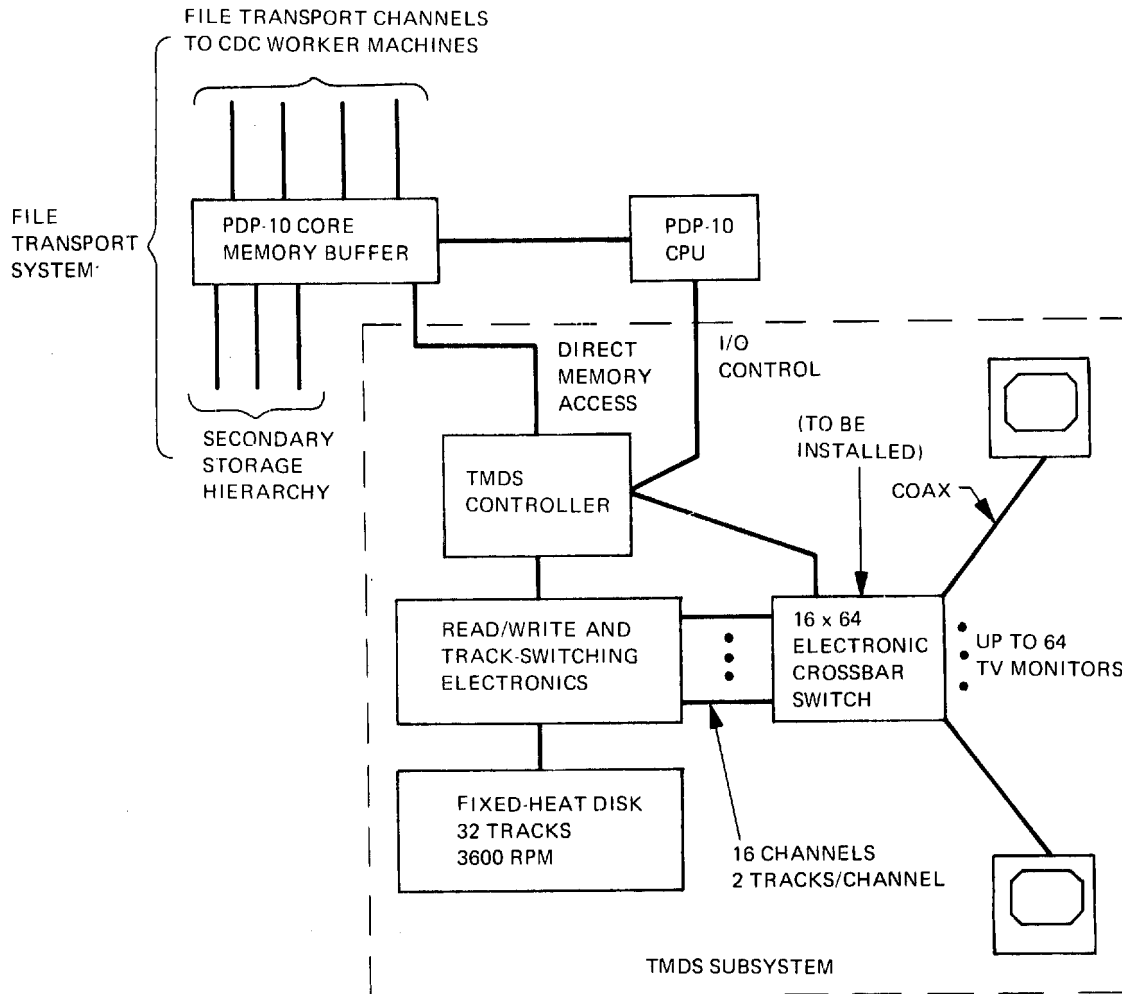


Figure 17 The Octopus Network

PROCESSOR	CORE SIZE	OTHER HARDWARE
<p>WORKER COMPUTERS¹</p> <p>6600 L</p> <p>6600 M</p> <p>7600 R</p> <p>7600 S</p> <p>START</p>	<p>128K</p> <p>128K</p> <p>65K SMALL CORE 500K LARGE CORE</p> <p>65K SMALL CORE 500K LARGE CORE</p> <p>500K (64-BIT) WORDS</p>	<p>225 MILLION 60-BIT WORDS (DISK))</p> <p>225 MILLION 60-BIT WORDS (DISK))</p> <p>167 MILLION 60-BIT WORDS (DISK) .8 MILLION 60-BIT WORDS (DRUM))</p> <p>167 MILLION 60-BIT WORDS (DISK) .8 MILLION 60-BIT WORDS (DRUM)</p> <p>167 MILLION 64-BIT WORDS (DISK)</p>
<p>TRANSPORT CONTROL COMPUTER</p> <p>PDP-10 (DUAL PROCESSOR)</p>		<p>IBM 2321 DATA CELL (3.2 x 10⁹-BIT CAPACITY)</p> <p>IBM 1360 PHOTOSTORE (10¹²-BIT CAPACITY)</p> <p>GENERAL PRECISION LIBRASCOPE DISK (8.8 x 10⁸-BIT CAPACITY)</p> <p>TMDS (TELEVISION MONITOR DISPLAY SYSTEM)</p> <p>EVANS & SUTHERLAND LINE DRAWING SYSTEM</p>
<p>PDP-8 CONCENTRATORS</p> <p>3 PDP-8s</p>	<p>8K</p>	<p>UP TO 128 TELETYPES ON EACH PDP-8</p>
<p>REMOTE I/O²</p> <p>PDP-11 (DUAL PROCESSOR)</p>		<p>TERMINALS TYPICALLY CONSISTING OF A PDP-8 CONNECTED TO I/O DEVICES (READERS AND PRINTERS)</p>
<p>¹THE LETTER TO THE RIGHT OF THE WORKER COMPUTER IS AN INTERNAL LRL DESIGNATION. THESE LETTERS ALSO APPEAR IN FIGURE 17.</p>		
<p>²NOT YET OPERATIONAL.</p>		

Figure 18 Octopus Hardware



SOURCE: PEHRSON, D. L., "AN ENGINEERING VIEW OF THE LRL OCTOPUS COMPUTER NETWORK," NOVEMBER 17, 1970, P. 24.

Figure 19 Television Monitor Display System (TMDS)

LRL has designed and built much of their hardware and almost all of their software (including the operating systems for their computers); for example, they have a special multiplexer enabling the PDP-8's to handle 128 teletypes each, whereas DEC permits a maximum of 32. LRL has implemented their own versions of COBOL, FORTRAN (called LRLTRAN), APL and PL/I. They are currently developing an APL compiler (their current version of APL is interpretive). In addition, they provide CDC FORTRAN, SNOBOL, debugging routines, a text editor, LISP and linear programs.

Communications


The File Transport Subnet connects the PDP-10 system and its central data store to the worker computers. Because of the inherent differences between the CDC 6600 and the CDC 7600, two distinct file transport channels have been developed, one for each machine type. However, there are two channel characteristics which are identical in both cases. The first of these is the PDP-10 interface, which uses a data demand protocol to synchronize the transmission rate between the computers. The second is the maximum channel transfer rate, which is about 10 million bits per second.

The 6600/PDP-10 file transport channel is shown in Figure 20. The principal components involved in the transmission process are as follows:

- the 6600 Peripheral Processor Unit (PPU), a 12-bit, 4K-word programmed I/O processor;
- a Channel Switch, which connects one of the ten available PPU's to one of the twelve available data channels;
- the 6000-Series Data Channels, which transfer data on 12-bit parallel lines;
- the Adapter unit which interfaces the 12-bit CDC channel to the standard LRL Octopus Data Channel, a 12- or 36-bit wide transmission system;
- the LRL Data Channel, which performs synchronized, half-duplex digital transmission; and
- the PDP-10 and its channel interfaces.

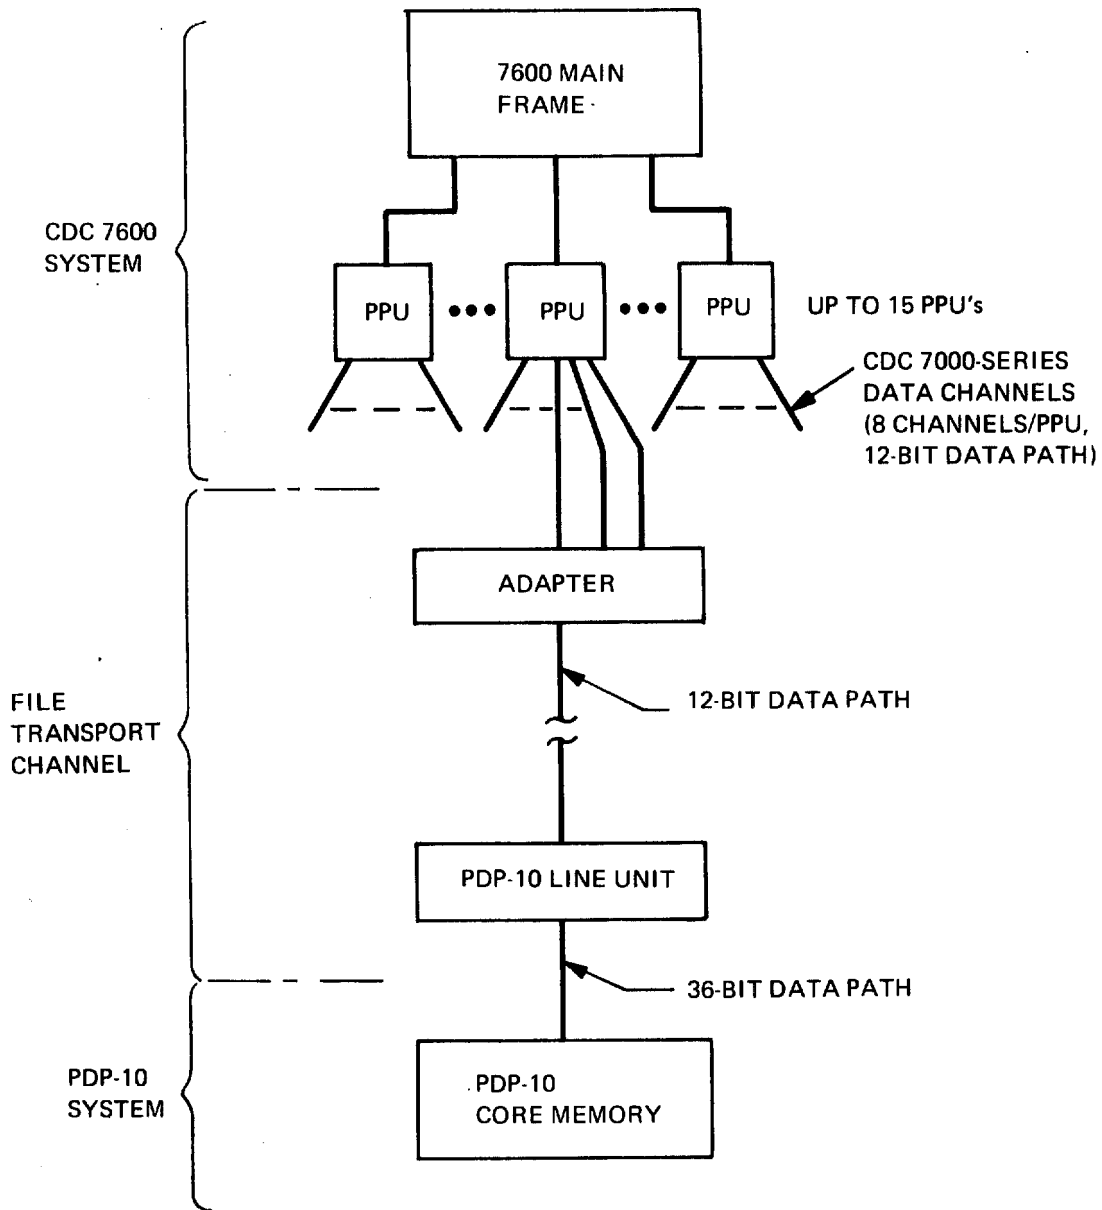
The operation of the file transport channel is initiated by a request from the 6600, requesting either a read from or a write to the central data store. These requests normally involve the transfer of a complete file, with the average transmission comprising more than 500,000 bits of data. A processor dialog is subsequently established to transfer data between the two computers. At the PDP-10, data is transferred between the GPL disk and core, and then between core and the transmission line. The 6600 uses two PPU's in a buffer-switching scheme, alternating between the transmission line and the local disk, bypassing the main core of the 6600.

As Figure 21 illustrates, the 7600 configuration is somewhat different. Each of the fifteen PPU's on the 7600 has eight dedicated channels available, eliminating the need for the channel switch. Moreover, because the 7600 CPU and PPU's are much faster than those of the 6600, a more classical transport protocol is used, with the PPU acting as a programmable interface between the PDP-10 and the 7600 CPU, which controls the data transfer.



SOURCE: PEHRSON, D. L., OP CIT, P. 14.

Figure 20 6600/PDP-10 File Transport Channel



SOURCE: PEHRSON, D. L., OP CIT, P. 15.

Figure 21 7600/PDP-10 File Transport Channel

The Teletype Subnet, shown in Figure 22, is designed to efficiently route short messages of approximately 80 characters between user terminals and the worker computers. Although this subnet functions independently of the File Transport Subnet, the two are interconnected, primarily for enhanced reliability.




Figure 22 Octopus Teletype Subnet

The central element of the Teletype Subnet is the PDP-8 computer. Each PDP-8 has 8K 12-bit words of core memory and is capable of supporting up to 128 full-duplex Teletype terminals, operating at 110 bps. A special operating system has been developed for the PDP-8 to support the Teletype Multiplexer, to manage core buffers, and to forward messages in the subnet. This system requires 4K words of core memory, leaving 4K words for line message buffers.

Each PDP-8 accepts characters from its terminal until a complete message has been formed. If the message destination is a worker which is directly connected to the PDP-8, the message is transmitted using links similar to those described in the File Transport Subnet discussion, but operating at about one-tenth the speed. If the worker is not directly connected, the message is forwarded to a neighboring PDP-8, where a similar process is repeated. An analogous protocol is followed for output messages traveling from a worker computer to the user terminal.

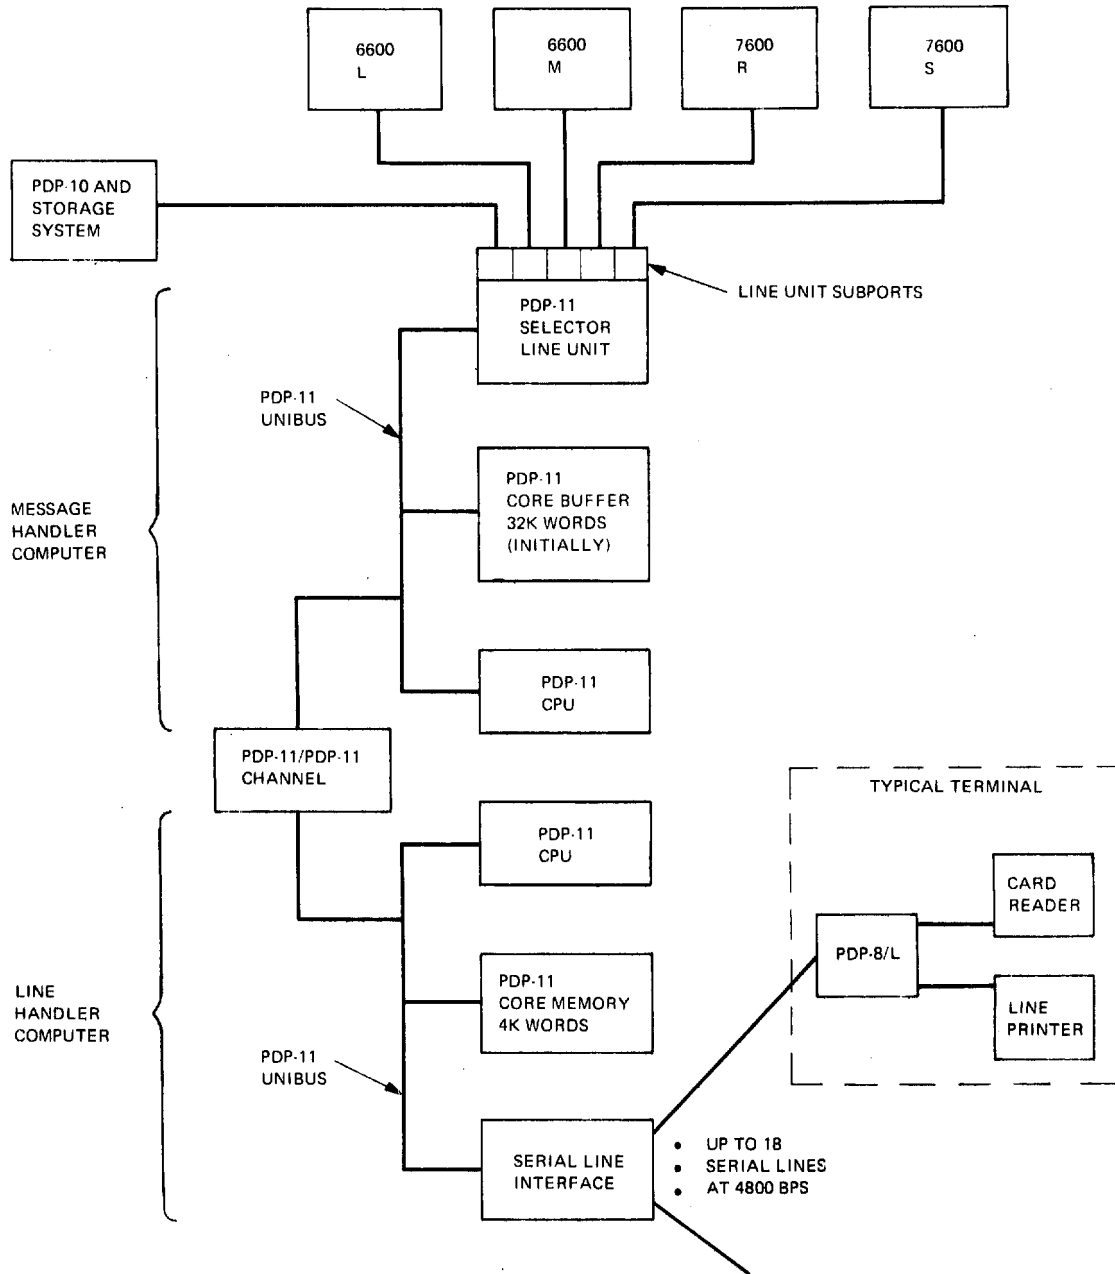
In the event that a PDP-8/worker link becomes inoperative, messages can be forwarded to the affected worker computer via the File Transport Subnet. Although the intermixing of short Teletype messages with long file transfers does downgrade system performance, the enhanced reliability that is achieved is adequate compensation.

A Remote Job Entry Terminal (RJET) Subnet is currently under development for inclusion in the Octopus Network. Its purpose is to provide a capability for card reader input and line printer output at remote sites throughout the facility. The proposed RJET configuration is shown in Figure 23. The controlling element of RJET is a pair of PDP-11 computer systems. One PDP-11 performs a role similar to that of the Teletype Subnet Computer, routing messages between the workers and local buffer areas. The other PDP-11 acts as a line handler, providing interface capabilities for eighteen 4800 bps, half-duplex terminal lines.

Usage

The Octopus Network has increased the overall effectiveness and efficiency of the computing facilities provided by LRL's large computers. The multicomputer complex is treated as a single resource, enabling all terminals to access all worker computers and providing the Octopus user with several advantages:

- easy accessibility to any worker computer from any teletype terminal;
- man-machine interaction with a high-speed computer while executing a program; and
- rapid turnaround time during debugging.



SOURCE: PEHRSON, D. L., OP CIT, P. 28.

Figure 23 Remote Job Entry Terminal (RJET) System and Network Connections

Most of LRL's computer work consists of long-running, compute-bound problems requiring many hours of computer time. Of their 1000 users, 20 to 40 are generally on line at one time, running concurrently with batch background jobs. User-controlled data sharing, program sharing and remote service are possible on the network; load sharing, however, is hampered by the fact that the worker machines are different.

Interactive multi-programming on LRL's giant computers generates a requirement for massive on-line storage. Tapes are inefficient in this type of environment, and for this reason the concept of the shared data base has been employed. A hierarchy of storage is composed of a Librascope fixed-head disk (807 million bits, rapid access, high transfer rate) and an IBM Data Cell (3.24 billion bits), both supporting the IBM Photostore (over one trillion bits), the major media for mass storage. Economics and flexibility make the sharing of these storage devices advantageous.

The large-capacity Photostore provides an economical means of storing long-life files; such a large storage device is reasonable only if it is shared by several large systems. Writing the Photostore is a time-consuming activity and it is therefore not amenable to files that change frequently. The storage hierarchy balances and smooths I/O loads in supporting the Photostore and also provides an indexing mechanism for this device. The shared data base concept instills flexibility and operational advantages into the system since files transported to the Photostore by one worker system can be subsequently accessed from another worker system, eliminating the need for unique copies of public files on each worker system.

The PDP-10 Transport Control Computer and the appropriate worker computer handle file transport. Data is copied from a file controlled by the PDP-10 and written into a file controlled by the worker computer; the source file is not altered or destroyed, although it can be rewritten while in the worker computer.

Maintaining a centralized data base has some disadvantages: since all worker computers depend on the shared storage hierarchy, reliability requirements are greatly increased; and a major effort is required to implement the centralized file storage subnet.

File access codes enable a file to be read by others, but written only by those with the correct access code. Worker computers have their own access codes which may inhibit file transport in some cases. Various types of files include the following:

- private files accessible to one user,
- shared files accessible to a group of users, and
- public files accessible to all users.

Each user is identified with a maximum security level at which he is permitted to operate; these include: unclassified, protected (data cannot be carried off the site), administrative, and secret. Each file and I/O device is identified with an operating level, and a user's access to them must not exceed the maximum operating level allowed him. No Top Secret work is done on the Octopus network.

Management

The Octopus network was developed to provide computer support exclusively for LRL. Consequently, its management is centralized, vested in the Computation Department of LRL. The Computation Department, which comprises over 300 staff members, is managed by a Department Head and a Deputy Department Head, supported by three Assistant Department Heads for Administration, Research, and Planning.

Octopus is managed in a fashion similar to that of any research computational center. Management is responsible for acquiring, developing and maintaining hardware and software; authorizing system access; allocating computer resources; and assisting the user community in achieving effective computer utilization.

The applications programming support functions of the Computation Department are necessarily extensive and varied. Applications programmers, working in one of six main groups, provide programming support throughout the Laboratory in areas including administrative data processing, engineering, physics, medicine, and nuclear weapons research.

In addition to the applications programming staff, the Computation Department maintains six project groups, each tasked with a specific support role. These groups, and their respective functions are as follows:

- the Systems Development Section, which designs and develops all of the systems software for the network;
- the Systems Operation Section, which performs software maintenance and consultation services;
- the Numerical Analysis Group, which designs, develops, and evaluates mathematically-oriented computer algorithms;
- the Computation Project Group, which engineers additions and modifications to the network hardware;
- the Computer Information Center, which obtains, edits, writes, publishes, and distributes all system documentation; and
- the Computer Operations Section, which is responsible for the operation of the computer systems.

THE TSS NETWORK¹

Introduction

The TSS network was first conceived in 1967 when Princeton, IBM Watson, and Carnegie-Mellon University decided to interconnect their computer facilities. The purpose of the network is to advance research in the applications of computer networks, particularly in the areas of cooperative system development.

The network is currently operational at all of the nodes. Moreover, experimentation programs are well underway at several sites, particularly those on the east coast.


Configuration

The TSS network is a distributed network of homogeneous (IBM 360/67) computers using the TSS/360 operating system. Each node manages a local network of heterogeneous computers including some large 360's running under OS; these processors appear as devices to the network. Nodes are located at IBM Watson Research Center, Carnegie-Mellon University (CMU), NASA Lewis Research Center, NASA Ames Research Center, Bell Telephone Laboratories (Naperville), and Princeton University. The nodes are interconnected by 2000 bps voice-grade auto-dial lines and 40,800 bps leased lines. Figure 24 presents an overview of the nodes participating in the TSS Network, and Figure 25 presents configuration information for each node. Additional facilities that may become TSS nodes are Chevron Oil Corporation and Northern Illinois University.

Modifications to the TSS/360 operating system were necessary to enable processors to initiate tasks on and communicate with other processors; one processor appears as a terminal to another. A major consideration was given to minimizing these modifications. By using like processors and software many of the usual obstacles of network design were avoided. TSS/360 provides enough flexibility for expansion, and the modular network design allows for the inclusion of other operating systems in the future.

Languages available over the network include PL/I, FORTRAN H, ASSEMBLER H, ALGOL, SNOBOL 3, SNOBOL 4, APL, BASIC, WATFOR, LISP, CSMP, GPSS, JOSS, LC², and LEAD. Other software includes NASTRAN, a structural analysis program, and

¹Time Sharing System—a time-sharing operating system for the IBM 360/67 computer.



SOURCE: IBM WATSON RESEARCH CENTER

Figure 24 An Overview of the TSS Network

a text editor. In addition, software is available to convert FORTRAN source code from the TSS format to the OS format automatically, for example between the 360/67 with the TSS operating system and IBM's 360/91 with OS/MVT.

NASA Lewis will be the Network Information Center for the TSS Network. They will keep records on machine configurations and available programs, will maintain up-to-date source code for the network software, and will keep a history of usage requests, identifying the user and the reason for the request. All changes to programs available over the network will be recorded for other users.

Communications

Communications among the 360/67's of the TSS network are carried out using voice-grade switched lines operating at 2000 bps. The lines are driven by Western Electric 201A modems in a half-duplex configuration. The 360/67 interface is provided by an IBM 270 or 2703 connected to the 2870 Multiplexer Channel.

Because of the lack of programmable interface hardware, all communications software support is resident on the host computer. In order to avoid extensive changes

NODE	NODE CONFIGURATION	LOCAL NETWORK CONFIGURATION
IBM (WATSON RESEARCH CENTER)	360/67/TSS	IBM 360/91/MVT, 1800, 1130, SYSTEM 7
CMU	360/67/TSS	EAI/680-PDP-9, UNIVAC 1108, PDP-8, PDP-11, DDP-116, PDP-10 ¹
NASA LEWIS	DUPLEX 360/67/TSS	SMALL COMPUTERS (XDS, DEC), ON-LINE CDC MICROFILM UNIT, 2321 DATA CELL, 2 2301 DRUMS, 2 2314 DISK UNITS, 3 MAINCORE UNITS (256K EACH), CALCOMP PLOTTER, SENSOR EQUIPMENT.
NASA AMES	DUPLEX 360/67/TSS	3 2314s, 2 2301 DRUMS, 6 2780s, 1800 WITH A 2250, IMLAC PDS 1, SC 4020, 30 TERMINALS (2741s, TTYs)
BTL/NAPERVILLE	DUPLEX 360/67/TSS	IBM 360/65/65/50 ASP
PRINCETON	360/67/TSS	IBM 360/91/MVT

¹PROPOSED.

Figure 25 TSS Network Hardware

to the standard TSS operating system, the communications software was designed to operate as a user program, contending for resources on the same basis as other user programs. As a consequence, the highest message throughput capacity which has been realized is less than the maximum possible with the present communications hardware.¹

¹In recognition of this problem, IBM is developing a communications computer concept which would use a 370/145 as a combination communications computer and data base manager.

In order to provide the user with access to the communications subsystem, the TSS network employs the Computer Access Method (CAM), a specially developed set of procedure calls and software which effect the intercomputer dialog. CAM is capable of supporting voice-grade lines operating at 2000 bps, and wide-band leased facilities operating at either 40,800 or 50,000 bps. A more generalized version of CAM, called Table Driven CAM, has also been developed. In Table Driven CAM, the characteristics of the communications subsystem and the receiver are defined by means of table entries, permitting a wider range of computers and communications equipment to be used on the network.

Upon receipt of a CAM request, the communications software must first determine whether a connection exists to the destination computer; if not, one is established. The message, which may be up to 1024 bytes long, is compacted prior to transmission. The message is subsequently transmitted to its destination by a special software task, which time-multiplexes all messages destined for any particular site. The receiving system software effectively performs the same process in the reverse order. Error checks are performed; retransmissions are requested if errors occurred, while acknowledgments are returned otherwise.

Usage

The first goal of the TSS network was to investigate the uses and advantages of a general purpose network of computers. The experience gained is to be used in determining future avenues of expansion in designing and implementing other networks. The nodes will use the network for experimentation and research rather than for production work.

The TSS Network provides a convenient means of exchanging programs and system modifications since like computers are used in the network. Use of the network for program sharing and data sharing saves duplication of programs and data at foreign sites. Load sharing, remote service and dynamic file access are among the features provided by the TSS Network. Figure 26 gives an example of how the network is used.

Since one processor appears as a terminal to another and since all devices in the network appear to be on each processor, the terminal user can command the full resources of the network as though he were dealing with a single system. After the user has gained access to the target processor, he may initiate processing activity, disconnect, or connect to another node; he may even have many jobs executing at various nodes simultaneously. Specialized facilities, such as graphic subsystems and large core memories, are available over the network.




Figure 26 Usage of the TSS Network

Passwords and keys serve to maintain the privacy and integrity of the user files. As an added precaution, however, IBM does not allow outsiders to connect to their 360/91 when proprietary information is being processed. The network has a "copy protect" feature that enables a person to use a file or a program without allowing him to copy it.

A Network Control Language enables the user to perform the following functions: connect to a specified node, initiate a computational process, disconnect from a specified node, test for any outstanding responses, and send and receive data sets, and display process responses. The language is simple since the designers concentrated on making the system easy to use.

Management

The TSS network is an interconnection of several independent research facilities. A consequent goal of the network is the establishment of an experimental environment which interferes with the other activities of the nodes as little as possible. The homogeneity of the network has been instrumental in establishing this environment by minimizing the amount of effort required to develop all of the network software.

The technical development of the network has been carried out informally. Representatives from each of the sites meet periodically to discuss technical proposals and ideas. Planned experiment activities are also discussed and coordinated at these meetings.

The collection and dissemination of network documentation is the responsibility of the Network Information Center (NIC), located at the NASA Lewis Research Center. Systems which will permit retrieval of appropriate documents by a network user are currently under development at the Lewis Center.

No formalized procedure has been developed for intersite billing. At the present time, the informality of the project and the nearly equal intersite utilization of resources has obviated the need for such procedures. However, usage statistics are gathered to monitor this situation and to uncover any heavy one-sided usage patterns.

THE TUCC NETWORK

Introduction

The Triangle Universities Computation Center (TUCC) was established in 1965 as a cooperative venture among three major North Carolina universities: Duke University, North Carolina State University (NCSU), and the University of North Carolina (UNC). Its

incorporation was a response to the saturation of existing local computer facilities and the unavailability of funds to permit the expansion necessary at each of the universities.

The network was developed to satisfy three primary goals:

- to provide each of the institutions with adequate computational facilities as economically as possible;
- to minimize the number of systems programming personnel needed; and
- to foster greater cooperation in the exchange of systems, programs, and ideas among the three universities.


Network operation was begun in 1966. Since that time, a continual growth in both the central computing capability and that at each of the universities has been necessitated. Throughput on the central computer has grown from 600 jobs per day in 1967 to a present peak volume of about 4200 jobs per day. Present plans call for a 100% increase in the central computer capacity by September 1971.

Configuration

The TUCC Network is centralized with homogeneous computers at its three nodes: UNC, Duke and NCSU. Through the North Carolina Educational Computer Services, TUCC also serves some fifty smaller schools within the State and provides general computing services to a small number of research-oriented organizations. Figure 27 gives an overview of the TUCC network.

The center of TUCC is a well-equipped IBM 360/75 with one million bytes of high-speed core and two million bytes of Large Capacity Storage, operating under OS/MVT (see Figure 28). There are approximately 100 terminals (high, medium, and low speed) in the network. The high-speed terminals are a 360/50 and an 1130 at UNC, a 360/40 at NCSU and a 360/40 at Duke. The 360 systems are multiprogrammed with a partition for local batch work and a telecommunications partition for TUCC remote I/O services. The medium-speed terminals are IBM 2780's (or equivalents) and 1130's, and the low-speed terminals are teletypes, IBM 2741's (or equivalents) and IBM 1050's. Less than 10% of TUCC's work is submitted at the card reader at the central computer.

Software facilities provided by TUCC include FORTRAN (E, G, H and WATFIV), ALGOL, APL, COBOL, CPS, BASIC, PL/I, SNOBOL, CSMP, ECAP, GPSS, MPS, FORMAT, TEXT/360, and Assembler G.



NOTE: IN ADDITION TO THE PRIMARY TERMINAL INSTALLATION AT DUKE, UNC, AND NCSU, EACH CAMPUS HAS AN ARRAY OF TERMINALS INCLUDING 2780, 2741, 1050, 1130 AND TELETYPE TERMINALS DIRECTLY CONNECTED TO THE TUCC 360/75.

Figure 27 An Overview of the Tucc Network




Figure 28 Configuration of the 360/75 at TUCC

Communications

The central TUCC computer connects to each local-site computer by means of a single, leased, wide-band line. Operation of the lines is half-duplex at 40,800 bps, using Western Electric 301-B modems. The computer interface is provided by an IBM 2701 Data Adapter connected to the 2870 Multiplexer Channel.

The lack of programmable interface hardware requires that each of the computers provide the necessary communications software support. However, because of the straightforward message flow patterns, the local sites need supply only a relatively simple software interface. The central computer employs HASP, which treats each local computer as a card reader, card punch, and line printer terminal.

The central computer also provides facilities supporting medium speed terminals, operating at 1200 to 2400 bps. These devices employ switched or full-period voice-grade lines, using Western Electric 200-Series modems. The computer interface is provided by a 2703 Data Adapter, capable of terminating 24 such lines. These devices, which are typically IBM 2780's, 2770's, and 1130's, are treated by the communications software in a fashion similar to the high speed devices.

The 2703 Data Adapter also terminates 64 switched voice-grade lines, operating at 100 to 300 bps using 100-Series modems. The low-speed terminals, which include Model 33 Teletypes, IBM 2741's, and 1050's are interfaced to the Conversational Programming System (CPS) or the APL system on an interactive basis, or the remote job entry system.

Usage

The TUCC network realizes substantially more power than would the three sites operating separately because of the economy of a large computer system and the sharing of personnel and programs. The network is used primarily for remote service; jobs are sent from the satellite computers and terminals over communications lines to the central computer, the 360/75, for processing.

Up to 4200 jobs per day can be handled by the network. Although the average running time for a job is 20 seconds, 50 percent of the machine time is devoted to jobs running 4 minutes or longer. The network is used for scientific and instructional work and for some administrative function. The workload includes all types of jobs: small student jobs, large jobs requiring a substantial amount of core, jobs with no set up requirements, jobs with large data bases, compute-bound jobs, and I/O-bound jobs. Turnaround time on batch processing ranges from 5 to 20 minutes for the slow- and medium-speed terminals to 4 hours for the high-speed, computer-based terminals.

TUCC offers interactive service through the Conversational Programming System (CPS) and A Programming Language (APL); both of these systems reside in LCS, except for some frequently used code which is in fast core. Of the six fast core MVT regions, one is used to provide express service for small non-setup jobs that can run in 100K.

Jobs with the shortest running times are generally scheduled first; for more efficient utilization, however, like jobs are scheduled to run concurrently to take advantage of compilers and library routines that have been brought into storage. This is accomplished through a scheduling algorithm that is part of TUCC's software.

Management

TUCC is a non-profit corporation owned jointly by the three participating universities. The corporation is governed by a nine-man Board of Directors, three appointed by the executive officer of each institution. Typically, one of the individuals from each university represents that institution's business interests; another represents the local user community; the third represents the computer sciences interests.

The Board of Directors meets once each month to discuss and dispose of pending matters. Most questions are decided by a majority vote of the members. However, "questions of fundamental importance," which include the selection of the President of TUCC, the annual budget, and major equipment purchases, are decided with each university delegation having a single vote.

Each of the universities maintains its own staff and is required to respond to the problems of its own users. This permits the relatively small central staff to focus its attention on the installation and maintenance of operating systems and documentation in accordance with the needs of the universities. At the same time, this dual staff arrangement allows the user to interface with his own local campus staff, contributing to the political stability of the overall system.

The problems of interuniversity billing have been eliminated in the TUCC billing system. Each university pays 1/3 of the budgeted operating costs of TUCC, plus its own local computer and terminal costs. In exchange for this payment, the universities are assured of receiving equal consideration for the allocation of resources by means of a "usage-levelling" scheduler which allocates resources first to whichever university has used the least. Each of the institutions is then free to bill its users based upon the payment made to TUCC and on the detailed usage statistics gathered at the central computer.

SECTION III MATRIX OF NETWORK FEATURES

This section tabulates the basic features of the surveyed networks. Features are listed on the left side of the chart and the networks are listed across the top. The terms used in this section are defined in the Glossary.

CONFIGURATION

Network topology, hardware, and software are shown. Features are listed only when they have been implemented or when definite plans call for their implementation. Footnotes indicate more tentative plans and further explain Matrix entries.

COMMUNICATIONS

This section presents an overview of the communications characteristics of each of the networks surveyed. Emphasis has been placed on hardware attributes, since these are the major factors determining the overall communications subsystem performance.

NETWORK USAGE

This section indicates network features available to the user and areas of investigation being pursued by the network management.

FEATURES	NETWORKS			
	ARPA	COINS	CYBERNET	DCS
Configuration				
Organization	Distributed	Centralized	Distributed	Distributed
Composition	Heterogeneous	Heterogeneous	Heterogeneous	Heterogeneous
Number of Nodes	19 ¹	3	36	6
Nodes	UCLA SRI UCSB Utah Rand BBN MIT SDC Harvard Illinois Stanford Case CMU Lincoln Labs AWS ² ETAC ² London University ² MITRE/ Washington ² National Bureau of Standards ² NASA Ames ² NCAR ² OCAMA ² RADC ² SAAC ² SAAMA ² University of Southern California ²	Classified	Palo Alto ³ (2) Los Angeles ³ (3) Minneapolis Boston ³ (2) New York ³ (3) Washington ³ (5) Houston ³ (2) Honolulu Atalanta Detroit San Francisco Hartford Seattle Phoenix Albuquerque Omaha ³ (2) Chicago Dallas Cleveland Cincinnati Richmond Baltimore Philadelphia ³ (2)	Nodes located in various buildings on the UCI campus.
¹ Of these approximately 14 are actually connected to the network. ² Not yet connected to the network. ³ The number in parentheses represents the number of computers in the indicated geographical area.				

NETWORKS					
DLS	MERIT	NETWORK/440	OCTOPUS	TSS	TUCC
Distributed	Distributed	Centralized ¹	Centralized ⁵ Distributed ⁶	Distributed	Centralized
Homogeneous	Heterogeneous	Heterogeneous	Heterogeneous	Homogeneous	Homogeneous
2	3	6 ²	10 ⁷	9	4
NMCSSC ANMCC	MSU WSU UM	IBM/Watson ³ (2) IBM/Boulder ⁴ NYU ⁴ YALE ⁴ IBM/San Jose ⁴	All nodes located at LRL.	IBM/Watson Princeton NASA Lewis CMU BTL (Naper- ville) NASA Ames Chevron Oil ⁴ Northern Illinois University ⁴	DUKE UNC NCSU TUCC
¹ Plans to become a distributed network. ² Only two nodes are connected at the present time. ³ Two different nodes at same location. ⁴ Entry into the network under negotiation. ⁵ Centralized File Subnet. ⁶ Distributed Teletype Subnet. ⁷ Includes the CDC STAR which has not yet been delivered.					

FEATURES	NETWORKS			
	ARPA	COINS	CYBERNET	DCS
<p>Configuration (continued)</p> <p>Types¹ of Host Computers Available over the Network</p>	Burroughs B6500 Burroughs ILLIAC IV DEC PDP-10 (TENEX) DEC PDP-1 DEC PDP-10 DEC PDP-10 Data Machine DEC PDP-11 GE 645 (MULTICS) Honeywell DDP-516 IBM 360/75 IBM 360/91 IBM 1800 IBM 360/65 IBM 360/65 (ADEPT) IBM 360/67 TX-2 XDS-940 XDS SIGMA 7	GE 635 UNIVAC 494	CDC 6600 (SCOPE) CDC 3300 (MASTER) CDC 6400 (KRONOS)	Micro 800 Teletype Varian 620/i Varian 620/a
<p>Special Hardware or Software Available Through the Network</p>	NLS TODAS ILLIAC IV PDP-10 Data Machine Laser Memory Culler-Fried System CPS MATHLAB CONVERSE ORBIT Artificial Intelligence Projects LOGOS	---	Structural Analysis Packages Linear Programming Packages	---
<p>¹While only one of each type of computer configuration is listed many of the same configurations may exist in any one network. In some instances the listed hosts have not yet been installed but they have been included in the tabulation because of definite plans to connect them to the network.</p>				

NETWORKS					
DLS	MERIT	NETWORK/440	OCTOPUS	TSS	TUCC
IBM 360/50 (OS MVT) IBM 360/65 (OS MVT)	CDC 6500 (SCOPE) IBM 360/67 (MTS) IBM 360/67 (MTS)	CDC 6600 ¹ IBM 360/91 (OS MVT) IBM 360/67 (CP) IBM 360/65 ¹ (OS MVT) IBM 360/44 ¹	CDC 6600 CDC 7600 CDC STAR	IBM 360/67 (TSS)	IBM 360/75 ² IBM 360/50 IBM 360/40
---	MTS	CP	Photo Store STAR	CP Graphics	CPS APL
¹ In negotiation. ² The grid node in a centralized net — other nodes may process jobs on their own machine or on the 360/75, but not on any of the machines at other nodes.					

FEATURES	NETWORKS			
	ARPA	COINS	CYBERNET	DCS
Communications				
Interface Device	Modified DDP-516	GE Datanet IBM 2701 UNIVAC CLT	CDC-3300 PPU	Special Design
Communications Protocol	Message Switch	Message Switch	Message Switch	Message Switch
Transmission Medium	Leased Line	Leased Lines	Leased Line Satellite DDD FX WATS	Coaxial Cable
Data Rates (bps)	50,000	2400	100-300 2000 2400 4800 40,800	2,000,000
Transmission Mode	Analog	Analog	Analog	Digital
Synchronization	Synchronous	Synchronous	Asynchronous	Synchronous
Link Protocol	Full-Duplex	Full-Duplex	Full-Duplex Half-Duplex	Simplex
Data Compression Used	No	No	No	No
Message Format	Variable Length	Variable Length	Fixed Length	Fixed Length
Message Size¹	8095 bits	15,000 characters	1024 characters	240 bits
Segment Format	Variable Length 1000 bits max	Variable Length 150 characters max	N/A	N/A
Security Level	None	Top Secret	None	None
¹ For variable length messages, this is the maximum message size.				

NETWORKS					
DLS	MERIT	NETWORK/440	OCTOPUS	TSS	TUCC
IBM 2701	Modified PDP-11/20	IBM 2701	PPU	IBM 2701	IBM 2701
Point-to-Point	Message Switch	Message Switch	Message Switch ¹ Point-to-Point ²	Point-to-Point	Point-to-Point
Leased Line	Telpak	Leased Line	Coaxial Cable	DDD	Telpak
40,800	2000	40,800	1,500,000 ¹ 12,000,000 ²	2000 40,800	100-300 2400 40,800
Analog	Analog	Analog	Digital	Analog	Analog
Synchronous	Synchronous	Synchronous	Synchronous	Synchronous	Asynchronous
Half-Duplex	Full-Duplex	Half-Duplex	Full-Duplex ¹ Half-Duplex ²	Half-Duplex	Half-Duplex
Yes	No	No	No	Yes	No
Variable Length	Variable Length	Variable Length	Variable Length	Variable Length	Variable Length
32,760 ³ bytes	240 characters ⁴	8192 bits	1208 bits ¹ 3,780,000 bits ²	8192 bits	1000 bytes
N/A	N/A	N/A	N/A	N/A	N/A
Top Secret	None	None	AEC Restricted Data	None	None

¹The Teletype Subnet.
²The File Transport Subnet.
³Absolute maximum. Limited core memory imposes an operational limit of 7294 bytes.
⁴May be varied by changing software parameter.

FEATURES	NETWORKS			
	ARPA	COINS	CYBERNET	DCS
Network Usage				
Load Sharing	No ¹	No	Yes	No ³
Program Sharing	Yes	No	Yes	No ³
Data Sharing	Yes	Yes	Yes	No ³
Remote Service	Yes	No	Yes	No ³
Dynamic File Access	No ²	No	No	No ³
Experimentation	Yes	Yes	No	Yes
Measurement	Yes	Yes	No	Yes
Data Description Language	Yes	No	No	No
Control Language	No ²	No	Yes	Yes
<p>¹While load sharing is technically possible, the network was not designed for this purpose. ²Technically feasible, still in research stage. ³This network is experimental and has not yet developed plans for software to assist users.</p>				

NETWORKS					
DLS	MERIT	NETWORK/440	OCTOPUS	TSS	TUCC
Yes	No	Yes	No	Yes	No
Yes ¹	Yes	Yes	Yes	Yes	Yes
Yes	Yes	Yes	Yes	Yes	Yes
Yes	Yes	Yes	Yes	Yes	Yes
No	Yes	No	No	Yes	No
No	Yes	Yes	No	Yes	No
No	Yes	Yes	Yes	Yes	Yes
No ²	No ³	Yes	No ⁴	Yes	No ²
Yes	No	Yes	Yes	Yes	No

¹While program sharing is possible, it is an unlikely use because of the large data bases involved.
²Unnecessary because of like processors.
³MICIS has proposed a data description language, but it has not been adopted as the standard.
⁴Not required because of centralized data base.

SECTION IV EXCLUDED NETWORKS

Six additional networks were considered for inclusion in the survey but were rejected either because they did not satisfy the definition of a computer network given in the Introduction or because insufficient information was available:

ARS (Advanced Record System) is a General Services Administration network used by various departments and civilian agencies of the Federal government. ARS was not included in the survey since it is primarily concerned with transmitting data bases and responding to queries, rather than resource sharing through connected computer systems.

ASP (Attached Support Processor) is an IBM-developed system providing facilities for simultaneously controlling the operation of several support processors and main processors. The support processors collect work from input stations, monitor library requests, schedule work and distribute output, while the main processors execute jobs exclusively. Because the processors do not operate under independent operating systems, ASP does not conform to the definition of a computer network given in the Introduction.

CLETS (California Law Enforcement Telecommunications System) connects over 450 different law enforcement agencies at the state, county and local levels to files in Sacramento and Washington, D. C. CLETS is the nation's largest state-wide computerized law enforcement system. A query system, this network does not meet the criteria used in this survey.

NCIC (National Crime Information Center) is an information network including 49 states and servicing law enforcement agencies at all levels of the government. NCIC is a query system and, hence, does not conform to the survey definition of computer network.

NTDS (Navy Tactical Data System) is a secure network which performs the routing of tactical data from ship-to-ship and from ship-to-shore. The computer facilities used in NTDS perform message switching and local decision-making functions, but do not support intersite resource sharing; consequently, NTDS does not satisfy the definition of a computer network.

The RTCC (Real Time Computing Complex) at the Manned Spacecraft Center in Houston, Texas, has linked their computers to achieve the reliability necessary for manned mission support. During a mission, two computers are actively engaged in processing. The primary computer is termed the missions operation computer, and

the backup computer is designated the dynamic standby computer. As long as the missions operation computer functions satisfactorily, the output of the dynamic standby is unused. If a failure occurs in the primary system, the dynamic standby computer becomes the missions operations computer; and another computer is activated as the dynamic standby. The usage of linked computers does not comply with the definition of a computer network.

SECTION V SUMMARY

The surveyed networks were designed to meet very different needs, making any comparison extremely difficult; therefore, this paper examines the current philosophies on networking and provides extensive information on those networks included in the survey.

From the survey it is apparent that computer networking is in its infancy. Only CYBERNET, Octopus, and TUCC can be considered fully operational, and they are in a state of continuous development of added capabilities. The other networks must all be considered either in an experimental phase or under development leading to an initial operational capability. Since computer networking is within the state-of-the-art, most of these networks can reasonably be expected to become operational in the foreseeable future.

This survey reflects how features of computer networks are employed in actual or planned networks. The information herein can serve network designers in determining how some of these features can be used to meet their networking requirements.

GLOSSARY

Acknowledgment	A process whereby the transmitter of a message is notified of its receipt by the receiver
ADEPT	A time-sharing operating system developed by System Development Corporation for use on IBM 360 computer systems.
ALGOL	Algorithmic Language, a processing language oriented to the arithmetic specification of numerical procedures.
Analog Transmission	A method of data transmission requiring the original information signal to be converted into an ac form at the proper frequency for transmission.
ANMCC	The Alternate National Military Command Center, a protected backup facility for the National Military Command System.
APL	A Programming Language, a highly interactive programming system particularly suited to the manipulation of vectors and matrices.
ARPA	The Advanced Research Projects Agency of the Department of Defense.
ASCII	The American Standard Code for Information Interchange.
ASP	Attached Support Processors, an operating environment for multiple IBM 360 installations, which optimizes allocation of local resources among waiting jobs.
Asynchronous Transmission	A mode of data transmission in which the time spacing between bits, characters, or message segments is not regular.
AWS	Air Weather Service, a scheduled node on the ARPA network.
BASIC	Beginner's All-purpose Symbolic Instruction Code, a simple, mathematically-oriented interactive programming language.
BBN	Bolt, Beranek and Newman (Cambridge, Mass.), the developers of the ARPA IMP, and a node of the ARPA network.
BCC	Binary-Synchronous Communication Controller, the communications module of the Data Link Support software.
bps	Bits per second.
BTAM	Basic Telecommunications Access Method, a data communications protocol supported by OS/360.

BTL	Bell Telephone Laboratories, the research arm of the American Telephone and Telegraph Company.
CAM	Computer Access Method, a specially developed communications protocol used in the TSS Network.
CC	Communication Computer, a modified PDP-11/20 computer used as a communications interface in the MERIT Network.
CCA	Computer Corporation of America, contractors for the development of the PDP-10 Data Machine to be attached to the ARPA network.
CCAM	Continuous Communications Access Method, a portion of the Data Link Support control software.
CCOS	The Communications Computer Operating System a specially developed communications software system used on the MERIT Network.
CDC	Control Data Corporation, a manufacturer of large computer systems and the operators of the CYBERNET Network.
Centralized Network	A network configuration in which each user node connects to a central node which performs all message switching functions.
CLT	Communications Line Terminal, a computer/link interface used on UNIVAC and ARPA systems.
CMU	Carnegie-Mellon University, Pittsburgh, Pa., a node in the TSS and ARPA networks.
COBOL	Common Business Oriented Language, a processing language designed for data processing applications.
COINS	Community On-Line Intelligence System, a network developed for the intelligence community.
Communications Protocol	A specification of whether intercomputer communication is achieved by message-switching or point-to-point lines.
COMPASS	Comprehensive Assembler System, a name generically applied to CDC assemblers.
Composition	A specification of whether a network is composed of similar (homogeneous) or dissimilar (heterogeneous) computer systems.
Computer Network	An interconnected group of independent computer systems which communicate with one another and share resources such as programs, data, hardware, or software.
Control Language	A special language which allows the user to directly specify functions that he wishes the network to perform.

CONVERSE	A natural language data management system developed by System Development Corporation.
COPE	Communications Oriented Processing Equipment, a series of terminals produced by University Computing Corporation.
CP	Control Program, a virtual memory operating system developed for the IBM 360/67.
CPS	Conversational Programming System, an interactive monitor used on IBM 360 computer systems.
CPU	Central Processor Unit, the central arithmetic and logic unit of a computer system.
CSMP	Continuous Systems Modeling Program, an applications program which provides the time response of physical systems defined by simultaneous differential equations.
Culler-Fried	An interactive system developed at the University of California at Santa Barbara to perform quick-response mathematical computations.
Data Definition Language	A control language which allows a user to define the structure of a data base. This is particularly useful in transferring data between machines with dissimilar data formats and structures.
Data sharing	A mode of network operation in which programs are sent to a particular node in order to utilize data resident at that node.
DCA	Defense Communications Agency, operators of the Data Link Support System.
DDD	Direct Distance Dialing, a class of switched, voice-grade service provided by the American Telephone and Telegraph Company.
DEC	Digital Equipment Corporation, manufacturers of the PDP series of computer systems.
Decentralized network	A network configuration in which groups of local user nodes are interconnected in a centralized fashion. The grid nodes are then connected to form a distributed network.
Digital Transmission	A method of data transmission in which unmodulated digital data is transmitted directly.
Distributed Network	A network configuration in which all node pairs are connected either directly, or indirectly through intermediate nodes and shared links.
DLS	Data Link Support, a communication facility interconnecting two of the elements of the National Military Command System.

Dynamic File Access	A mode of network operation in which programs executing at one node can access data at a remote node as if the data were locally resident.
EAI	Electronic Associates, Incorporated, a manufacturer of analog and hybrid computer systems.
EASE	Elastic Analysis for Structural Engineering, a CYBERNET applications program which performs structural analyses.
ECAP	Electronic Circuit Analysis Program, an applications program which performs analyses of linear and non-linear electrical networks.
ETAC	Environmental Technical Applications Center, Air Weather Service, Washington, D. C., a scheduled node on the ARPA network.
Experimentation	A specification of whether network research experiments are being carried out on the network.
FORTTRAN	Formula Translation, a programming system designed to solve problems which can be expressed in algebraic notation.
Full-Duplex	A communications link which permits simultaneous use of the link in both directions.
FX	Foreign Exchange, a class of leased line service offered by the American Telephone and Telegraph Company in which a termination in one central office is assigned a number belonging to a remote central office.
GPL	General Precision Librascope, manufacturers of small computers and peripheral equipment.
GPSS	General Purpose Simulation System, a discrete system modeling program.
Half-Duplex	A communications link which may be operated in either direction, but in only one direction at a time.
HASP	Houston Automatic Spooling Priority, a unit record support program operated in conjunction with IBM 360 OS.
Heterogeneous	A network characteristic denoting the use of dissimilar computer systems.
Homogeneous	A network characteristic denoting the use of only similar computer systems.
Host	A computer system which provides the user with an appropriate interface to the network.

ILLIAC Illinois Automatic Computer, a specially-designed, multiple-processor computer system under development by the Burroughs Corporation.

IMP Interface Message Processor, a modified Honeywell PDP-516 used as a communications computer on the ARPA network.

JOSS Johnniac On-Line Support System, an interactive algebraic interpreter originally developed by the RAND Corporation.

JOVIAL Jules' Own Version of the International Algebraic Language, an algebraically oriented computer language.

K An abbreviation used to represent the number 1024, e.g., 32K is 32,768.

KRONOS A CDC 6400 time-sharing operating system.

LC² An interactive ALGOL-like language in use at Carnegie-Mellon University.

LCS Large Capacity Storage, a large, relatively slow core memory extension used on IBM 360 computer systems.

LEAP A graphics language used at Lincoln Laboratories.

LIL Local Interaction Language, an interpretive system for interactive graphics in operation at Lincoln Labs.

Link A communications channel which interconnects a pair of nodes.

Link Protocol A specification of whether the link is operated in a simplex, half-duplex, or full-duplex fashion.

LISP List Processing, a programming system designed to facilitate the manipulation of linked lists.

Load Sharing A mode of network operation in which a given workload is distributed among the computer systems of the network in order to achieve equal use of resources.

LOGOS A design tool capable of producing systems that are certifiable as being secure; currently under development at Case Western Reserve University.

LRL Lawrence Radiation Laboratory, Livermore, California, a part of the University of California operated as a research facility for the Atomic Energy Commission; the location of the Octopus

LRLTRAN An extended version of the FORTRAN programming system written and used by the Lawrence Radiation Laboratory.

MARC	Multiple Access Remote Computer, a designation given to terminals on the CYBERNET Network.
MASTER	A CDC 3300 multiprogramming operating system.
MATHLAB	A mathematically-oriented computation system in operation at the Massachusetts Institute of Technology.
Measurement	A specification of whether statistics reflecting the performance of the network are being collected and analyzed.
MERIT	Michigan Educational Research Information Triad, Incorporated, a non-profit corporation responsible for the development and operation of the MERIT network.
Message	A logical unit of communication between two hosts.
Message-Switching	A process of accepting a message for the purpose of relaying it toward its destination.
MICIS	The Michigan Interuniversity Committee on Information Systems, the predecessor of MERIT and the MERIT network.
MIT	Massachusetts Institute of Technology, a node in the ARPA network.
Modem	Modulator/Demodulator, a device used to convert digital signals to an appropriate frequency for analog transmission, and vice versa.
MPS	Mathematical Programming System, a library of scientific subroutines available under OS 360.
MSU	Michigan State University, East Lansing, Michigan, a node in the MERIT network.
MTS	Michigan Terminal System, a specially developed time-sharing operating system used on the IBM 360/67 at the University of Michigan.
MULTICS	An experimental time-sharing system in operation at MIT.
Multiplexing	A method of "packing" several data streams into a single link in order to achieve a higher effective transmission rate. Two methods are commonly used: time multiplexing, in which the bits or characters of the data streams are interleaved to form a single stream; and frequency multiplexing in which each data stream is assigned a particular frequency slot in the overall link bandwidth, and all the streams are transmitted in parallel.
MVT	Multiprogramming with a Variable Number of Tasks, an IBM 360 operating system.

NASA The National Aeronautics and Space Administration. Two of NASA's research facilities, the Ames Center in Mountainview, California, and the Lewis Center in Cleveland, Ohio, are nodes in the TSS network.

NASTRAN A special programming system developed by NASA for performing structural analyses.

NCP Network Control Program, an ARPA host program which interfaces the host with the network.

NCAR National Center for Atmospheric Research, Boulder, Colorado, a scheduled node in the ARPA network.

NCSU North Carolina State University, Raleigh, North Carolina, a node in the TUCC network.

NETFLOW A CYBERNET applications program which solves network flow problems.

NIC Network Information Center, a network node which acts as a central repository for network documentation.

NLS A system developed by SRI which is similar to TODAS, but employs graphic display terminals.

NMCS National Military Command System, an aggregation of agencies and systems which provide communications and processing support for the Joint Chiefs of Staff.

NMCSSC National Military Command System Support Center, the ADP support arm of the NMCS.

NMCSTS National Military Command System Technical Support, the directorate of the Defense Communications Agency responsible for the operation of the NMCS.

Node A data transmission terminal point.

NWG Network Working Group, the technical advisory board for the ARPA network.

OCAMA Oklahoma City Air Materiel Area, a projected node on the ARPA network.

Octopus A computer network developed and used by the Lawrence Radiation Laboratory.

OPHELIE II A linear programming system available on the CYBERNET network.

OPTIMA A linear programming system available on the CYBERNET network.

ORBIT	On-Line Retrieval of Bibliographic Information, a system developed by SDC.
Organization	A specification of the basic topology of the network.
OS	Operating System, the primary multiprogramming monitor used on IBM 360 computer systems.
Packet	An ARPA term used to denote a message segment.
PCM	Pulse Code Modulation, a form of digital transmission.
PDQ/LP	A linear programming system available on the CYBERNET network.
PL/I	Programming Language I, a higher-level programming language capable of performing a wide range of algebraic, data processing, and system-level functions.
Point-to-Point	A network topology in which conversant nodes communicate directly without intervening message switching.
PPU	Peripheral Processing Unit, a small, programmable computer which typically controls data flow on CDC 6000 and 7000 series computer systems.
Program Sharing	A mode of network operation in which data is sent to a particular node to be processed by programs resident at that node.
RADC	Rome Air Development Center, Rome, New York, a scheduled node on the ARPA network.
Remote Service	A mode of network operation in which programs and data resident at a given site are manipulated by a remote user.
RJET	Remote Job Entry Terminal, a planned subnet of the Octopus network.
SAAC	Seismic Array Analysis Center, Alexandria, Virginia, a projected node on the ARPA network.
SAAMA	Sacramento Air Material Area, a projected node on the ARPA network.
SCOPE	Supervisory Control of Program Execution, a CDC 6600 batch operating system.
SDC	System Development Corporation, Santa Monica, California, a node in the ARPA network.
Segment	The portion of a message which serves as the basic unit record of information interchange between two nodes.

Simplex A communications link which can be operated in one direction only.

SIMSCRIPT A discrete system simulation program.

SIMULA A simulation system available on the CYBERNET network.

SNOBOL A string manipulation system developed by Bell Telephone Laboratories.

SRI Stanford Research Institute, Menlo Park, California, the Network Information Center for the ARPA network.

SRS Start-Restart Synchronizer, a portion of the Data Link Support operating software.

STAR String Array, a powerful, high-speed computer system recently developed by the Control Data Corporation.

STARDYNE A dynamic structural analysis system available on the CYBERNET network.

Store-and-Forward A form of message switching in which the message is stored at the intermediate node prior to being forwarded. This allows the previous node in the chain to release its copy of the record.

Synchronous transmission A mode of data transmission in which the time spacing between bits, characters, and message segments is regular.

SYSTEM 2000 A Data Management System available on the CYBERNET network.

TSS Time Sharing System—a time-sharing operating system for the IBM 360/67 computer.

TX-2 A specially developed computer in operation on the ARPA network at Lincoln Labs.

UCC University Computing Corporation, manufacturers of COPE-Series communications terminals.

UCI The University of California at Irvine, developers of the Distributed Computer System.

UCLA The University of California at Los Angeles, the Network Measurement Center for the ARPA network.

UCSB The University of California at Santa Barbara, a node in the ARPA network.

UM University of Michigan, Ann Arbor, a node in the MERIT network.

UNC University of North Carolina, Chapel Hill, a node in the TUCC network.

WATFIV	An in-core FORTRAN compiler, particularly well suited to processing small, student jobs.
WATS	Wide Area Telephone Service, a voice-grade leased line service offered by the American Telephone and Telegraph Company.
WHINOT	An in-core PL/I compiler developed by TUCC, used primarily for short student jobs.
WSU	Wayne State University, Detroit, Michigan, a node in the MERIT network.
XDS	Xerox Data Systems, manufacturers of the Sigma computer series.

APPENDIX

The following is a list of individuals who can provide a central point of contact for each of the surveyed networks.

ARPA	Mr. Steve Crocker, Chairman Network Working Group ARPA/IPT 1400 Wilson Boulevard Arlington, Virginia 22209 202-694-5922
COINS	Classified
CYBERNET	Mr. Gabe Bustamante Control Data Corporation 4550 West 77th Street Minneapolis, Minnesota 55435 612-920-8600 X5956
DCS	Dr. David Farber Information and Computer Science Department University of California Irvine, California 92664 714-833-6891
DLS	Lt. Richard P. Quintana, B-111 NMCSSC The Pentagon Washington, D. C. 20330 202-695-4789
MERIT	Dr. Bertram Herzog, Director MERIT Computer Network 611 Church Street Ann Arbor, Michigan 48104 313-764-9423
NETWORK/440	Mr. Doug McKay Watson Research Center Yorktown Heights, New York 10598 914-945-3000 X1159

OCTOPUS

Mr. Sam Mendicino, L-61
Box 808
Livermore, California 94550
415-447-1100 X8582

TSS

Mr. Al Weis
Watson Research Center
Yorktown Heights, New York 10598
914-945-3000 X1593

TUCC

Dr. Leland Williams, President and Director
Triangle Universities Computation Center
P. O. Box 12175
Research Triangle Park, North Carolina 27709
919-549-8291

BIBLIOGRAPHY

- ARPA Network, *Current Guide to Network Facilities*, NIC 5148, Network Information Center, Stanford Research Institute, Menlo Park, California 94025.
- Aupperle, Eric M., "MERIT Computer Network Hardware," Notes from presentation at the computer network seminar at the Courant Institute of Mathematical Sciences, 30 November 1970.
- Bolt, Beranek and Newman, "The Interface Message Processors for the ARPA Computer Network," Report No. 2103, January 1971.
- Braden, R. T., "Request for Comments: 90," ARPA Network Working Group, NIC 5707, 25 January 1971.
- Brooks, Frederick P.; Ferrell, James K.; and Gallie, Thomas M., "Organizational, Financial and Political Aspects of a Three-University Computing Center," *Proceedings*, IFIP Congress 68, Edinburgh, August 1968.
- Carr, Stephen C.; Crocker, Stephen D.; and Cerf, Vinton G., "HOST-HOST Communication Protocol in the ARPA Network," *Proceedings, Spring Joint Computer Conference* 1970.
- Cocanower, A. B.; Fischer, W.; Gerstenberger, W. S.; and Read, B. S., "The Communications Computer Operating System—Initial Design," MERIT Computer Network, Manual 1070-TN-3, October 1970.
- Control Data Corporation, "CYBERNET Service."
- Crocker, Stephen D., "HOST-HOST Protocol," Document No. 1, ARPA Network Working Group, 3 August 1970.
- Davis, M. S., "Economics — Point of View of Designer and Operator," *Proceedings*, the University of Texas and the MITRE Corporation Interdisciplinary Conference on Multiple Access Computer Networks, Austin, Texas, 20-22 April 1970.
- Donaldson, Marylyn; Robinovitz, Stewart; and Wolfe, Barbara, "Proposed MICIS Standard for Data Description," Michigan Interuniversity Committee on Information Systems, 4 December 1970.
- DuBois, Pierre, et al., Lawrence Radiation Laboratory, "Livermore Time-Sharing System Part I: Octopus, Chapter 4: Files," 14 October 1970.
- Fletcher, John G., Lawrence Radiation Laboratory, "Livermore Time-Sharing System Part I: Octopus, Chapter 0: Introduction to the OCTOPUS Network," December 1970.
- Freeman, David N.; and Pearson, Robert R., "Efficiency versus Responsiveness in a Multiple-Services Computer Facility," *Proceedings — 1968 ACM National Conference*.

- Freeman, David N.; and Ragland, Joe R., "The Response-Efficiency Trade-Off in a Multiple-University System," *Datamation*, March 1970.
- Heart, F. E.; Kahn, R. E.; Ornstein, S. M.; Crowther, W. R.; and Walden, D. C., "The Interface Message Processor for the ARPA Computer Network," *Proceedings, Spring Joint Computer Conference 1970*.
- Herzog, Bertram, "MERIT Proposal Summary," MERIT Computer Network, February 1970.
- IBM, "Data Link Support (DLS) Final Report (Draft)," 25 September 1970.
- IBM, "Data Link Support (DLS) Program Description Volume I," 25 September 1970.
- Karp, Peggy M., The MITRE Corporation, Event Report, "February 1971 ARPA Network Working Group Meeting," 23 February 1971.
- Lawrence Radiation Laboratory, "A Tour of the Computer Facility," January 1971.
- Lawrence Radiation Laboratory, "Mathematical Programmers"
- Luther, W. J., "Introduction to CYBERNET," Presentation made at the Symposium on Computer Networks at the New York University Courant Institute of Mathematical Sciences, 30 November -- 1 December 1970.
- Management Consultants Bulletin, "State and Local Government," October 1970.
- MERIT Computer Network, Slides of the MERIT Network.
- Pehrson, David L., Lawrence Radiation Laboratory, "An Engineering View of the LRL Octopus Computer Network," 17 November 1970.
- Peterson, J. J., The MITRE Corporation, Event Report, "Visit to the Watson Research Center," 18 December 1970.
- Peterson, J. J., The MITRE Corporation, Event Report, "Visit to the University of California at Irvine," 18 January 1971.
- Peterson, J. J., The MITRE Corporation, Event Report, "Visit to the Stanford Research Institute," 19 January 1971.
- Roberts, Lawrence G.; and Wessler, Barry D., "Computer Network Development to Achieve Resource Sharing," *Proceedings, Spring Joint Computer Conference, 1970*.
- Rutledge, Ronald M.; Vareha, Albin L.; Varian, Lee C.; Weis, Alan H.; Seroussi, Salomon F.; Meyer, James W.; Jaffee, Joan F.; and Angell, Mary Anne K., "An Interactive Network of Time-Sharing Computers," *Proceedings of 24th National Conference Association for Computing Machinery, 1969*.
- Triangle Universities Computation Center, "TUCC Telecomputing Memorandum," TUCC Documentation Index -- 18, 10 February 1971.
- Veit, S. A., The MITRE Corporation, Event Report, "Visit to Lawrence Radiation Laboratory," 18 January 1971.

Veit, S. A., The MITRE Corporation, Event Report, "Visit to NASA Ames Research Center," 18 January 1971.

Weis, Alan, IBM Watson Research Center, Slides on the TSS Network.

Whitfield, Williard D., Veterans Administration, Text from Presentation at the National Communications System Computer/Communications Symposium, 27-29 October 1970.

Wood, D. C., The MITRE Corporation, Event Report, "Interim SHARE Network Project Meeting," 9 December 1970.

STAT

Approved For Release 2004/09/23 : CIA-RDP79M00096A000500010017-0

Approved For Release 2004/09/23 : CIA-RDP79M00096A000500010017-0