introduction to minicomputer networks

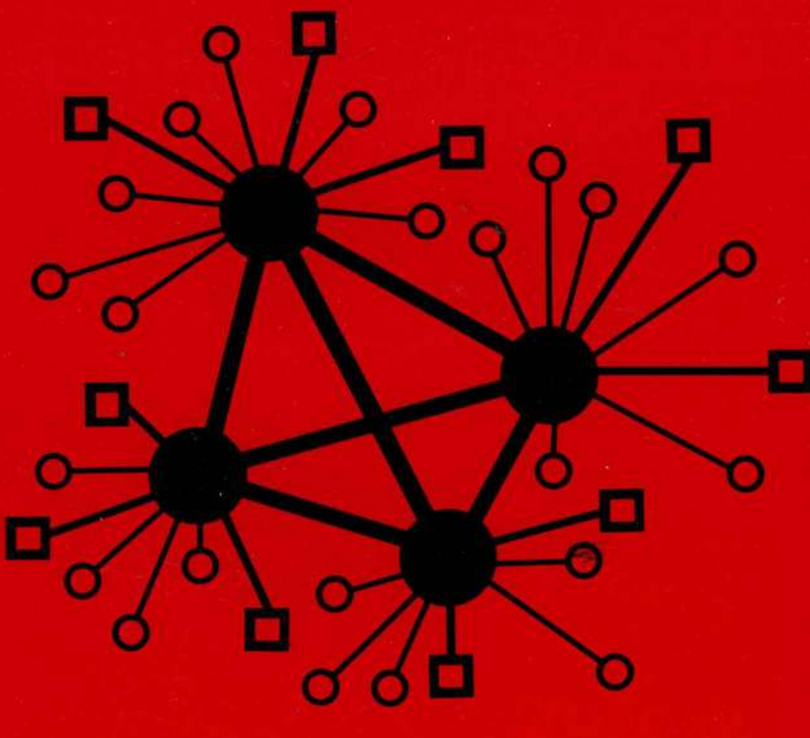

digital

introduction to minicomputer networks

digital

digital and data communication

DIGITAL has been involved with data communications almost as long as it has been involved with computers. In fact, some of our very first computers (the PDP-1 series, for example) were used in communications applications. Since then, we've been busy providing flexible and economical solutions to data communication applications. Because we were able to draw on that early experience, data communications features were an integral part of the PDP-11 minicomputer design. So it wasn't surprising to see the PDP-11 become the most widely used minicomputer in data communications.

Shortly after the PDP-11's introduction, a separate product line group—called DECcomm—was established to specialize in data communications. Our job was, and still is, to develop a large repertoire of hardware and software modules specifically to exploit the communications capabilities of the PDP-11 family of processors.

This Handbook was produced by the DECcomm Product Line of Digital Equipment Corporation. Its intent is to assist in technical communications. This task is becoming increasingly difficult as computer networking efforts have expanded from pure research within a closely knit community towards the development of commercial networks encompassing a diversity of applications.

Foreword

The use of computers is expanding at an ever increasing rate and with it, efficient data communications is becoming a necessity. To satisfy this need for better data communications, minicomputers are being used more frequently because of three important considerations: they are programmable, they have memory, and they are capable of supporting a wide variety of peripheral devices. These features allow the minicomputer to perform many of the functions that traditionally have been performed by hardwired controllers or by large EDP computers.

This Handbook focuses on those data communications systems referred to as "networks" or "distributed computer systems." These systems consist of multiple computers capable of communicating with each other in a variety of operating environments. There is a rapidly growing trend in the establishment of such networks. This trend is a logical extension of the traditional computer timesharing networks in which many remote terminals are linked to one large central computer.

At Digital Equipment Corporation, for example, a growing percentage of orders for our PDP-11 family of processors involves the control of various functions within data communications networks of all sizes and shapes. These network control functions are discussed in Chapter 3. They include front-end communications processing, data concentration, message switching, and the control of multiple remote terminals. In many network configurations, PDP-11's are functioning as main-site, or "host," computers, remote computing systems, and Remote Job Entry (RJE) or batch processors.

In recognition of this growing trend in multi-computer network design, Digital Equipment Corporation has placed a great deal of emphasis in the design and development of computer hardware and software products tailored specifically for use in such networks. In addition to the compatible family of PDP-11 processors, these products include synchronous and asynchronous line interfaces, message protocols, and a variety of communications software packages. These products and their implications to data communications users are summarized in Appendix A.

With continuous progress in the development in communications and computer technology, applications for computer-based data communications systems have become quite diversified. Today's applications are not confined to the business world; they span all user groupings—financial, industrial, scientific, and government. Some indication of this variety is presented in Chapter 6.

This Handbook assumes a familiarity with the basic concepts of data communications covered in INTRODUCTION TO DATA COMMUNICATIONS, published by Digital Equipment Corporation.

contents

illustrations

chapter 1
introduction

The increasing integration of computers and communications into a single system has led to a new and rapidly growing industry —the computer-based data communications industry. Although less than a decade old, the technological accomplishments within the industry have been significant. At universities, in industrial complexes, in financial institutions—wherever many users require the services of a computer—there is a rapidly growing demand for data communications facilities to link the central computer with remote users. This rapid growth trend is not unique to the United States. Considerable technical achievement along with a marked increase in the availability of communication services has been experienced on a worldwide basis.

The advances in technology allow communications to take place over long distances with increasing ease. Computers are "talking" to computers, people are "talking" to computers, and computers are "talking" to people. The telephone has become a necessity and the remote computer terminal is becoming a common administrative tool in many places of business. Government agencies, business concerns, and private individuals have all become heavily dependent upon the ability to communicate rapidly. Many businessmen have come to expect faster access to complete and timely data bases. The salesman wants to improve his competitive edge with audio response networks and computerized order entry systems that can be accessed from any major city. The scientist or engineer wants to interact with sophisticated computers from remote work sites. And everyone wants their travel and hotel reservations to be confirmed by computer quickly and reliably.

Brief History

Development of modern digital data communications can be traced directly to the advent of person-to-person telegraph systems.

Such systems have been in existence since Samuel Morse invented the telegraph in 1844. Torn-tape telegraph systems that transmit data at the rate of around 10 characters per second, have found wide use through the years. Many a railroad way station boasts continuing good service from a 30-year old teletypewriter. Even today considerable data communications traffic is generated between teletypewriters without any computer involvement.

The Emergence of Timesharing

The large-scale use and development of data communications began only in the mid-1950's when it was recognized that the teletypewriter circuits and machines could also be used for transmitting data to and from the computer. This uncovered many new application areas, permitting people to interact with the computer, thus utilizing its resources for information processing. Foremost among these new application areas is computer timesharing—a technique that enables many people to use one computer simultaneously on completely different problems whether the computer is within their sight or hundreds of miles away. A typical timesharing network is illustrated in Figure 1.

With the emergence of timesharing, the computer became an accessible, understandable tool that could be used in the solution of management's most common daily business problems. An engineer or scientist can use the central computer to solve complex research or mathematical problems quickly. Finance, manufacturing, and marketing managers can obtain up-to-date status reports on key control measurements of their business. Educational institutions can use the computer as a training tool. Government agencies can apply timesharing service to a wide variety of business and scientific problems, ranging from inventory control to aerospace computations.

Timesharing encouraged people to use the computer as a "thinking aid"—for casual use in real-life scientific and engineering problems—for the solution of fragments of problems encountered in research—for solving a host of other problems associated with product design, manufacturing, marketing, and finance.

COMPUTER NETWORKING

Since the advent of timesharing, the large central computer has traditionally been the heart of most data communications systems.

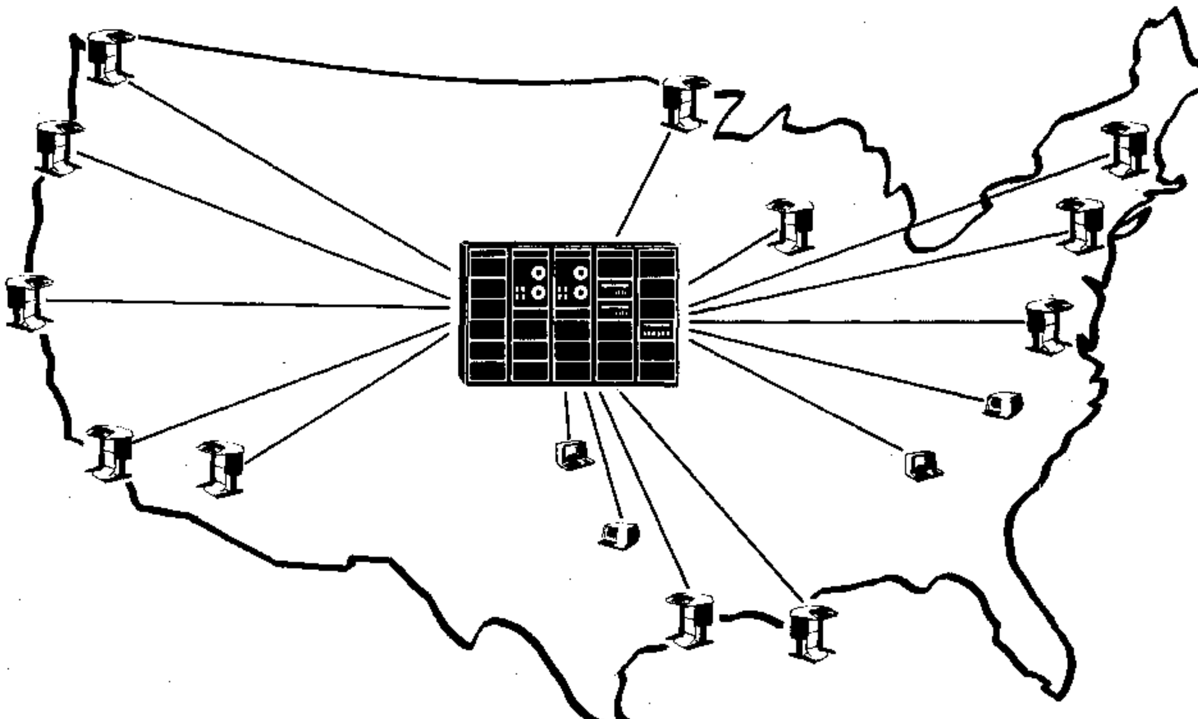

Figure 1   Computer Timesharing

Because of increasing demands for service, however, many users have been forced to upgrade to larger central computers only to find that increasing usage necessitated another upgrade in a few years. This type of growth could not go on indefinitely, and it soon became apparent that the only practical solutions were either to off load processing from the large central computer to multiple minicomputers or to replace the large central computer with a network of minicomputers.

A computer network can be defined as an interconnected group of computers, some of which are designated as "processing systems" and others as "communications control systems." The processing systems may be broken down into two types:

1. Main-site, or host, computers
2. Remote computing systems

The host processors in the network perform major computation, control data bases, and generally supervise operation of the network. They can share such resources as programs, data bases, and memory space.

Remote computing systems are local computing facilities with access to the host processors in the network. They perform processing that would otherwise have to be performed by the host processor. Thus, they relieve not only the communications load, but also that portion of the processing that can be carried on remotely, often closer to the point of entry.

In a true network situation, the host processors and remote computing systems can operate in a local mode under their own individual operating systems, can participate in network activity under the direction of a higher-level network supervisory program, or they can do both.

The communications control computers are devoted primarily to network control functions. These functions include line control, error checking, message formatting, message switching, and data concentration.

In addition to the processing and control computers, a typical network might consist of a wide variety of remote terminals to accommodate users at varying locations in an organization. Today's terminal equipment offers a wide variety of choices, from simple TOUCH-TONE devices with a 10-key numerical keyboard, to CRT display terminals, or large remote batch entry systems.

Thus, a large computer network might consist of several host processors, several remote computing systems, several communications control computers, a variety of remote terminals, and the transmission paths or channels that link all the components together. Figure 2 illustrates a sampling of some typical network components.

Advantages of Minicomputer Networks

A minicomputer network offers a number of significant advantages. By segmenting the application problem, a network of smaller computers can handle a distributed processing load more economically, and can provide more processing power than could possibly be built into one large central computer. A network of minicomputers has the ability to grow fairly easily as increasing demands are made upon it; minor variations for specific parts of the network can be handled conveniently, and software development is significantly facilitated.

The topology of the network is geared to individual jobs. Minicomputers are placed where the work is. With this approach, each processor performs a clearly defined job, communicates with its neighbor processors, as needed, and also responds to commands from a higher ranking, or host, processor. In this way, the software breaks down to a manageable level, programming time and costs are reduced, and data are available where needed.

The system can be made very tolerant to failures since the failure of one part of the network can be made to have limited, or even zero, effect on the total network operation. It is also fundamentally simpler to develop such a system since it can be built, tested, and put into service one leg at a time. It is not necessary to have a large central software complex running correctly before any service can be provided.

Resource Sharing

Among the primary reasons for the growing interest in minicomputer networks are the economies and conveniences that can be achieved through resource sharing. By linking several host and remote computers together through communications nodes—themselves little computers—researchers have come up with a community of computers that share vast computational power, massive data base files, and specialized software.

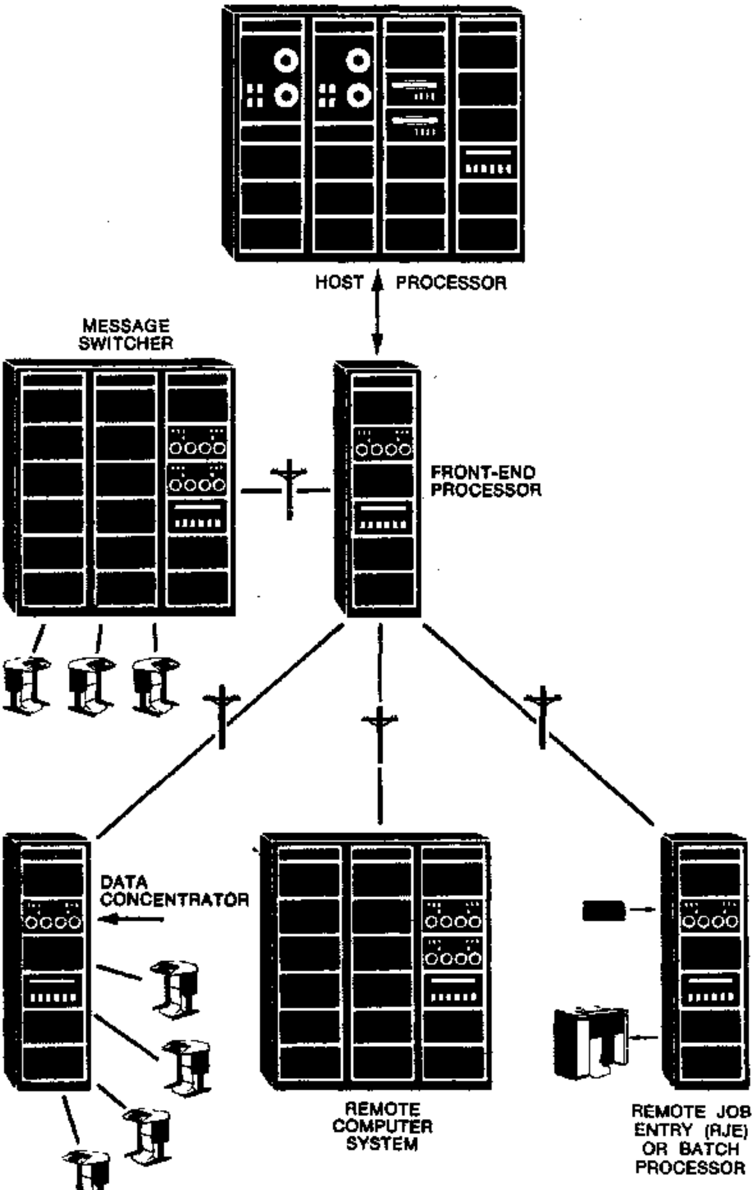

Figure 2 Typical Network Components

Resources available at each location in the network are available for use at other locations in the network. The following aspects are involved:

- DEVICE SHARING: The ability to connect to and use the resources of a remotely located computing device as if it were local.
- FILE SHARING: The ability to read from, write to, or update files on a remotely located computing system as if they were local.
- PROGRAM SHARING: The ability to send a loadable program to a remotely located computing system to be loaded and executed by that system.
- PROGRAM DATA SHARING: The ability to open a data path between programs on an interactive basis, so that large tasks may be divided into smaller units for execution at different computing sites in the network.

chapter 2

network topology and switching techniques

GENERAL

Topology refers to the geometric arrangement of links and nodes of a network. A link is the communications path between two nodes. The term link is used synonymously with circuit and channel. The term "node" can be defined as an end point of any branch of a network, or a junction common to two or more branches of a network. A variety of computer hardware and software may be installed at network nodes depending on their major function. Hence, a node might be a remote computing system, a host computer, or a computer devoted exclusively to network control functions, such as data concentration or message switching.

The network topology is related to network design, operations, reliability, and operating cost. A fully connected distributed network has more links for the same number of nodes than either a partially connected network or a simple star network. Differing forms of network control are reflected in centralized or decentralized networks, in hierarchical networks, as well as in multipoint or point-to-point connections of terminals and computers. These various network types are discussed in this chapter.

In designing a network, many factors must be evaluated in choosing the most suitable topology. However, one major factor can exert a pronounced influence on this choice: the type of participation by the nodes. Any node can be a provider of resources exclusively, a user of resources exclusively, or some combination of resource provider and resource user. Centralized structures are advantageous when a relatively large number of the sites need only reach one or a few sites which participate merely as resource providers. On the other hand, fully distributed structures are most appropriate when all sites participate both as resource users and resource providers.

A sampling of some typical network types is presented in the remainder of this chapter. This is followed by a brief discussion of the techniques used in switching and routing messages through a network.

BASIC NETWORK TYPES

Point-to-Point

The simplest possible network structure is that of a point-to-point connection shown in Figure 3, in which the host processor is connected to one communications input/output device per line. The communications I/O device may be a terminal or another processor. Whenever the host processor or other communications I/O device has data to transmit, the line is available. The connection may be fixed as a private line or switched as on the dial-up network. If the dial-up method is used, the receiving device has to be available to receive in order for transmission to take place. If a busy signal exists, the sender has to wait before sending his traffic across the system.

Multipoint

Next in complexity is the multipoint (or multidrop) connection, a party-line structure in which several users share the same line. In multipoint operation, one station in the network (normally the host processor) is always designated as the control station as shown in Figure 4. The remaining stations are designated as tributary stations. The control station controls network traffic by means of polling; that is, it invites (or polls) the tributary stations (which may be terminals or computers) to send messages as distinguished from free-for-all basis (referred to as contention). Multipoint networks are usually established over leased (non-switched) lines.

Multipoint networks may be centralized or distributed. In a centralized multipoint network, messages are transmitted between the control station and tributary stations but not between tributary stations. In a distributed multipoint network, messages are transmitted between the tributary stations, or between the control station and tributary stations.

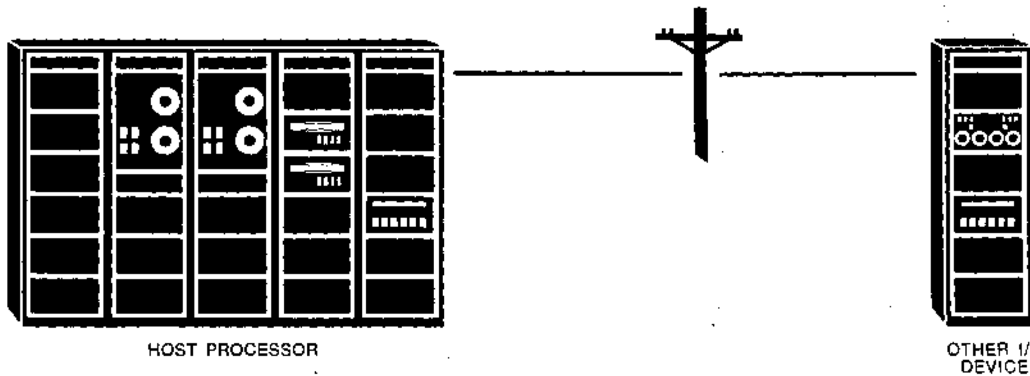

Figure 3 Point-to-Point Network

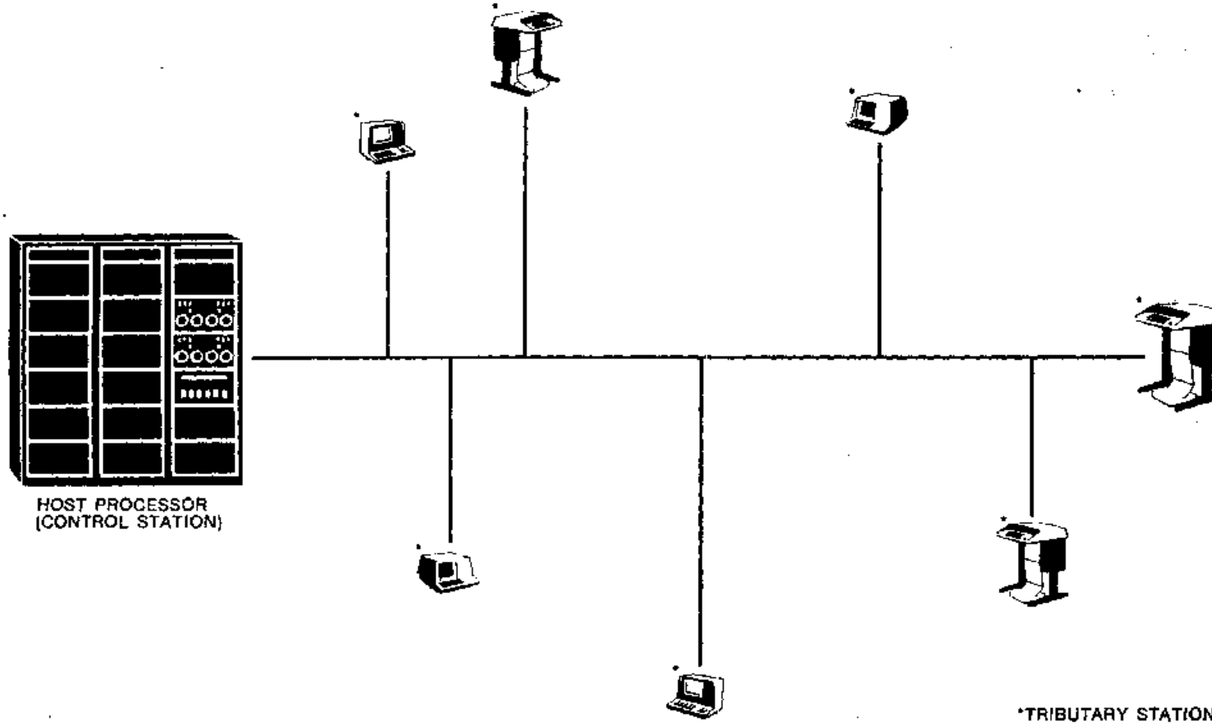

Figure 4 Multipoint Network

MIXED STRUCTURES

Centralized or Star

A centralized system with a star configuration is illustrated in Figure 5. In this type of system all users communicate with a central point that has supervisory control over the system. Users communicate with each other only by permission of this central processor. Data movement is outward from or inward toward the host. If communciation becomes necessary between the remote processors or terminals, the host acts as a central message switcher to pass data between them. This configuration makes the network simple to control.

Hierarchical or Tree Structure

In industrial environments, a hierarchical or tree structure (see Figure 6) is often used to supervise and control a variety of real-time, process-control applications. In such systems, a hierarchy of computers is used to control processes, synchronize them, and report their status. Small sensor-based systems provide real-time process control handling while recording the occurrences of events at each process and reporting them to a supervisory level. The computers at the supervisory level coordinate the sensor systems and report status, parts count, etc. to a host computer for corporate planning, inventory control, etc.

Loop or Ring Structure

Many organizations design their computer networks in the form of loop or ring structures. With this arrangement, many of the remote stations (terminals or computers) connected to the ring do not communicate with the main-site or host processor individually. Instead, the data to be transmitted is looped around the stations as illustrated in Figure 7.

The loop or ring structure is economical when several remote stations and host processors are located near each other. When remote stations are geographically dispersed over long distances, line costs would be very expensive if a loop structure were used. In such cases, it would be more economical to use a distributed-type structure as illustrated in Figures 8 and 9.

Distributed or Multistar

Figure 8 shows a distributed system with a multiple star configuration in which there are several supervisory or exchange

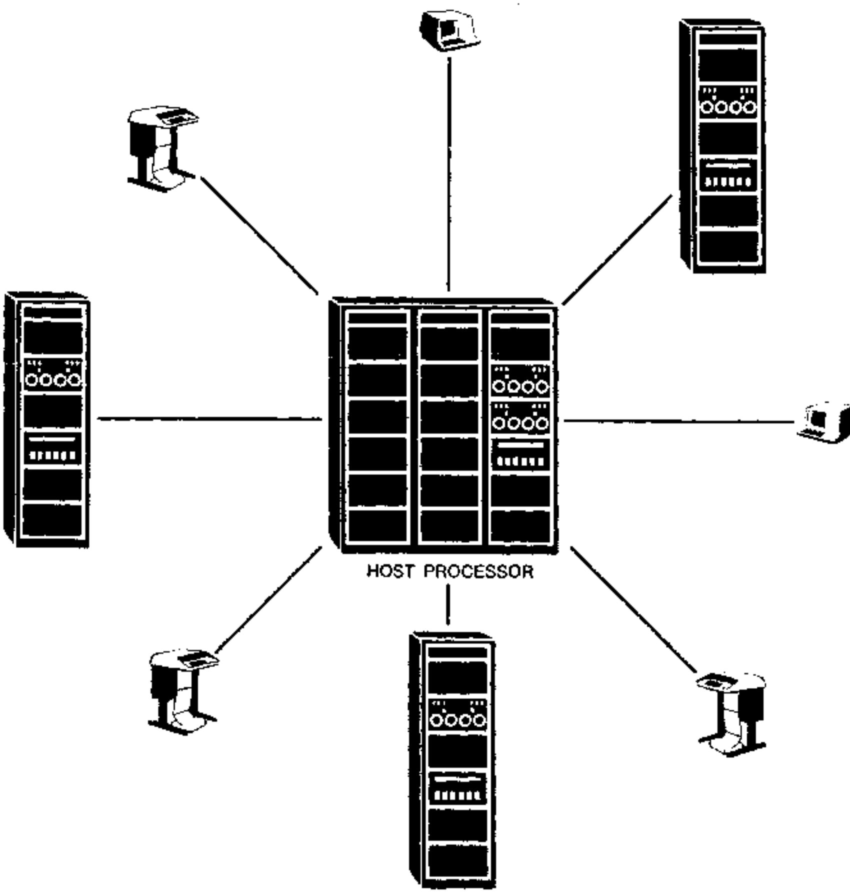

Figure 5 Centralized or Star Network

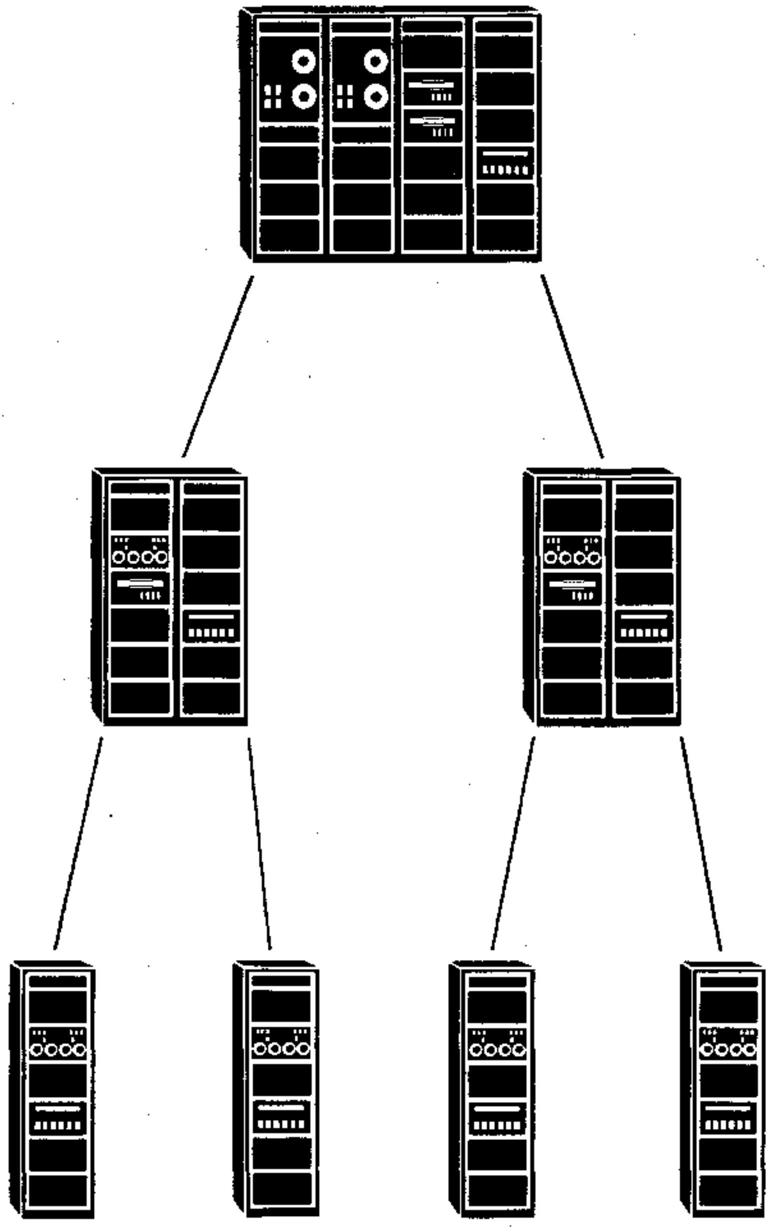

Figure 6 Hierarchical or Tree Structure

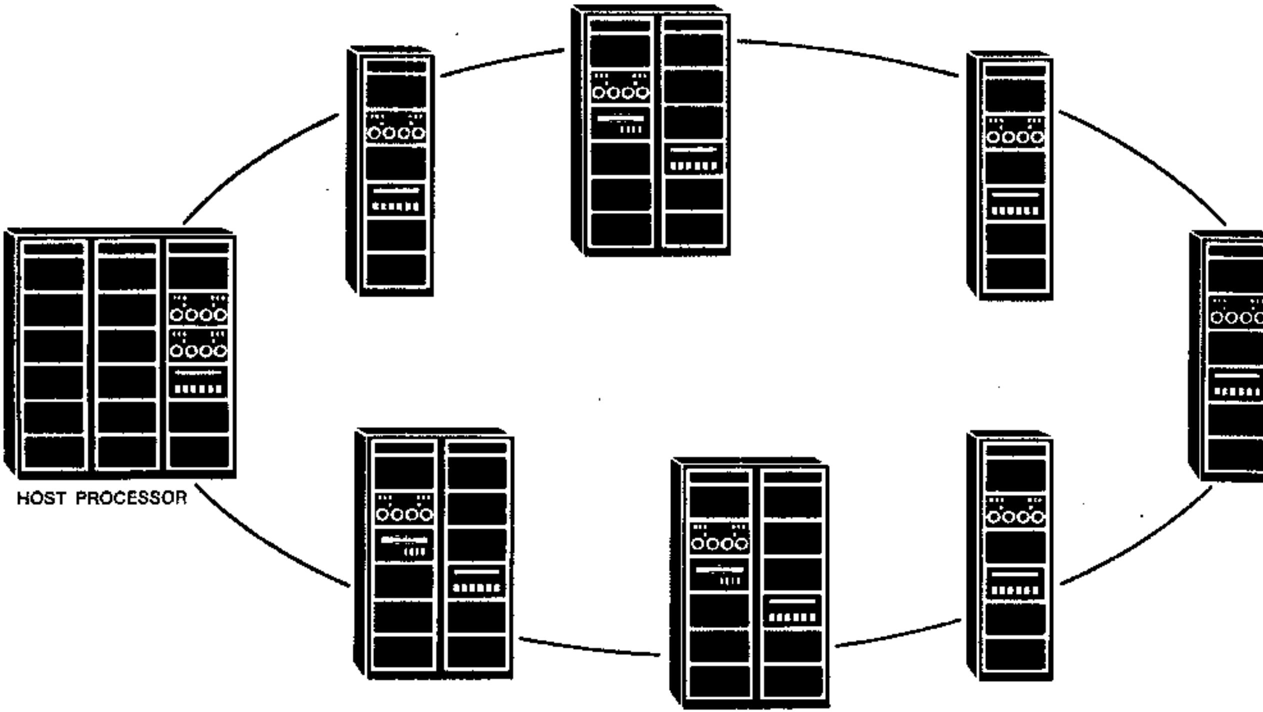

Figure 7 Loop or Ring Structure

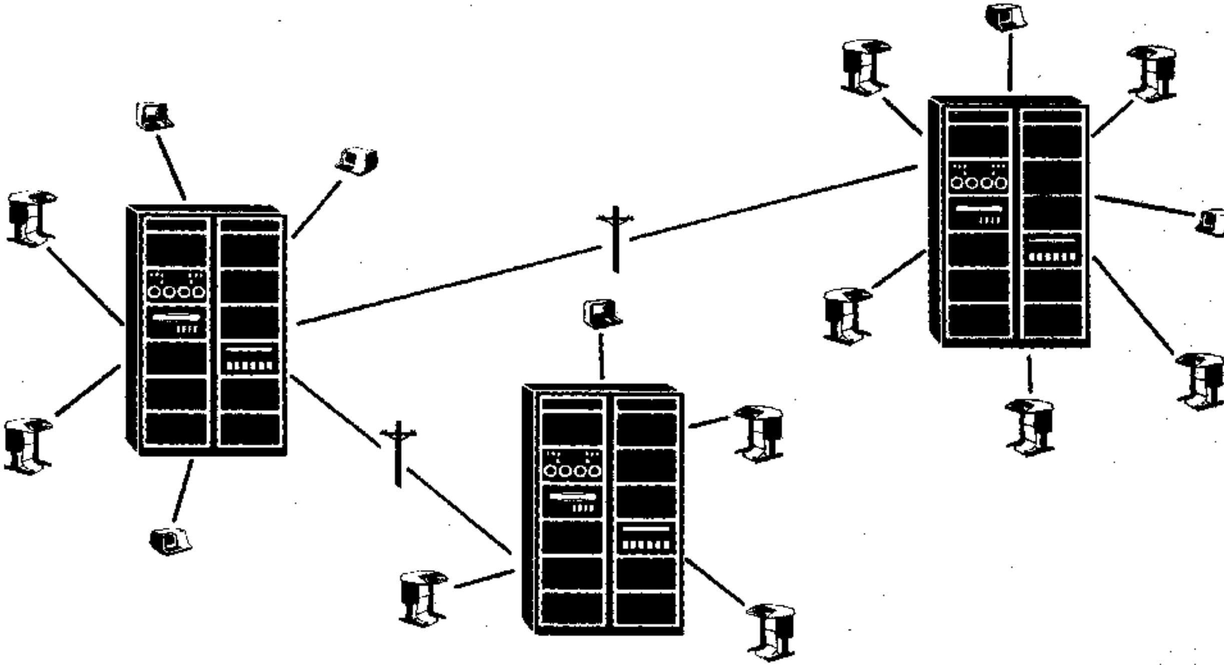

Figure 8 Distributed or Multistar Network

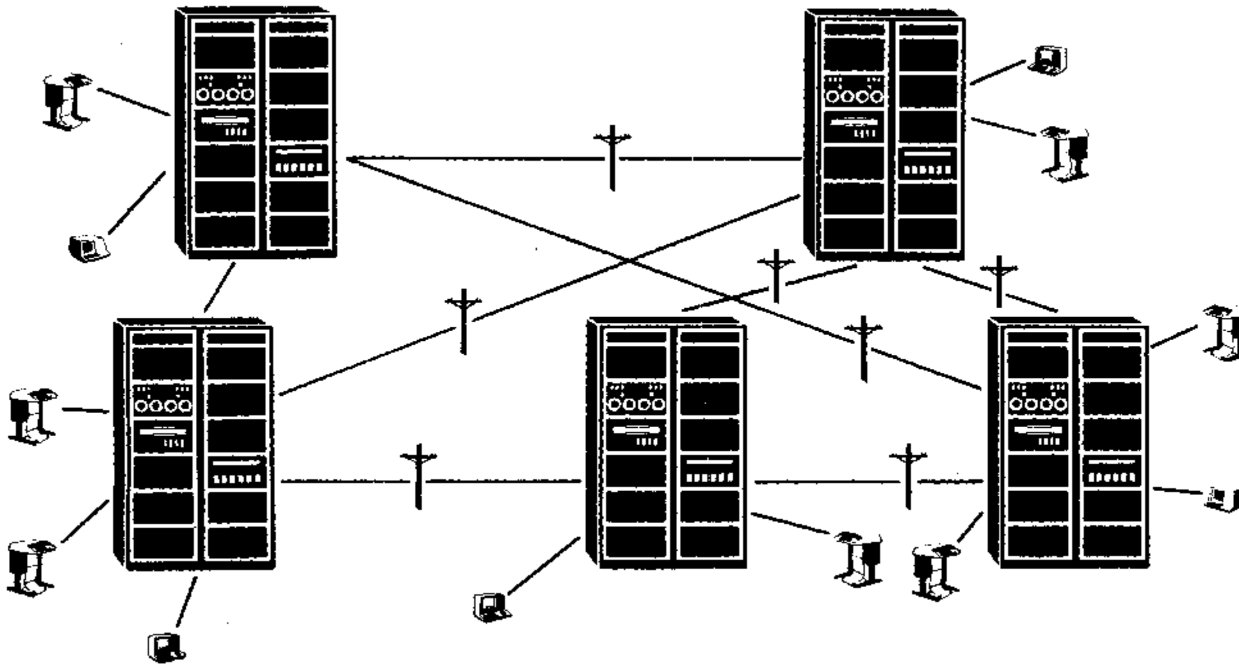

Figure 9 A Fully Distributed Network

points, each with its own set of users and a means for direct communication between the central points. Such a distributed structure appears to offer considerable advantage in reducing the cost of terminal communications by permitting installations to be located near concentration of terminals. If properly designed, distributed networks can offer significant reliability advantages, since a failure at one node does not affect the rest of the network.

Fully Distributed Network

In applications where the reliability of continuous communications is important, a fully distributed network (see Figure 9) in which every point is connected to several neighboring points may be preferred. The additional transmission paths provided by this type of structure improves the overall performance of the network. In designing this type of network, a traffic analysis must be performed to determine where the links are required.

Few fully distributed networks have been implemented because distributed data bases are difficult to handle and communications cost has been minor when compared with computation investment. With solutions of the data base management problem and the decreasing cost of computer systems and terminals, it is likely that many more distributed systems will emerge in the future.

MESSAGE SWITCHING AND ROUTING

There are three basic techniques for switching (and routing) communications traffic from source to destination—circuit or line switching, message switching, and packet switching.

Circuit Switching

In a circuit-switching network, the role of the switching center is to establish a direct connection from one terminal or computer to another. After the connection is established, the terminals or computers carry on their one-way or two-way communications. When communication is completed, the switching centers disconnect the circuit and restore the system to readiness for other connections.

The public switched telephone network and the TWX and TELEX teletypewriter services operate as circuit-switching networks. In these systems, calls are routed from one terminal or computer to another through different switching centers. Every time a call is originated, the proper electrical path must be established and switched into the network to provide an interconnection between the calling and the called terminal or computer. Interexchange and long-distance calls may, and usually do, establish a totally new path between the same terminals and computers every time a new call is originated.

Direct Distance Dialing (DDD) uses an Electronic Translator System (ETS) in the switched-dial network to automatically direct long-distance calls to their destination by the most readily available route. In the event all long-distance trunks are in use, the caller will receive a busy signal or a prerecorded announcement.

Circuit switching currently requires connect times and ties up transmission capacity for long periods. However, the connection, once made, has no delay. There are also networks that use circuit-switching techniques for routing of message traffic. Once a path is determined through the network nodes, all traffic between a source and destination pair follows the same path.

Message Switching

In message switching, each message is sent into the network and is routed as a unit to its destination. Subsequent messages for the source/destination pair may take a different route. Thus, we have message switching instead of circuit switching. The message makes its way through the network with the destination address given in the header preceding the message text telling each station in the network where to forward the message. Since some stations may be busy, the message must often be stored at intermediate stations; thus, it is also called a "store and forward" system.

Packet Switching

Packet switching is essentially a form of message switching. The major difference is that long messages are divided into fixed-length segments—called packets—at the source, and each packet of perhaps 1,000 to 8,000 bits is treated individually and forwarded along the best available route; that is, the route with the shortest transmission delay. Each packet is checked for errors at each node along the way, and at the destination another minicomputer reassembles related packets into complete messages for the subscriber's use. Short messages may fit into one packet. Long messages arrive at their destination faster because they are broken up and sent over different routes.

Packet-switching networks more fully utilize the relatively expensive wideband lines, and users enjoy the resulting reduction in cost. Since the intermediate nodes in a packet-switched network see only parts of messages (packets) and are thus unable to assemble the entire message, transmission is more secure.

chapter 3
network functions

INTRODUCTION

As illustrated in Figure 10, a data communications network consists of an interconnected group of computers, each of which may have human as well as data base interfaces. They can range from as simple as an intelligent terminal linked to a host processor, to a complete network of terminals, concentrators, remote processors, and multiple host computers. Each processor in such a system is selected on the basis of its ability to perform the task required at its location. The individual processors may vary considerably in their speed, interfaces, languages, and other characteristics. It is only important that they be efficient in the performance of their assigned tasks.

HOST PROCESSING

The host processor might be the large central computing facility in the typical timesharing system. It can receive and process a general class of requests from a variety of local or remote devices, including computers and terminals. The host usually has a large internal memory and is connected to an array of peripheral devices, including disk storage devices capable of storing vast amounts of information that make up the data base. The host software usually includes a sophisticated operating system supported by communications control software designed specifically for participation in network activity.

In addition to performing the bulk of the main data processing, the host processors comprising a network control data bases and generally supervise operation of the network. They can share such resources as processing power, programs, data bases, and peripheral devices with each other.

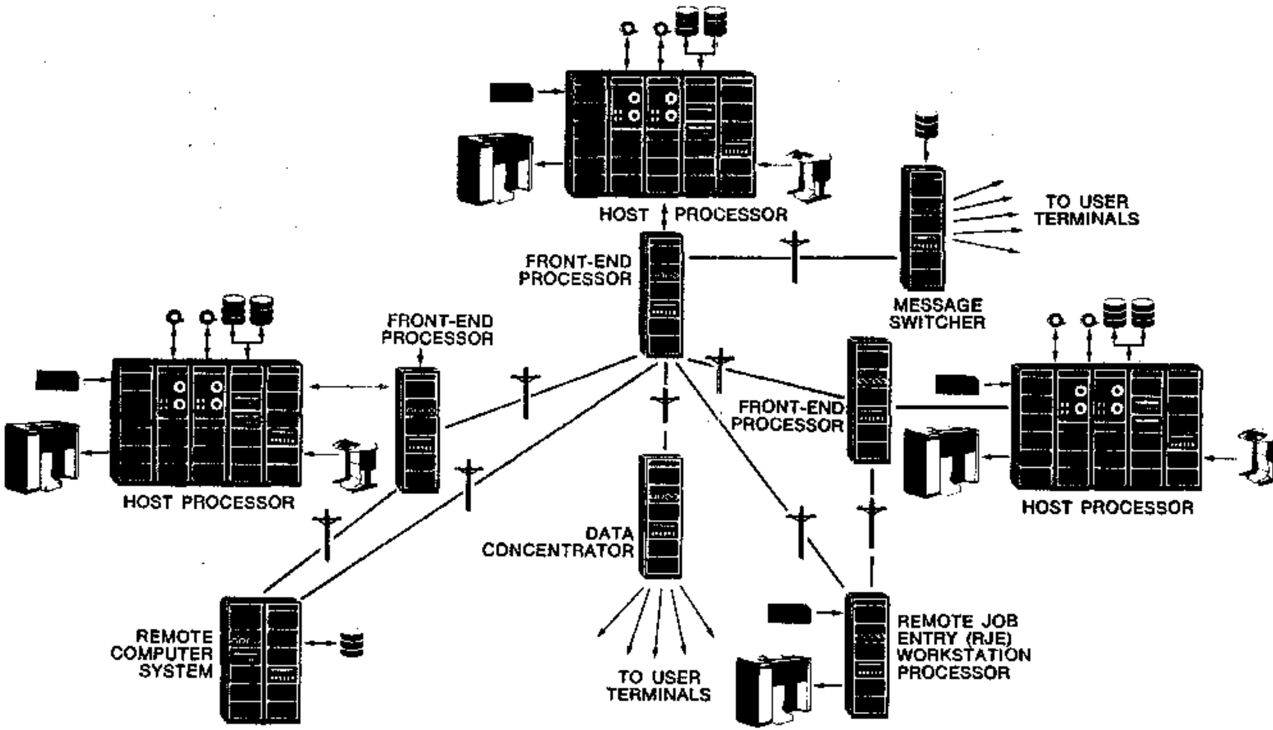

Figure 10 A Data Communications Network

A host processor connected to a front-end communications preprocessor is shown in Figure 11.

FRONT-END COMMUNICATIONS PROCESSING

The concept of front-end processing involves removing the data communications control function from the host processor and setting it up as an external, largely self-contained system. The minicomputer used for this function is placed in the data stream between remote data terminals or computers and the host processor. Unlike the remote data concentrator or remote message switcher, the front-end is local to the host processor.

The front-end relieves the host processor of such communications tasks as line control, error checking, code conversion, automatic answering, polling and addressing, and character-to-message assembly and disassembly. A communications system with a front-end processor can incorporate, without changes in the host processor, different types of lines, speeds, codes, and devices. Since it controls communications functions independently, the front-end reduces processing time and main memory requirements in the host processor. The amount of memory and processing times used for all of these tasks are lower in cost, easier to implement, and more flexible in the front-end than in the host processor. By freeing significant amounts of memory and processing time, the front-end extends the working life of the host processing facility.

General-purpose minicomputers have gained wide acceptance as front-ends to larger computers because they are more flexible and more economical than fixed or hardwired systems. The general-purpose minicomputer possesses many of the features inherent in larger general-purpose computers, such as large memory capacities, arithmetic units, large mass storage, and the ability to use high-level languages. Because of these features, the capabilities of the modern programmable minicomputer are not limited to communications control functions. For example, the front-end can maintain network operation in the event of host system failure. By utilizing its peripherals, it can continue to collect or route traffic while maintenance is performed on the host configuration. If the front-end is performing a message switching function, there may be little or no effect on the rest of the network when the host system fails. If processing transactions from remote stations, it can capture and store data for later processing by the host processor.

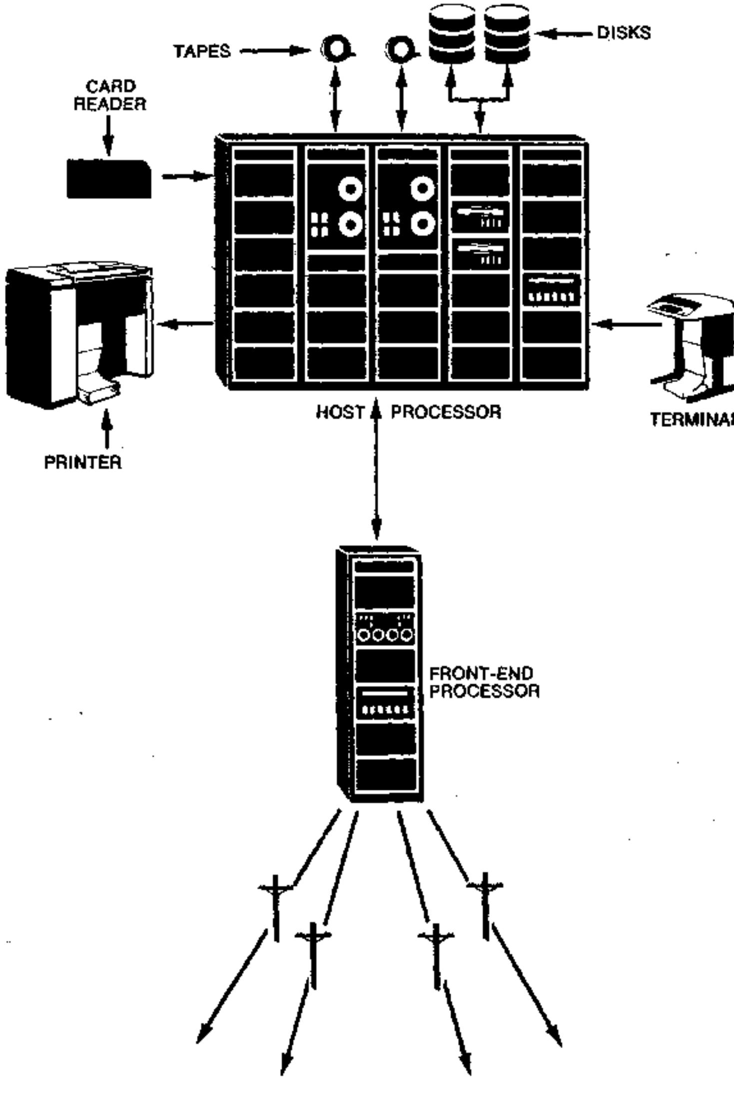

Figure 11 Host and Front-End Processors

Since front-ends can be configured with their own peripherals, they are capable of storing a portion of the system load. This can be a distinct advantage during peak load situations and can help smooth system throughput. The programmable front-end can also be configured as a standard data processor. When not functioning in a communications mode, it can be used to perform data processing.

Special performance features can be provided by a front-end. For example, in many applications it is desirable to identify specific terminals as they generate messages. A front-end can be programmed to identify the caller and to provide the proper interface between the terminal and host processor. Another communications feature allows direct access to the network for monitoring and evaluating common carrier performance, providing instant notification of network faults.

Because of its flexibility, the front-end minicomputer can easily be adapted to handle future network modifications. New system inputs, generated from changes in the number of terminals supported or higher traffic loads, can easily be handled by programmable front-ends. The ease with which such changes can be made relates to the over-all system being broken down into more manageable, functional segments.

REMOTE CONCENTRATION

The programmable remote concentrator, shown in Figure 12, helps reduce line costs by accepting messages from many terminals via slow-speed lines and transmitting them to the host processor via a single high-speed line. A concentrator is essentially a front-end processor that is not local to the host processor. It may be used wherever several communications paths exist over long-distance lines. If, for example, there are several locations remote from each other but in the same part of the country, it is more cost effective to multiplex slow or medium-speed data from these locations into one high-speed line with a concentrator rather than have a dedicated high-speed line for each location.

When used as a remote concentrator, the minicomputer functions as a storage buffer with input/output capabilities. Thus, the concentrator may alter the signalling speed of the message in addition to regulating the flow of messages over the lines.

The processing speed of the concentrator is sufficiently rapid so

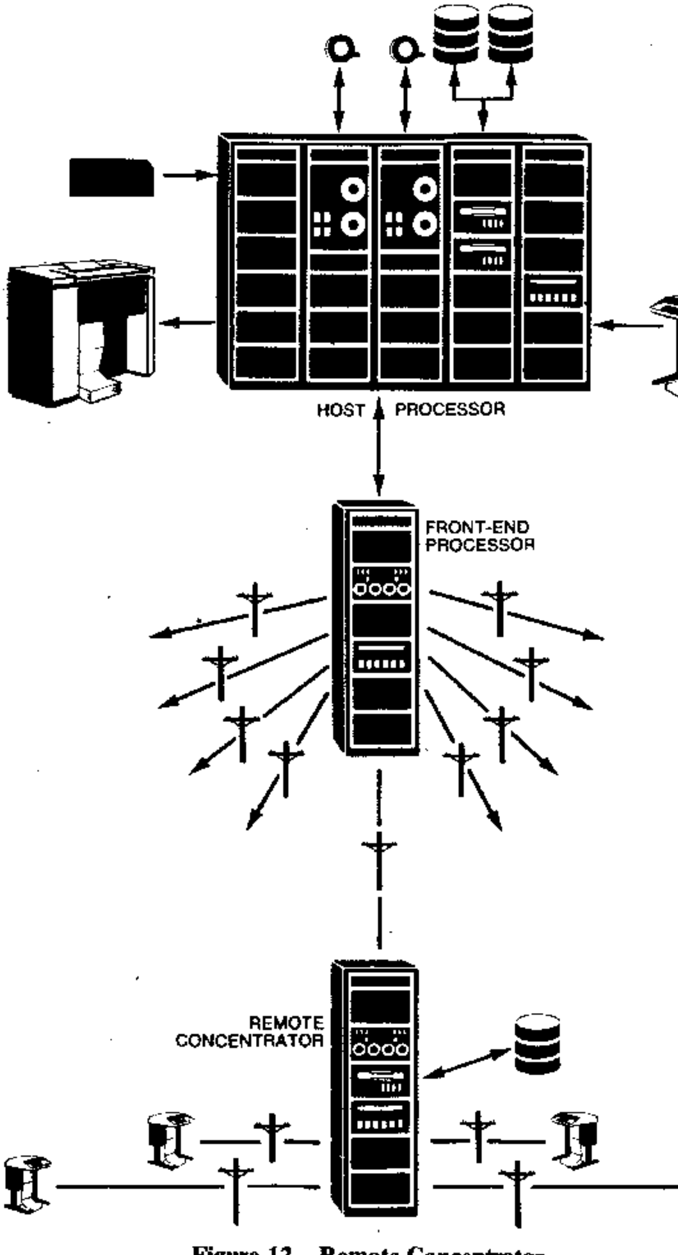

Figure 12 Remote Concentrator.

that it can accept messages simultaneously from several slow-speed terminals, thus reducing terminal delays in waiting for an available circuit to the host processor. Like the front-end processor, the concentrator can be programmed to perform such functions as character-to-message assembly and disassembly, communications line control, message buffering, code conversion, error detection, and automatic answering.

Because of its inherent stored program flexibility and modular expandability, the general-purpose minicomputer offers significant advantages when used as a remote concentrator. These advantages include accommodation of interfaces to special terminals; buffering capability which, by smoothing out peak loads, allows a higher concentration ratio than is available in a hardwired multiplexer; and accommodation of changes in data rates, formats, codes, communication procedures, and number of terminal devices.

The modern minicomputer concentrator has sufficient storage capacity and processing power to perform applications-oriented processing in addition to message receipt and transmission. In inquiry/response systems, for example, the minicomputer concentrator can be used to receive inquiry messages from remote terminals, process them to determine the specific information required, retrieve the information from on-line random-access disk storage units. and send it back to the inquiring terminal.

REMOTE COMPUTING

The term "Remote Computing" refers to two remotely located but functionally different classes of processing systems, as illustrated in Figure 13. One is the traditional Remote Job Entry (RJE) terminal. The other is a local general-purpose processing facility that may operate on a stand-alone basis or participate in network activity. This latter type is called a "Remote Computer System."

Remote Job Entry (RJE)

The Remote Job Entry or remote batch terminal enables the remotely located user to submit batch or production jobs to the host computer for processing and receive job output from the host. Typically, the RJE terminal receives batch input (usually from a card reader), processes certain control cards, transmits the job to the host processor, receives output results from the host, and routes these results to designated peripherals or files. The host processor

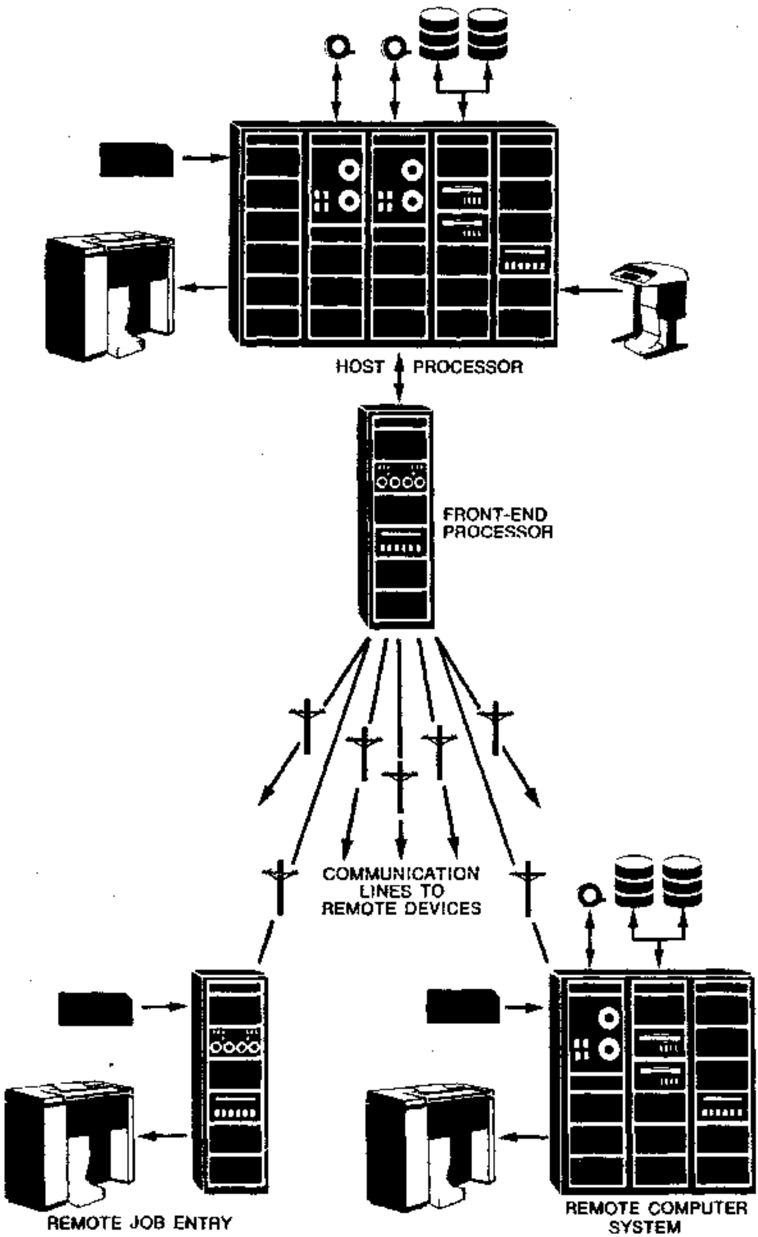

Figure 13 Remote Computing

accepts, queues, processes, and returns job results to the RJE terminal.

RJE terminals have traditionally been non-programmable hardwired devices generally limited to a card reader and line printer. However, there has been a trend toward remote batch terminals with more extensive local capabilities, including:

- Alternate methods for entering data other than cards;
- Local mass storage;
- Alternate devices for receiving data other than line printers;
- Local processing facilities.

One increasingly successful variant of Remote Job Entry is the use of a minicomputer to perform the functions normally associated with the hardwired RJE terminal. Such an arrangement can substantially enhance operations in a facility that requires some RJE activity. A programmable minicomputer-RJE system can give a facility flexibility and extended features. When not being used in a "dedicated" function for Remote Job Entry, the minicomputer can be used as a remote computer system in its own right.

Remote Computer System

Unlike the traditional hardwired RJE terminal, the remote computer system is not dedicated to processing batch or production jobs, although it may perform this function.

The programmable remote computer system is capable of processing a variety of applications that would otherwise have to be performed by the host processor. Thus, it relieves not only the communications load, but also that portion of the processing that can be carried on remotely, often closer to the point of entry. With a remote computer system, the user can maintain a local data base, perform local processing, and still communicate with the host processor when necessary.

In a true network situation, the remote computer system can directly access the data base files of the host computer. This puts a powerful resource at the disposal of the remote computer user and greatly extends the range of applications available.

When the host processor is operating at peak activity, the remote computer system may relieve the host of some of the processing load. This permits more effective utilization of each of the interconnected computers, because the slack time in one computer's schedule may be used to help smooth out the peaks in the other's.

MESSAGE SWITCHING

Message switching provides an alternative to line switching for handling messages over communications networks. The entire message is transmitted to an intermediate point (i.e., a switching computer), stored for a period of time, and then transmitted to its destination. The destination of each message is indicated by an address integral to the message.

A minicomputer-based message switcher does considerably more than route traffic. It lets a terminal send a single message with multiple addresses, eliminating retransmissions for each address. It can temporarily store data on disks or magnetic tape, convert codes, edit, log, and poll and address terminals. To do these functions without a minicomputer would require substantial investment in host processor core space and software.

Message switching concepts parallel many of the functions described previously under the description of concentrators, the major difference being the long-term store-and-forward characteristics of the message switcher. This means that the message switch processor is capable of storing messages it receives on on-line auxiliary storage units, such as disks or magnetic tape. The length of time messages are stored can range from a few seconds to an entire day or more, depending on the specific application needs and traffic volumes.

Typical functions performed by the programmable message switcher include assembly and disassembly of messages, polling and addressing terminals, line control, error control, code and speed conversion, time and date stamping, message routing, and message header analysis.

Savings in line costs can be considerable, since the message switcher can transmit to the receiving terminals at top speeds, several terminals can share one line, and messages may be packed together to make good use of the line.

In a multiprocessor network, the message switcher may transmit between remote terminals without going through the host processor. In so doing, it conserves host processor time and storage capacity and saves the communications line costs involved.

Without a message switcher, users must send data to all other locations either one at a time, or utilize a separate line for each

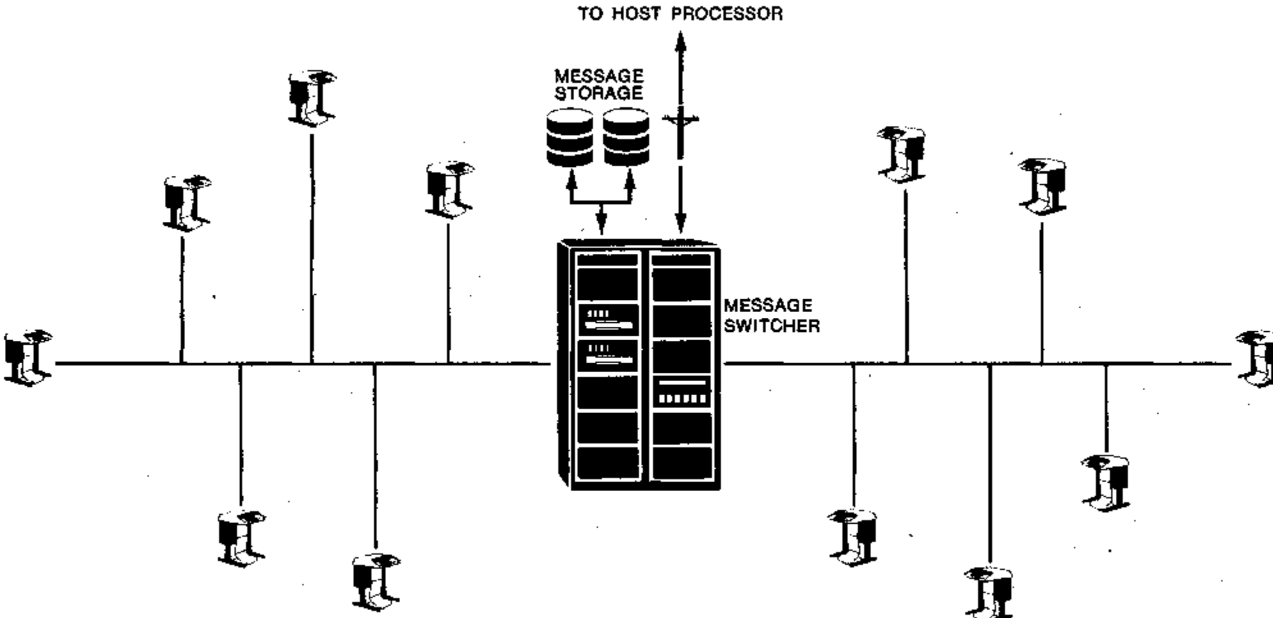

Figure 14 Computerized Message Switching

location. With a message switcher, they can send the data once over a single line and have the information automatically distributed by the switcher. Message switching also benefits users on busy traffic routes, where it might take some time to contact the message recipient. In such situations, the user stands a better chance of transmitting to the switcher without delay, where the message can be stored for transmission to the receiving terminal at the most appropriate time.

chapter 4
network building blocks

HARDWARE COMPONENTS

General

Many hardware components comprise a computer communications network. First, there are the host computers themselves, with their local peripheral equipment. The host computers serve as the primary processing resource for users of the network. In addition to the host computers, a typical network might consist of several minicomputers devoted exclusively to network control functions, such as data concentration or message switching. Connected to the host computers is the local communications hardware, which may be a programmable minicomputer used to remove the high communications traffic overhead from the host processors. Then there are the line interfaces, the modems, and the communications channel or facility. Several remote computing systems and a variety of terminals may be included in a typical network. The components required depend upon a number of factors, including volume of data to be handled, functions to be performed by the system, responsiveness required, susceptibility of the system to errors, transmission speeds, etc. Figure 15 illustrates some typical hardware components at one node in a network. Other nodes in the net will have identical or similar components.

Communications Channels

The communications channels are the paths for transmitting signals. They are typically obtained from the common carrier; channels may access the public switched telephone network or the switched telegraph network, or may be obtained on a leased private line basis.

A variety of communications channels are available from the carriers to meet data transmission requirements. These facilities can be divided into the following three classes.

- Narrowband—These facilities provide data communication capabilities at up to 300 bits per second.
- Voiceband—These facilities make use of communications channels having effective bandwidths of about 3,000 Hertz. Equipment is available from the carriers and independent suppliers for data transmission at speeds up to 9,600 bits per second.
- Broadband—These facilities provide data-communication rates higher than those of voiceband channels. They also provide higher reliability. Current facilities can provide transmission rates up to several million bits per second.

See Appendix B for additional information on available communications channels.

Modems

A modem is a device capable of changing or converting information-bearing signals from one form to another. In data-communications applications involving digital data and analog telephone lines, the modulator portion of a modem converts digital d-c pulses, originated by computer or terminal equipment, to analog, wave-like signals acceptable for transmission over telephone lines. The demodulator reverses the process, converting the analog telephone signal back into a pulse train acceptable to the computer or terminal. If a modem was not used to convert/reconvert data signals, and the computer or terminal was directly connected to the telephone lines, the signal would be degraded and the data garbled or made unintelligible by the electrical characteristics of the line.

Modems transmit data in spurts (asynchronous modems) or in steady streams (synchronous modems). Asynchronous modems are usually associated with keyboard entry terminal devices (teletypewriters and CRT display terminals) where the time between information segments (characters) is random. Synchronous modems are used with continuous data sources (punched card readers, paper or magnetic tape equipment, and computers), where each character follows the preceding one at a fixed interval.

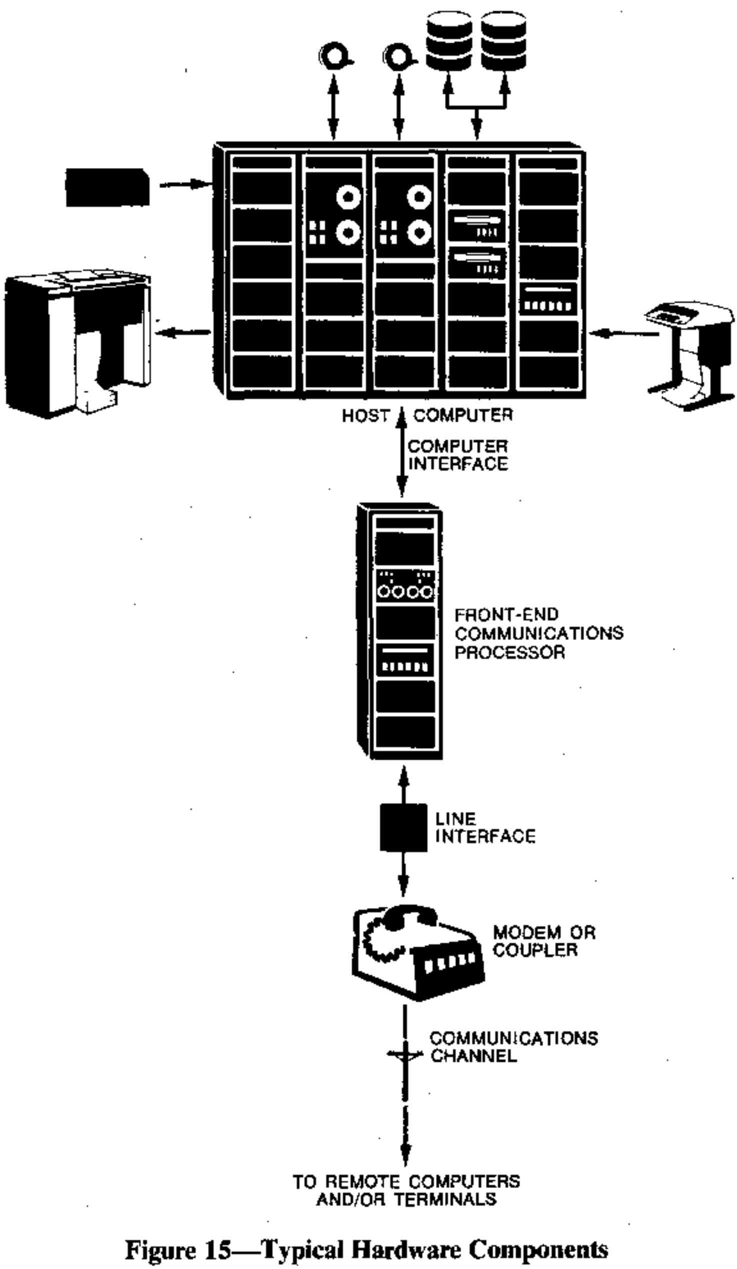

Figure 15—Typical Hardware Components

Modems may operate in three different modes: in the simplex mode, where data is either transmitted or received; in the half-duplex mode, where data can be both sent and received—but not simultaneously—similar to an ordinary voice telephone conversation; and in the full-duplex mode, where data can be simultaneously transmitted and received.

The most commonly referenced operational characteristic of a modem is its speed or data rate, expressed in Baud or bits per second (bps).

Within the Bell System, modems are referred to as Data Sets or DATA PHONE SETS. For data communications users, modems are available from the common carriers or, subject to certain restrictions, may be purchased outright from independent equipment suppliers.

Acoustic Couplers

An alternative to the modem is the acoustic coupler, a device that accepts a serial data stream from a data-processing device, modulates it in the audio spectrum, and produces the modulation as an audible tone. Acoustic couplers are equipped with cradles or fittings that accept a conventional telephone handset and couple the acoustic energy directly into the mouthpiece. At the receiving end, a similar device picks up the audible tones from the telephone earpiece and demodulates them to a serial data stream.

Acoustic couplers enable any conventional telephone to be used as a data terminal.

The principle limitation of acoustic couplers is data rate; because of the properties of carbon microphones, they cannot pass as wide a band as the faster voice-line modems. Typical maximum rates are 300 and 600 bps, although a few units can handle 1,200 or 1,800 bps. Acoustic couplers are widely used with teletypewriters and other relatively slow-speed terminals.

Line Interfaces

Most of the elements in a communications network—such as terminals, modems, and communication lines—are available in a wide variety of types and capabilities and thus offer ample flexibility for precisely tailored systems, initially and as the user grows. However, in order for a computer to communicate over a particular line and with a particular terminal or another computer, the computer manufacturer must provide a communications interface to handle that specific line-terminal or line-computer combination. The interface is therefore one of the keys to flexibility and the cornerstone to economical systems design. The greater the variety of communications facilities that a host or remote computer can handle, the greater the chances for a system of optimum design.

The line interface provides a way for the remote device (terminal or computer) to talk to the host computer site. It usually conforms to ASCII code and discipline, meets the electrical and logical requirements of the Electronic Industries Association (EIA)RS 232 standard, and connects to a modem or acoustic coupler at speeds between 110 and 9600 bits per second.

In some cases (where the terminal-to-computer distance is not too great) modems can be eliminated and the terminal can be connected to the computer directly.

The computer interface connects the communications hardware to the host computer. The interface will differ depending on the type of host and the particular input/output path used. In some cases, the interface may be an integral part of the communications hardware. Direct connections are often parallel and provide very high data transfer rates.

Digital's PDP-11 family of processor's are available with a full range of programmable single and multi-line asynchronous and synchronous interfaces. Other data communications options are availble for error detection, automatic calling, and signal conditioning. The characteristics of these options and line interface units are summarized in Appendix A.

Network Computers

The selection of computers to serve as host or remote processors or to perform control functions in data communications networks is of vital importance. Many computers were not designed for easy interfacing to a variety of communications lines and devices. Others are far too expensive to justify the establishment of a network involving more than one computer.

A rapidly growing trend is to design a network around an array of general-purpose minicomputers. One reason for this trend is the new series of interconnecting hardware and software that minicomputer manufacturers are offering, off-the-shelf, to their customers. Previously, the custom-built interface designs that were needed to implement multi-computer networks made such systems, in the main, economically unfeasible.

Because of these new standard interfaces, more and more multi-minicomputer systems are being used in a variety of applications. In some systems, for example, minicomputers are combined simply to add redundancy when high reliability is needed.

In other applications, multi-minicomputer systems allow the kind of distributed processing that is desirable in widely dispersed data-communcations networks. And in still other applications, arrays of minis have been linked together to provide raw processing power surpassing that of a medium or large-scale computer—at significantly lower cost.

In summary, minicomputers have clearly established their ability to manage sophisticated communications networks with high performance and great reliability while significantly reducing operating costs. They are finding increasing use as host and remote processors, as data concentrators, as message switchers, and as front-end communications processors.

One of the most popular choices has been DIGITAL's 16-bit PDP-11 computer, whose unique features render it a natural for data communications applications. These features are summarized in Appendix A.

Communications Terminals

The many different types of terminals employed in computer networks include typewriter-oriented terminals, Cathode-Ray Tube (CRT) display terminals, and intelligent terminals.

The most significant consideration affecting terminal selection is matching the functional capability of the terminal to the functional requirements of the application. The best evidence of this importance lies in the variety of special-purpose terminals developed for such applications as airline reservations, on-line banking, industrial data collection, and brokerage transaction processing. The application primarily sets the requirements of speed, permissible error rates, input/output medium, information codes and formats, and storage in a terminal.

TYPEWRITER-ORIENTED TERMINALS

Typewriter-oriented terminals are so called because the interface with the operator is much like that of the standard office typewriter. Ordinarily these terminals consist of a keyboard, a printer, and the necessary control and interface circuitry. Often they include auxiliary storage devices for off-line message preparation.

Typewriter-oriented terminals were originally designed for telegraphic communications with similar devices over wire or radio communications facilities. Today, their usefulness includes data communication with a remote computer. In all cases, a typewriter-oriented terminal permits an operator to communicate at the same time that he visually verifies the transmitted and received messages.

Today's typewriter-oriented terminals have many advantages over their earlier counterparts. They can operate at faster speeds, provide more efficient error control, handle a variety of prepared input media, and produce output in many forms. Specialized keyboards have been introduced for specific applications. In addition, some typewriter terminals are equipped with buffers and line-control units. Recent advances in packaging techniques as well as reductions in circuit size have permitted the design of truly portable typewriter-oriented terminals.

Typewriter-oriented terminals are used in such applications as computer timesharing, file updating, data retrieval, data dissemination, and message conversion.

CRT DISPLAY TERMINALS

CRT display terminals employ a TV-like display screen to record and display data. They have a keyboard similar to that of a typewriter. When a key is depressed, the associated numeric or alphabetic character appears on the screen. The character code may be simultaneously transmitted or it may be stored for subsequent transmission as part of a word or a complete message. Sequences of codes are received by the terminal which cause messages to appear on the screen. Typically, 1,000 characters may be simultaneously displayed on the screen, and character transmission rates are in the range of 110-to-1200 characters per second.

CRT display terminals are widely used in inquiry-response applications as "electronic blackboards" to provide rapid access to data stored in computer systems. These terminals find use in such diverse applications as timesharing systems, computer-aided design and instruction facilities, on-line reservation and banking systems, and inventory control systems. CRT terminals are preferred due to their higher speed resulting in more efficient use of the telephone line and human reading capabilities, quiet operation, and the ability to selectively alter the displayed data.

The distinguishing features of CRT terminals are: display size area, display arrangement, total number of characters storable, storage facility, transmission speed, and other special features such as editing and control and auxiliary outputs.

A description of DIGITAL's product offerings in the typewriter-oriented and CRT display terminal class is presented in Appendix A.

INTELLIGENT TERMINALS

Intelligent terminals contain their own processing and memory capabilities. Such terminals can be programmed to perform functions that otherwise would require the services of the host processor. These functions include local batch processing, listing and formatting, logging, editing, file maintenance, and the control of communications between itself and the host processor. Some people call these terminals "smart" or "programmable." A more precise definition is "intelligent terminal system," since it can include a number of modular components such as CRT's, printers, tape drives, and disk drives.

SOFTWARE COMPONENTS

General

Data communications software consists of the programs and routines necessary to send data, commands, messages, and status from one computer to another. The time and cost involved in developing this software represent a substantial percentage of the total resources required to implement a computer-controlled data communications network. The maximum efficiency of any network is obtained by making trade-offs between hardware and software. There are no set rules that apply to all types of systems. Each system must be thoroughly analyzed for each application in order to arrive at the optimum solution. Two common problems facing any user of a communications network are flexibility and expandability. Each system must be flexible enough to handle different line rates, code structures, etc. In addition, systems should be easily expandable to handle increased line loads or to provide additional options, such as auto-dialing and auto-answering.

Software Structure

Figure 16 shows the relationship between the basic software elements found at one point in a communications network.

Since the Input/Output (I/O) driver is the closest piece of software to the hardware, it controls a number of data transfer functions. The structure of the I/O driver is tailored to the characteristics of the hardware interface it controls. The drivers pass control information and data from the line control modules to the hardware interface and pass status information and data from the interface to the line control modules.

Built upon the I/O driver base are the line control modules that are dependent upon the protocols of the devices and communications lines attached to the interface. These control modules implement parameters specified by the user application program, such as line protocol procedures, mode control, code conversion, etc.

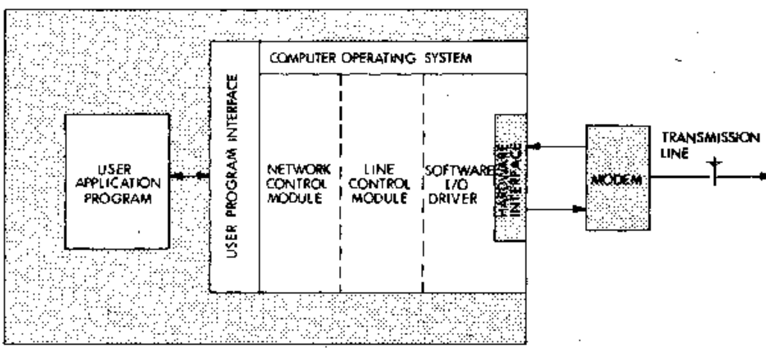

Figure 16—Basic Software Elements

The line control module provides an error-free communications envelope for exchanging blocks of data or packets between computers connected by a data link. To achieve meaningful exchanges of information in a network, a set of commands controlling connection and data flow must be provided to the user programs. A message structure must be imposed on the data transmitted over this envelope to facilitate implementation of user functions such as file transfers and remote terminal support. This message structure and its operation is provided by the network control module.

The basic function of the network control module is to support interprocess communication in the network. It is concerned with the creation of logical data links between processes, the routing of messages over these links, and the acknowledgment and diagnosibility of message flow in the net. The network control module constructs logical data links connecting user processes. It multiplexes these links into a single data stream when destined for transmission over the same line.

While the I/O drivers, line control modules, and network control modules may be part of the computer operating system provided by the computer manufacturer, the application program, that is, the software component that is unique to the application, is left to the user. These programs are usually written and maintained by professionals in their particular field, be it finance, government, transportation, or engineering. The user application program interfaces with the operating system via a set of easily understood commands.

Operating Systems

Operating systems can be defined in several ways. The choice of definition depends upon the functions the system is designed to perform. Basically, an operating system is an organized collection of programs which increases the productivity of a computer by providing common functions for all user programs.

Depending on its complexity, an operating system can be responsible for any or all of the following functions:

1. Control and monitor program execution
2. Manage system resources, such as memory and I/O devices
3. Control input/output devices
4. Store and retrieve data
5. Store and retrieve programs
6. Prepare programs in one or more programming languages

Operating systems monitor applications programs to ensure that processing and input/output operations are performed concurrently. To further maximize throughput capability, operating systems are designed to automatically direct the computer in job-to-job transition. The program segment of an operating system responsible for task and job management is called the "executive"; it is also referred to as the "supervisor" or "monitor." It is the most sophisticated of the programs comprising an operating system.

The preceding objectives of an operating system are equally important to any data processing application. Ensuring optimum system accessibility is of special significance to data communications networks. It is primarily the responsibility of the communcations control segments of the operating system. The control modules and

I/O drivers depicted in Figure 16 are examples of such control segments. It is the function of these control programs to do all the "housekeeping" tasks associated with communications systems, such as: polling and addressing of terminals (on a leased party line); automatic dialing of terminals on a public switched network; buffer management—the allocation of core storage for each incoming transaction; queuing, logging, and sequencing of messages (in a message switching system); character-to-message assembly and message-to-character disassembly; and error checking. The communications control segments of the operating system perform these functions independent of the applications program.

Operating systems can be classified in several different ways. For example, the systems may be available in core only or in core-disk versions. Some are designed primarily for use in terminal-oriented networks, such as in computer timesharing systems where many users communicate with one central computer from remote interactive terminals. In interactive processing, each unit of data may be as small as one character. Interactive systems are often referred to as "conversational" systems because the fast response time makes it appear that a terminal user is holding a dialogue with the system. Interactive processing usually implies "on-line," meaning constantly available system resources, particularly data files.

Other operating systems are designed to permit communications between terminals and computers or between two or more computers. These operating systems form the basic foundation of the software required in implementing resource-sharing networks where many computers and terminals are linked together.

"Real-time," "multiprogramming," "multitasking," "multiprocessing," "response time," and "throughput," are terms frequently used in discussions of operating systems. These terms are discussed in the remainder of this chapter.

Definition of Terms

REAL-TIME pertains to the performance of a computation during the actual time that the related physical process occurs so that the results of the computation can be used in guiding the physical process. Real-time systems collect data ,monitor instruments, and control laboratory experiments. For example, in an oil refinery, the computer measures temperatures and pressures using a multitude of sensors in liquids and atmosphere, and transducers connected to pipelines and containers. Real-time programs accept data, produce status reports, adjust physical factors in the external process, and notify operators in special cases.

Real-time systems are also finding increasing use in many banking and brokerage transaction processing systems, involving multiple terminals.

In real-time systems, the occurrence of an external event implies that the computer system has to take an action since something of importance has occurred. For example, a ring indicator signal from a modem indicates that a call should be answered by the system. The action taken by the computer system as the result of such an event is the execution of a piece of code or program which is referred to as a "task." Because real-time systems usually contain several programs to control different conditions or processes, several external events may require the attention of different programs simultaneously.

MULTIPROGRAMMING is a resource management strategy in which an executive routine allocates the resources of the computer (such as core, I/O channels, disk storage, and execution time) to many programs concurrently. Multiprogramming allows a single computer system to execute a multiplicity of real-time tasks concurrently. Multiprogramming increases the efficiency, response speed, and flexibility of the computer because logical delays in one program (e.g., for input) are used to execute other programs. As a result, a single processor interleaves the execution of two or more programs. The operating system has a scheduling module which decides when to execute the competing program. Each program has a priority level assigned by the user. When contention occurs, the operating system decides in favor of the higher priority program.

MULTITASKING is a program design strategy in which the various logical elements making up a program are written so that they may operate asynchronously with respect to each other. Multitasking is sometimes referred to as "asynchronous processing" (different from asynchronous communications). Splitting the program up into several asynchronous tasks allows a single program to exploit many of the benefits of multiprogramming for itself. In distributed computing networks, it is often possible to move some of the tasks that are part of a multitasking program into other computers in the network, thus increasing the power of the system.

MULTIPROCESSING normally implies multiprogramming, but actually means that there are two (or more) central processors performing as one system. This configuration enables the system to perform overlapped simultaneous processing to a greater extent than does a single computing system. Multiprocessing configurations may consist of a host computer and a front-end processor or both may be host computers. But in either case, both are processing (either on segments of the same job or on entirely different jobs) simultaneously.

RESPONSE TIME is the time interval between the terminal operator transmitting his message and the time the system responds with an answer. This can also be called the system turnaround delay. It is affected by transmission time, interrupt servicing time, message processing time, wait time, mass storage access and data transfer time. It is important that the operating system be optimized to minimize the delays caused by these fatcors via such design features as overlapping I/O, double buffering, multitasking, and efficient handling of data. The latter feature is accomplished by moving only the pointers to data and not the data itself within the processor's memory. The term "pointer" refers to data that is used to specify the location of other data.

THOUGHPUT is defined as the total amount of work a system can perform in a given period of time. Operating systems employed in modern computer networks are designed to ensure that the internal operating elements of the computers are utilized at their maximum efficiency.

DIGITAL's operating systems, described in Appendix A, include all the capabilities described earlier, including real-time operation, multitasking, multiprogramming, and multiprocessing.

chapter 5
line control procedures

GENERAL

In designing computer networks one of the basic considerations is the physical transmission of data from one computer to another. In the absence of transmission errors this becomes a relatively simple task. Once errors are introduced, however, problems of correct data sequencing and synchronization of the transmitter and receiver appear and must be solved. The solution consists of a data communications link protocol which assures the correct sequencing and integrity of data transmitted between computers and between computers and terminals in a network.

By using defined control characters, the link protocol provides an orderly and efficient way of assuring that, among other things, a remote terminal or computer is in a ready condition, and that the remote device will send data when instructed, will receive data when instructed, and will advise the sending terminal or computer when it receives erroneous data. Since the same physical link carries both data (text) and control characters, the protocol must be capable of distinguishing between the data and control characters available within the code set.

There are a number of protocols available today. To understand how they differ and to make judgments about the advantages of each, it is necessary to have a basic understanding of the functions line protocols perform. The next section describes these functions in detail. This is followed by discussions of three protocols, with emphasis on how each one handles the various functions. A brief introduction to Standards organizations and their activities with line control procedures is then presented. This is followed by a comparative summary of the functions performed by five different protocols.

PROTOCOL FUNCTIONS

The functions performed by protocols include:

- Controlling data transfers

- Error checking and recovery
- Information coding
- Information transparency
- Line utilization
- Synchronization
- Communications facility transparency
- Bootstrapping

Controlling Data Transfers

Data transfers are controlled by three elements: formatting, control information, and handshaking procedures. Formatting means reserving positions, or fields, in the transmission block for specific information. Control data and error checking data must be included in the transmission block. These are usually called the header, body, and trailer of a block as shown in Figure 17.

HEADER OR CONTROL FIELD	BODY OR TEXT FIELD	TRAILER OR ERROR CHECKING FIELD

Figure 17—General Block Format

Some protocols delimit these fields with control characters. For example:

SOH—Indicates Start-of-Header information

STX—Indicates Start-of-Text

ETX—Indicates End-of-Text and implies that the information following is the error-checking information.

Protocols use information code sets to distinguish control characters from information characters and to distinguish one information character from another (i.e., the letter "A" from the letter "B"). The various information codes in use today are discussed briefly later in this section.

To control the flow of information, the header usually contains addressing, block sequencing, control flags, and acknowledgment information. The addressing information identifies the destination (and possibly the source) of the data. It is used to route the data from source to proper destination.

Block sequencing ensures that the data arrives in the proper sequence and that no blocks are lost or duplicated.

Control flags indicate whether the information being transmitted is a data or control-only message, or whether it is the first, intermediate, or last block of a message.

Control messages are used to determine who transmits and who receives, to acknowledge good or bad reception of blocks, or to abort transfer sequences. These procedures are usually called "handshaking" procedures. Some of these procedures are recovery procedures, as discussed under "Error Checking and Recovery." Acknowledgment information gives good or bad status with respect to messages sent in the opposite direction. This information is part of the error recovery procedures. This field is often called the "response field."

Error Checking and Recovery

An important function of a line protocol is to assure correct reception of data. Communications facilities are error-prone. To compensate for this, line control procedures include the generation, transmission, and testing of check bits. These check bits (often called Block Check Characters—BCC) make up the trailer field of the transmission block. They are generated by a checking algorithm which is usually applied only to the information field of a block.

Each block of data transmitted is error-checked at the receiving station in one of several ways, depending on the code and the functions employed. These checking methods are Vertical Redundancy Checking (VRC), which is parity checking by character as the data is received; and either Logitudinal Redundancy Checking (LRC) or Cyclic Redundancy Checking (CRC), which check the block after it is received. After each transmission, the receiving station normally replies with a positive (ACK) acknowledgment (data accepted, continue sending), or with a negative (NAK) acknowledgment (data not accepted; e.g., a transmission error was detected—retransmit previous block). The acknowledgment may be in a special control message (ACK) or in the response field in the header of the next message to be sent in the opposite direction. Retransmission of a block of data following an initial NAK is usually attempted a number of times. Until the block is accepted, the data buffer cannot be released by the transmitting system.

VRC (Vertical Redundancy Checking) is an odd or even parity check performed on a per-character basis and requires a parity check bit position in each character. If individual characters are represented by eight bits, such as when using an eight-level code, seven may be used to represent actual numbers and letters, and the eighth may be reserved for checking purposes. The presence or absence of the eighth bit provides the inherent checking feature. For example, in an even parity check, the parity bit is used to make the total number of one bits in the character even. If the character contains four zeros and three ones, then a one bit is inserted as the parity bit.

Longitudinal Redundancy Checking (LRC) is a technique for checking an entire message or block of data. In this case, an exclusive "OR" logic is used for all the bits in the message and the resulting character, called the Block Check Character (BCC), is transmitted as the last character in the block. The receiving device independently performs the same counting procedure and generates a Block Check Character. It then compares its own BCC character with the one received. If they are not identical, an error condition exists, and the sending device is notified that an error condition exists within the block. LRC is frequently used in conjunction with VRC to increase the error detection capability within a system.

Cyclic Redundancy Checking (CRC) is a more sophisticated method of block checking than LRC. This type of error checking involves a polynomial division of the data stream by a CRC polynomial. The 1's and 0's of the data become the coefficients of the dividend polynomial while the CRC polynomial is preset. The division uses subtraction modulo 2 (no carries) and the remainder serves as the Cyclic Redundancy Check. The receiving station compares the transmitted remainder with its own computed remainder, and an equal condition indicates that no error has occurred.

There are many constants that may be used to perform the CRC division. Two of the most popular versions are called CRC-16 (which uses a polynomial of the form $x^{16} + x^{15} + x^2 + 1$) and CRC-CCITT (which uses a polynomial of the form $x^{16} + x^{12} + x^5 + 1$). Each generates a 16-bit BCC. CCITT, the International Consultative Committee for Telephony and Telegraphy, is responsible for usage standards.

Error checking also involves checking for sequence errors. Protocols handle this in different ways. These include alternating acknowledgments and block sequencing. The technique used depends on the protocol. The receiving station sends back an indication of a sequence error with a negative acknowledgment or some other control message.

Information Codes

To allow communications between various computers and terminals comprising a network, a uniform method of exchanging information is needed. This requires the establishment of a character code structure for interpretation of bits as characters, a message syntax to form characters into messages, and data communications control procedures for exchanging the messages.

A number of different coding schemes are used to represent characters in data communications systems. The codes differ primarily in the number of bits used to represent characters and the particular pattern of bits which correspond to the characters. Characters are divided into graphic characters, representing a symbol, and control characters which are used to control a terminal or computer function.

Among the many codes used in communications today is the 7-bit-plus-parity ASCII (American Standard Code for Information Interchange) code. ASCII was introduced by the USA Standards Institute and has been accepted as the U. S. Federal Standard. Techniques for transmitting "transparent" or binary data exist within the structure of the ASCII code. Special characters are set aside for the purpose of communications control. These control functions include synchronization, message heading and control.

A variation of the ASCII code is the Data Interchange Code. Primarily, this code differs from ASCII in that some printing characters are replaced by non-printing control characters, and the parity is specified to be odd. This code is now readily adaptable to computer-to-computer communications.

Of the other existing codes, the more widely used ones are the Extended Binary Coded Decimal Interchange Code (EBCDIC), the 5-bit Baudot code found in old teleprinter equipment, the Four of Eight Code, the IBM punched card Hollerith code, the Binary Coded Decimal (BCD) code. and the 6-bit Transcode.

EBCDIC is an eight-level code similar to ASCII except that while ASCII uses its eighth level for parity bits, EBCDIC uses it for information bits, thereby extending the range of characters to 256.

Information Transparency

It is often necessary to transmit binary data, floating point numbers, packed decimal data, unique specialized codes, or machine language computer programs. In order to do this, all data, including the normally restricted data-link control characters, are treated only as specific bit patterns. There must be a way to permit the use of all bit patterns in the information field and still control the transmission of the block. The technique used for achieving transparency differs for each of the protocols discussed in this chapter. The techniques used are summarized in Figure 25.

Line Utilization

The structure of a line control procedure has considerable impact on the utilization of the channel it controls. There are four factors affecting the utilization of the channel: direction utilization, control overhead, acknowledgment handling, and number of stations per line.

A line between two systems may physically permit one-way or two-way transmission. On a physical circuit this is called simplex or duplex operation. The two-way transmission may alternate in direction of transmission (called half-duplex) or provide simultaneous two-way transmission (called full-duplex). If the physical facility is full duplex, the line control procedure may or may not take advantage of this capability. In other words, it may be a half-duplex protocol (alternate data transmission), although the physical facility is full-duplex. To make most efficient use of a full-duplex facility, a full-duplex line control procedure is required.

In any transmission block there are either control bits, information, or error checking bits. All but the information bits are considered overhead bits. A physical facility is capable of transmitting a fixed number of bits per second in each direction, and the control bits detract from the effective rate of information transfer. The ratio of the information bits to the total bits determines the one-way line utilization. The more control, header and error checking characters needed by a protocol, the less efficient the line. Line efficiency is also affected by the way in which transparency is handled.

Acknowledgment handling can affect line utilization in two ways.

If the acknowledgment is a separate message, then it and the gaps between it and the data blocks are part of control overhead. Furthermore, there is more overhead if each message requires a separate acknowledgment. Acknowledgments in blocks containing information reduce the first overhead because it usually takes fewer characters for normal conditions. Only errors are indicated by separate blocks. If the line control procedure defines a way to acknowledge multiple blocks with one response, the number of overhead bits is further reduced.

When the activity from one station on a line is below full utilization including control overhead, the extra capacity can be regained by putting additional stations on the line. This is similar to telephone party lines and is called "multipoint" or "multidrop." When there are only two stations involved, it is called "point-to-point." Most protocols support both point-to-point and multipoint arrangements.

Synchronization

When transmitting a continuous stream of data bits (synchronous transmission), some form of synchronization is required between the sender and the receiver. The most efficient method is to precede a group of characters (message or block) to be transmitted with a unique group of bits called a synchronization sequence. The receiver searches for this unique sequence bit pattern so that it can properly frame and get into phase with the control or data characters in the incoming bit stream. Ideally the synchronization sequence and the transmission code should be so designed that the sequence could never be reproduced within the actual data stream, so that false synchronization might never occur. This is not always possible, especially when, as in ASCII code, all 128 combinations of the seven bits within characters are utilized. Protocols differ in the number of synchronizing characters used and the code of the characters.

Communications Facility Transparency

It would be extremely useful if the same protocol could be used regardless of whether the communications channel is serial asynchronous, serial synchronous, or parallel. Most line control procedures are designed for only one of these three facilities. The way in which data transparency is handled is the major factor affecting facility transparency.

Bootstrapping

Some of the computing systems at the end of lines for which the line control procedure is used will have their software loaded and the system restarted via the communications line. This starting up of an inoperative system is called "bootstrapping." The bootstrap procedure can be part of the line control procedure or it can be embedded in the text field.

BINARY SYNCHRONOUS COMMUNICATIONS (BSC)

One of the most widely used protocols in the industry is IBM's Binary Synchronous Communications (BSC). BSC, also known as BISYNC, has been in use since 1968 for transmission between IBM computers and batch and video display terminals. The way in which BSC handles the necessary protocol functions is explained in the following praragraphs.

Controlling Data Transfers: The format of a BSC transmission block is shown in Figure 18. BSC uses control characters to delimit the fields. The header is optional. If it is used, it begins with SOH (Start of Header) and ends with STX (Start of Text). The contents of the header is defined by the user. Polling and addressing on multipoint lines are handled by a separate control message and not by using the header field. The text portion of the field is variable in length and may contain transparent data. If it is defined as transparent, it is delimited by DLE STX and DLE ETX (or DLE ETB). The trailer section contains only the BCC. DLE is an abbreviation for Data Link Escape.

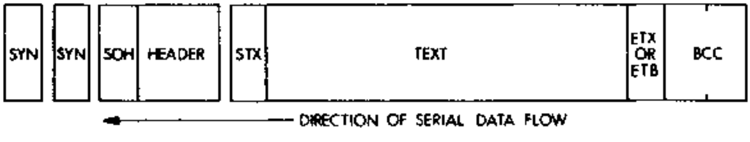

Figure 18—BSC Message Format

Binary Synchronous Communications employs a rigorous set of rules for establishing, maintaining, and terminating a communications sequence. A typical exchange between a terminal station and a computer on a point-to-point private line is illustrated in Figure 19.

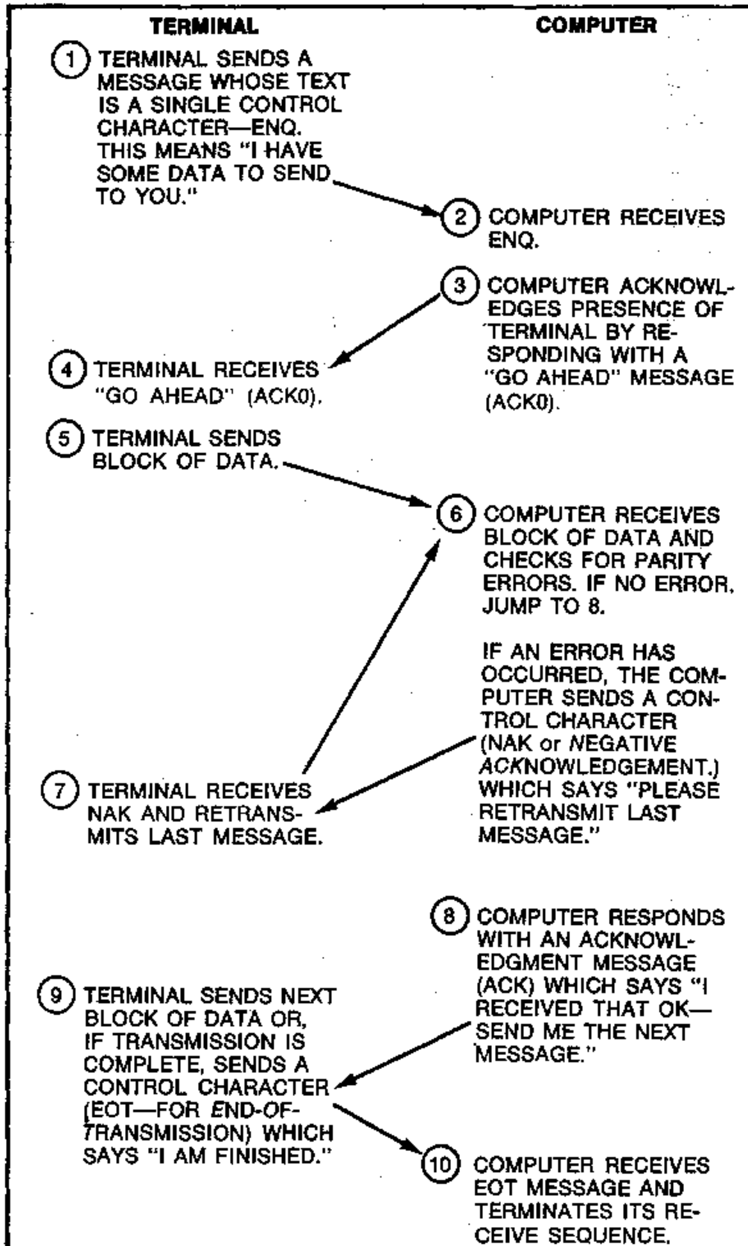

Figure 19—Typical Data Exchange Using BSC

Error Checking and Recovery: To detect and correct transmission errors, BSC uses either VRC/LRC or CRC, depending upon the information code. If the code is ASCII, a VRC check is performed on each character and an LRC on the whole message. The LRC becomes one 8-bit BCC. If the code is EBCDIC or 6-bit Transcode, a CRC is performed. Neither code has a bit available for parity, thus eliminating the possibility of a VRC check. With EBCDIC, CRC-16 is used, resulting in a 16-bit BCC. With 6-bit Transcode, CRC-12 ($x^{12} + x^{11} + x^3 + x + 1$) is used resulting in a 12-bit BCC. If the BCC transmitted does not agree with the BCC computed by the receiver or if there is a VRC error, then a NAK sequence (described in Figure 19) is sent back to the data source. BSC calls for the re-transmission of the block when an error occurs. BSC will typically retry several times before it assumes the line is in an unrecoverable state. BSC checks for sequence errors by alternating positive acknowledgments to successive blocks. ACK0 and ACK1 are the responses to the even-numbered and odd-numbered blocks in the message, respectively. These are sent as separate control messages.

Information Coding: As mentioned above, BSC supports ASCII, EBCDIC, or 6-bit Transcode for coding the information. Certain bit patterns in each set have been set aside for the required BSC control characters (i.e., SOH, STX, ETX, ITB, ETB, EOT, NAK, DLE, and ENQ).

Some BSC controls are two-character sequences (i.e., ACK0, ACK1, WACK, RVI, and TTD). These control character abbreviations are defined in the Glossary, Appendix C.

Information Transparency: In BSC the transparent mode is defined by starting the text field with DLE STX. Once in transparency, the only control character of significance is DLE (Data Link Escape).

In BSC, any data link control characters transmitted during the transparent mode must be preceded by a DLE control character to be recognized as a control function. When a bit pattern equivalent to DLE appears within the transparent data, a DLE is inserted to permit transmission of DLE as data. When received, one DLE is discarded; the other is treated as data. This technique may be referred to as "character stuffing."

Line Utilization: In BSC systems, transmission is half-duplex. The line must be turned around twice between each block (one for the acknowledgment sequence and one for the data block). All the fields are delimited by control characters, and acknowledgments are handled by separate control sequences. There is an acknowledgment sequence required for each block. BSC supports both point-to-point and multipoint lines.

Synchronization: BSC synchronizes on each block or control sequence by preceding the formatted block with a minimum of two synchronizing (SYN) characters. SYN is defined as a unique bit pattern in each of the three information exchange codes available with BSC.

Communications Facility Transparency: BSC was designed for serial synchronous lines. Because it is a character-oriented protocol, it could be implemented on asynchronous and parallel channels. IBM does not offer BSC for these facilities.

Bootstrapping: BSC does not include bootstrapping as part of its line control procedure.

DIGITAL DATA COMMUNICATIONS MESSAGE PROTOCOL—DDCMP

DIGITAL's link protocol, DDCMP (Digital Data Communications Message Protocol), is designed to operate over clocked (synchronized) full- or half-duplex channels, switched or direct links, point-to-point or multipoint networks, and serial or parallel transmission facilities. Further, it will accommodate both synchronous and serial start-stop (asynchronous) modes. DDCMP is capable of controlling message transfers over standard existing hardware and can be implemented on many operating systems.

Controlling Data Transfers: The format of a DDCMP transmission block is shown in Figure 20. The only control character used in DDCMP is the first character in a message. It is used to distinguish between data, control, and bootstrap messages. SOH, ENQ and DLE are used, respectively.

The header is required. It contains the count of 8-bit quantities (bytes) in the information field, some control flags, a response field for positive acknowledgment of received messages, a message sequence number (modulo 256), and an address. The latter field is mainly used for addressing tributary stations in multipoint configurations. The header is verified by having its own CRC. The information field is of variable length (up to 16,383 bytes) and it is followed by a BCC containing a CRC-16 calculated remainder.

Like BSC, DDCMP employs a rigorous set of rules for establishing, maintaining, and terminating a communications sequence. However, because it provides for simultaneous two-way transmission, its procedure is too extensive to cover in this book. The procedure outlined in Figure 21 only touches upon its simplest but least efficient method and from a unidirectional point of view. The same procedure can be occurring in the opposite direction in full-duplex operation.

SYN	SYN	SOH	COUNT 14 BITS	FLAG 2 BITS	RESPONSE 8 BITS	SEQUENCE 8 BITS	ADDRESS 8 BITS	CRC 1 16 BITS	INFORMATION ANY NUMBER OF 8-BIT CHARACTERS	CRC 2 16 BITS

Figure 20—DDCMP Message Format

Error Checking and Recovery: DDCMP uses CRC-16 for detecting transmission errors. When an error occurs, DDCMP sends a separate negative acknowledgment (NAK) message. DDCMP does not require an acknowledgment message for all messages. The number in the response field of a normal header or in either the special NAK or positive acknowledgment (ACK) message specifies the sequence number of the last good message received. For example, if messages 4, 5, and 6 have been received since the last time an acknowledgment was sent and message 6 is bad, the NAK message specifies number 5 which says "message 4 and 5 are good and 6 is bad." When DDCMP operates in the full-duplex mode, the line does not have to be turned around. The NAK is simply added to the sequence of messages for the transmitter.

When a sequence error occurs in DDCMP, the receiving station does not respond to the message. The transmitting station detects from the response field of the messages it receives (or via timeout) that the receiving station is still looking for a certain message and sends it again. For example, if the next message the receiver expects to see is 5 and it receives 6, it will not change the response field of its

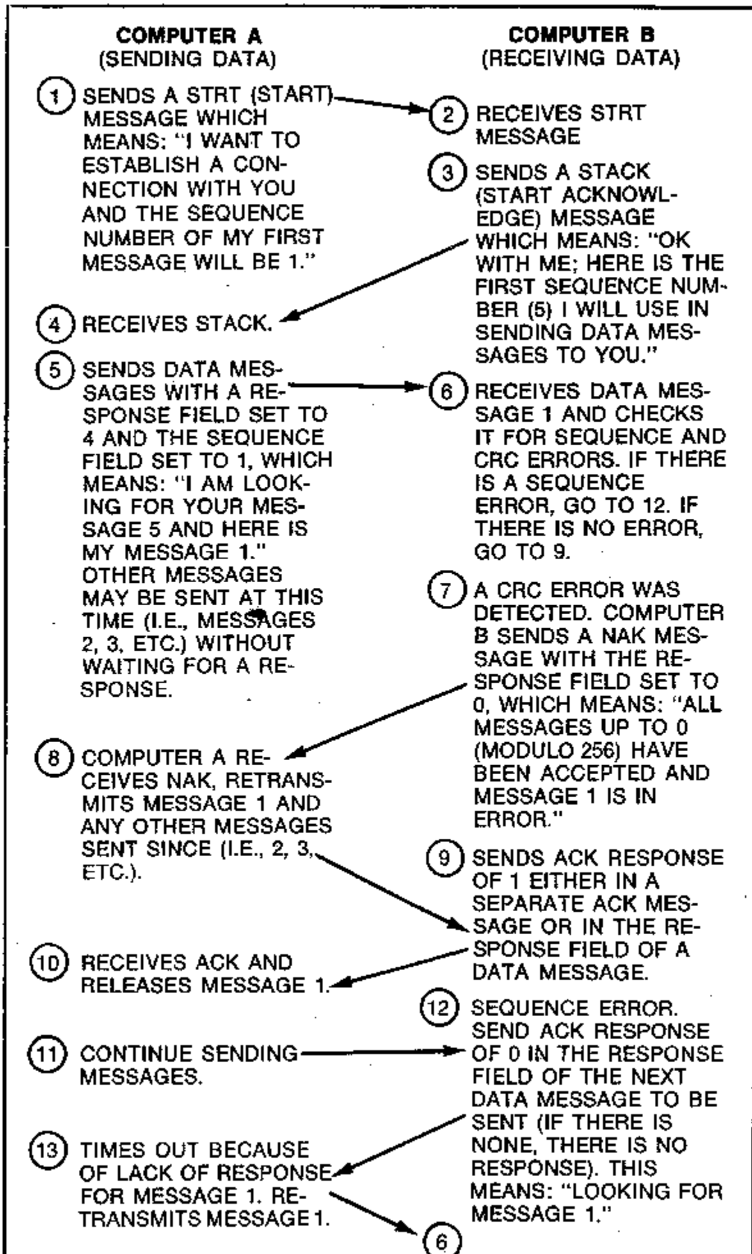

Figure 21—DDCMP Simple Handshaking Procedure (One Direction)

data messages which contains a 4. This says: "I accept all messages up through message 4 and I'm still looking for message 5."

Information Coding: DDCMP uses the ASCII control characters SOH (Start of Header) ENQ (Enquiry), and DLE (Data Link Escape) to distinguish between types of messages. The remainder of the message (including the header) is transparent. This means the data in the information field is encoded by some higher-level protocol according to its design standard for information interchange.

Information Transparency: DDCMP achieves transparency by use of a count field in the header. The header is of fixed length. The count in the header determines the length of the information field which can be up to 16,383 bytes. To validate the header and count field, it is followed by a 16-bit CRC-16 field; all header characters are included in the CRC calculation. Once validated, the count is used to receive the data and locate the second CRC-16 which is calculated on the data field.

Line Utilization: DDCMP uses either full-duplex or half-duplex circuits to their fullest extent. Both modes of operation are defined in a part of the protocol. In the full-duplex mode, it operates like two dependent one-way channels, each containing its own data stream. The only dependency are the acknowledgments which must be sent in the data stream in the opposite direction. To reduce the response overhead, separate ACK messages are unnecessary. They are simply placed in the response field of the next message for the opposite direction. If several messages are received correctly before its transmitter is able to send a message (because the previous message was a long one), all of them can be acknowledged by one response. Only when a transmission error occurs or if traffic in the opposite direction is light (no data message to send) is it necessary to send a special NAK or ACK message, respectively.

In summary, DDCMP line utilization features include:

1. The ability to run on full or half-duplex transmission facilities.
2. The ability to run on many existing hardware interfaces.

3. The ability to support point-to-point and multipoint lines.
4. Transmission mode independence (synchronous, asynchronous, or parallel).
5. No separate ACK's when traffic is heavy.
6. Multiple acknowledgments per ACK (up to 255 messages in one acknowledgment).

Synchronization: DDCMP achieves synchronization through the use of two ASCII SYN characters preceding the SOH, ENQ, or DLE. It is not necessary to synchronize between messages as long as no gaps exist. This feature helps reduce the control overhead because each SYN takes 8-bit times. Character synchronization is unnecessary when DDCMP is used for serial asynchronous and parallel channels.

Communications Facility Transparency: DDCMP can be used for serial synchronous, serial asynchronous, and parallel facilities. This is total facility transparency.

Bootstrapping: DDCMP has a bootstrap message as part of the protocol. It begins with the ASCII control character DLE. The information field contains the load programs and is totally transparent.

SYNCHRONOUS DATA LINK CONTROL—SDLC

IBM's SDLC (Synchronous Data Link Control) was announced in 1973 and operational on products in 1974. Unlike BSC and DDCMP, it is bit-oriented rather than character-oriented. It is designed for full and half-duplex operation.

Controlling Data Transfer: The format of SDLC is shown in Figure 22. The only control character used in SDLC is the flag character which has the bit pattern 01111110. There is a fixed length 24-bit header, a variable length information field, and a fixed-length 24-bit trailer. In addition, the transparency technique (explained below) can increase the size of any of these fields except for the two 8-bit flags that frame the message.

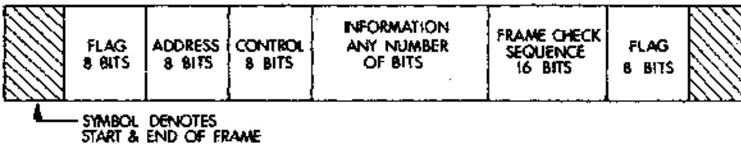

Figure 22—SDLC Message Format

Like DDCMP, SDLC may operate on full-duplex facilities and uses an efficient procedure for data exchange. The example in Figure 23 presents only a one-directional transfer of data in a full-duplex operation. Equivalent operation in the other direction also occurs because it is symmetrical.

Error Checking and Recovery: SDLC uses CRC-CCITT to detect transmission errors. It handles CRC with an inversion technique that differs from methods used by character-oriented protocols. This technique improves the range of detected errors in the domain of possibilities. Like DDCMP, SDLC has a response field (3 bits in the control field) and separate ACK and NAK messages. Unlike DDCMP, SDLC does not NAK transmission errors. The way it handles them is best explained by example: Station B receives messages 2, 3 and 4 from Station A and message 4 is bad. Station B in its next data message for Station A, responds with message 4 which says: "I have received and accepted messages 2 and 3 and I am still looking for message 4. Station A, knowing it has sent message 4, must wait for a period of time to see if it will be acknowledged (timeout) before it sends the message again.

Unlike DDCMP, SDLC responds to sequence errors with a NAK message (a less frequent occurrence than transmission errors). The response is similar to DDCMP's handling of transmission errors in that the NAK acknowledges the messages received in correct sequence condition, i.e., "received 4, 5 correctly and looking for 6" when it receives 7 instead.

Information Coding: SDLC is not concerned with information exchange codes. The only control character is the flag and it has a fixed pattern (01111110) regardless of the exchange code used by the system. Like DDCMP, a higher-level protocol defines the information exchange code used to transfer meaningful data.

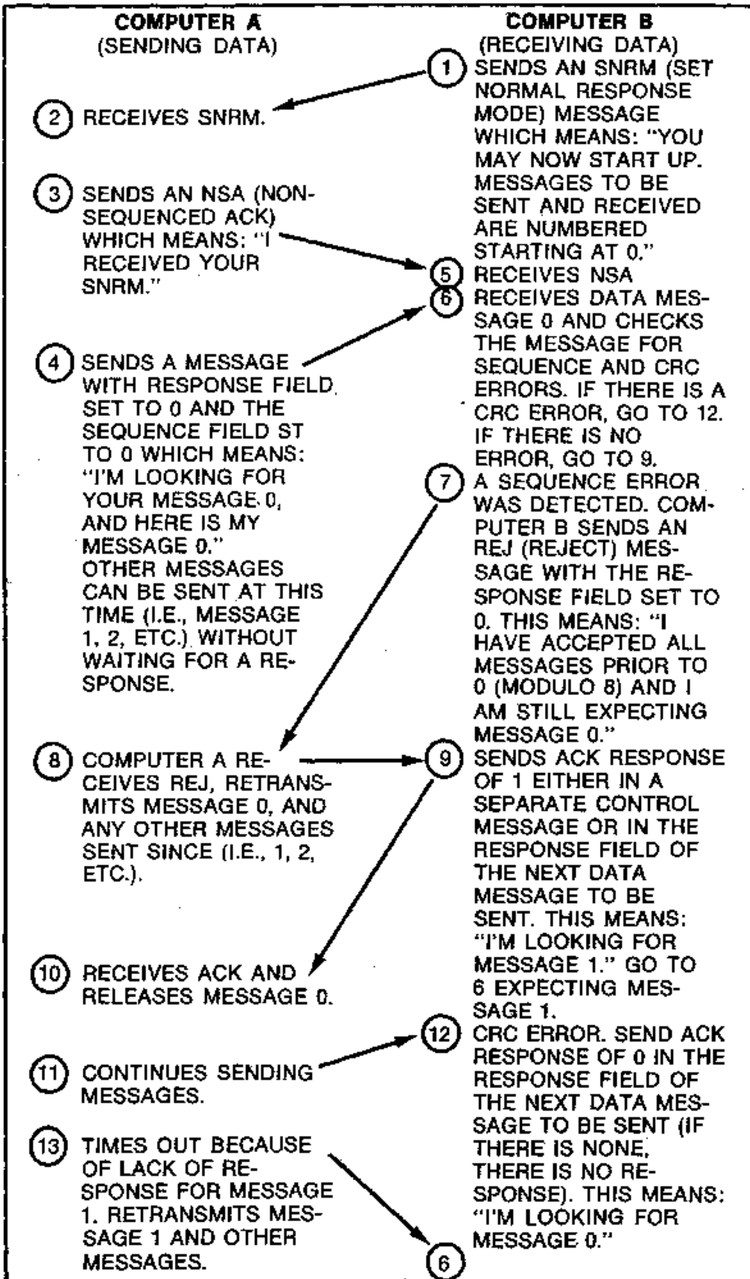

Figure 23—SDLC Simple Handshaking Procedure (One Direction)

Information Transparency: SDLC is a bit-oriented protocol and only has to be sure that a flag character bit pattern anywhere between frames does not arrive at the receiver.

The technique used by SDLC for achieving transparency may be referred to as "bit stuffing." SDLC uses a character at each end of the message called a "flag." The flag's bit sequence is 01111110. The only control characters are the flag characters. To ensure that a flag character does not appear in the data portion of the message, a zero (0) bit is inserted whenever five one (1) bits appear in a row. Thus, a 01111110 bit pattern meant to be data would appear in the data stream as 011111010, or 9 bits. The receiver also counts bits, and if it detects five one's in a row followed by a zero bit, it removes the zero bit. If it was a one bit, it is a legitimate flag and the end of message has been received.

Line Utilization: SDLC uses transmission facilities with similar efficiency to DDCMP because of the following features:

1. The ability to transmit on full or half-duplex facilities.
2. Low control character overhead (flag , header, and check bits total 6 characters.) BSC uses 8 characters for transparency (without header) and DDCMP uses 10 characters.
3. No "character stuffing"; SDLC uses "bit stuffing"; In contrast, DDCMP's transparency overhead is fixed because it uses a count method.
4. No separate ACK messages are necessary.
5. Multiple acknowledgments per ACK (up to 7 messages).
6. The ability to support point-to-point and multipoint lines.

Synchronization: SDLC synchronizes on the flag characters between messages. No SYN characters are required.

Communications Facility Transparency: Because of the "bit stuffing" requirement of SDLC, it cannot be used for serial asynchronous or parallel facilities. Asynchronous characters are fixed length and "bit stuffing" would destroy them. On parallel connections, a separate wire(s) would be needed for the stuff bit(s). SDLC is designed for efficient use of high speed serial synchronous full-duplex facilities.

Bootstrapping: SDLC does not provide for bootstrapping as part of its protocol.

OTHER PROTOCOLS

A number of national and international standards organizations have been cooperating over the past years to produce standard data-communication protocols. These organizations include the American National Standards Institute (ANSI) and the International Standards Organization (ISO).

The cooperative efforts of these organizations have resulted in the development of a variety of standard data communications control procedures and codes. Typical of these developments are the Advanced Data Communication Control Procedures (ADCCP) developed by ANSI, and the High-Level Data Link Controls (HDLC) developed by ISO. These protocols differ slightly from SDLC.

The features and characteristics inherent in the five protocols mentioned thus far (i.e., DIGITAL's DDCMP, IBM's BSC and SDLC, ANSI's ADCCP, and ISO's HDLC) are summarized in Figure 24.

DATA TRANSPARENCY

The way in which each of the five protocols achieves data transparency is illustrated in Figure 25. This illustration summarizes the transparency techniques of "character stuffing," "bit stuffing" and "count."

FEATURE	DDCMP	BISYNC	SDLC	ADCCP	HDLC
Full Duplex	Yes	No	Yes	Yes	Yes
Half Duplex	Yes	Yes	Yes	Yes	Yes
Serial	Yes	Yes	Yes	Yes	Yes
Parallel	Yes	No	No	No	No
Data Transparency	Count	Character Stuffing	Bit Stuffing	Bit Stuffing	Bit Stuffing
Asynchronous Operation	Yes	No	No	No	No
Synchronous Operation	Yes	Yes	Yes	Yes	Yes
Point-to-Point	Yes	Yes	Yes	Yes	Yes
Multipoint	Yes	Yes	Yes	Yes	Yes
Error Detection (CRC)	CRC-16	CRC-16	CRC-CCITT	CRC-CCITT	CRC-CCITT
Retransmit Error Recovery	Yes	Yes	Yes	Yes	Yes
Bootstrapping Capability	Yes	No	No	No	No

Figure 24—Protocol Features and Characteristics

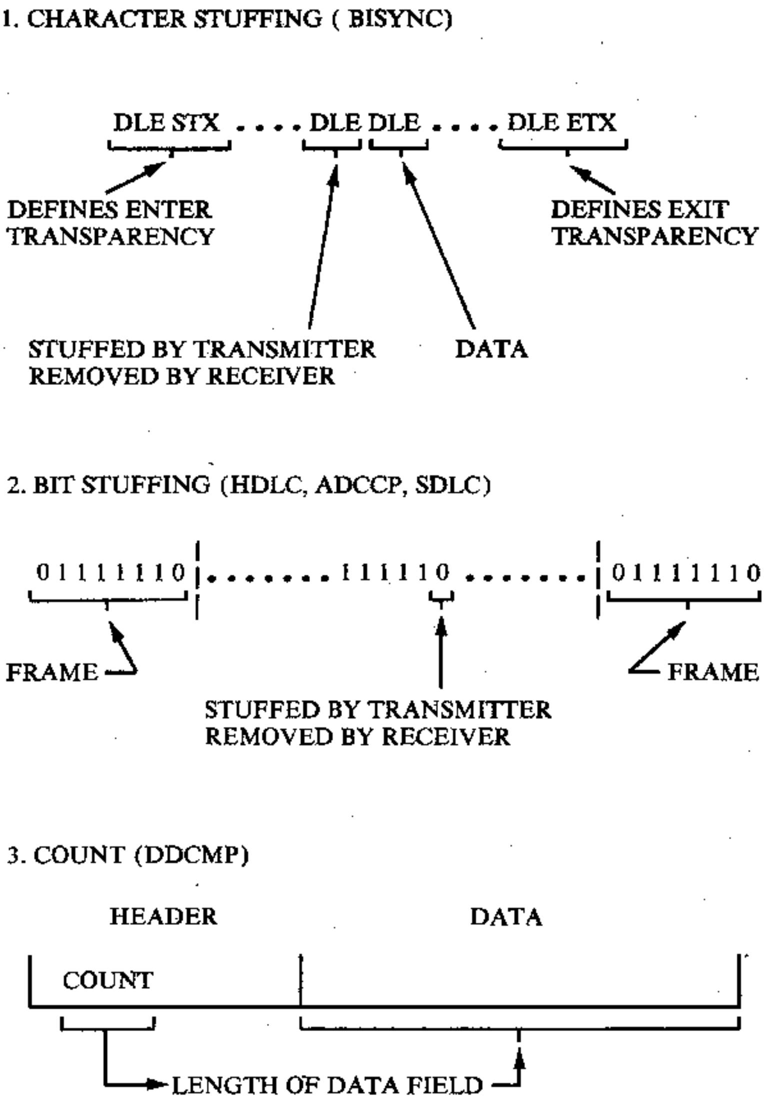

Figure 25—Techniques for Achieving Transparency

chapter 6
using pdp-11 minicomputer networks

INTRODUCTION

A growing number of commercial and industrial organizations, government agencies, and educational institutions are increasing their operational efficiency, reducing costs, and improving their organizational structure by redistributing their large centralized computing systems into smaller functional systems to serve specific regions.

As mentioned in Chapter 1, one of the primary reasons for the growing trend in computer networking is the economy that can be achieved through resource sharing. Resources, such as devices, files, programs, and processing power, at each location in the network are available for use at other locations in the network.

Typical examples of how DIGITAL's PDP-11 family of computers can and are being used in such networks are described in the remainder of this chapter.

THE INDUSTRIAL ENVIRONMENT

Computer networking has had a major impact on the structure of many industrial manufacturing organizations. A new concept of the automated factory called the "hierarchical" approach has come into vogue. With this approach, the minicomputers, programmed for simple, single tasks, direct the machines on the factory floor. And as they do, they feed information on what they are doing to successively higher levels of computers. These higher-level computers compile and analyze the data and provide management with the information it needs to run the business. Management is able to know, within short notice, how work is progressing at the various locations, where shop orders are, what machines are working, and the status of parts and materials—and be able to do something about them if necessary.

The discrete-manufacturing and continuous-process industries have applied this approach to all areas of industrial automation. By connecting multiple minicomputers and terminals in a hierarchical network, entire manufacturing processes have been coordinated, resulting in increased production capacity, reduced product cost, and improved product quality.

The hierarchical network depicted in Figure 26 employs three levels of computing and control responsibility. Level III in this network consists of three PDP-11/10 core-only minicomputers operating in hostile environments. These computers are used in real-time control of various processes and other on-line functions in the factory. An assortment of digital and analog acquisition devices are connected to the control computers. Their function is to perform limit checking, control algorithms, and directly control or alter the processs through the use of digital and analog outputs. Typical applications include control of automatic warehousing equipment, closed-loop process control, scan-log-alarm, and control of automated testing equipment.

Next in the hierarchy are the supervisory computers at Level II. These PDP-11/40 systems monitor and control the functions being performed by the Level III machines and load them with new programs as the need arises. The Level II supervisory mini's periodically interrogate the Level III machines for pertinent data about the process under control.

In addition to collecting and processing pertinent data received from the Level III control computers, the supervisory mini's also acquire data from various manual entry stations and terminals located throughout the factory. These manual entry stations are used to record material movement, machine utilization, production quantities, employee attendance, inspection results, etc. The terminals are used for real-time inquiry from various managerial levels about order status, equipment utilization, machine up-time and down-time, production capacity, product yield, raw material consumption, etc.

The large-scale PDP-11/45 minicomputer at Level I supervises the Level II machines and performs software development for Levels II and III. It also performs processing and file updating on data received from the Level II computers.

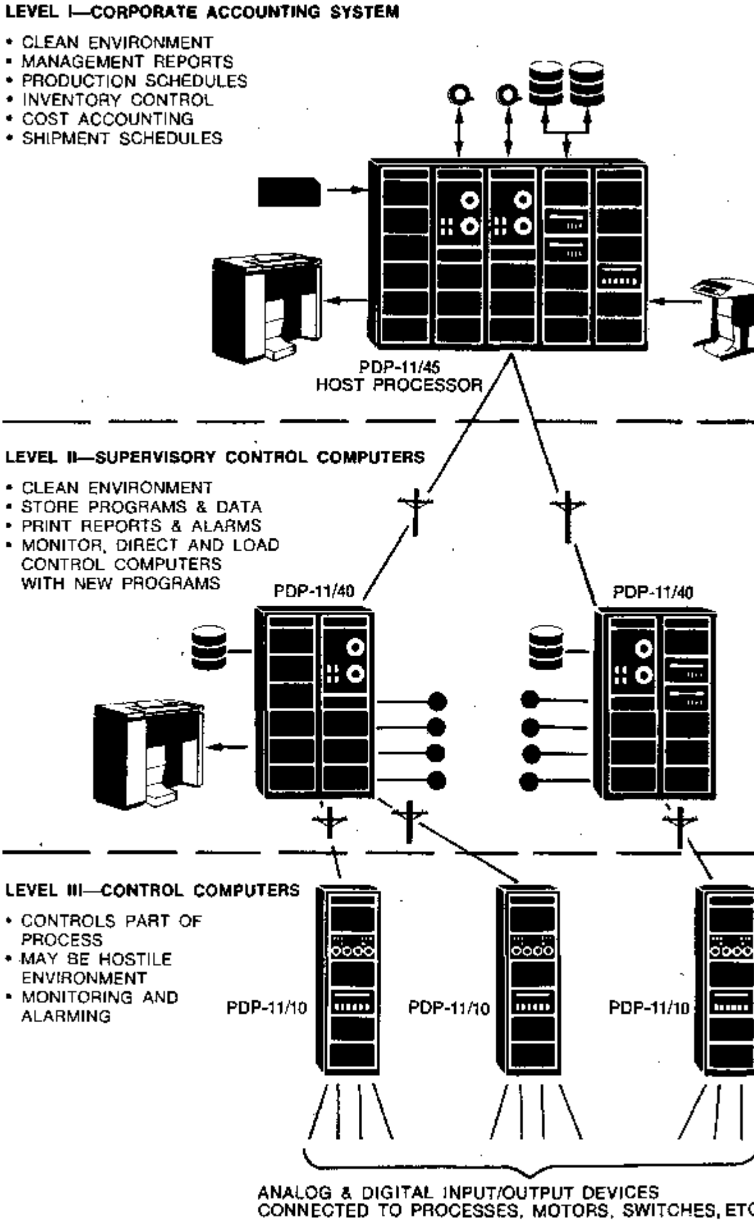

Figure 26—A Hierarchical Network in an Industrial Environment

The data base files reflecting the over-all operation of the factory are distributed between the Level I host processing facility and the Level II supervisory systems.

The most outstanding attribute of a hierarchical network of this type is the distribution of processing power at the right levels or locations within the plant; that is, the levels where the action takes place. By distributing processing power throughout the plant, the load on the host computer is reduced, response time at the remote locations is improved, and over-all reliability of the entire industrial operation is improved.

THE LABORATORY ENVIRONMENT

The number, kind and scope of experiments that are performed in many university laboratories have increased markedly over the last few years, while at the same time, the sophistication and complexity of many of the experiments has made the job of conducting them more and more difficult.

To alleviate the work load that has been thrust upon the experimenter, university laboratory complexes have implemented minicomputer-based resource-sharing networks of the type illustrated in Figure. 27.

In this example, each of the outlying university laboratories has a small PDP-11/10 minicomputer with 8K words of memory and a complement of standard laboratory peripherals attached to it. While each of these systems is sufficient for conducting experiments and collecting data, additional processing resources are required to analyze the data, generate reports, change existing programs, and create new programs.

Rather than have each of the outlying laboratories upgrade its existing system with more core memory, disk storage devices, line printers, etc., a centrally located PDP-11/45 minicomputer with 80K words of memory was installed. The PDP-11/45, which is front-ended by a PDP-11/40 system, provides all of the necessary processing resources required at the outlying laboratories. The front end buffers the incoming and outgoing data and acts as a spooler for the card reader and line printer.

A synchronous line interface connected to each of the PDP-11/10 systems enables the experimenters to transmit data over private lines to the front-end processor at speeds up to 4800 Baud. A higher speed interface connects the front-end processor to the host processor.

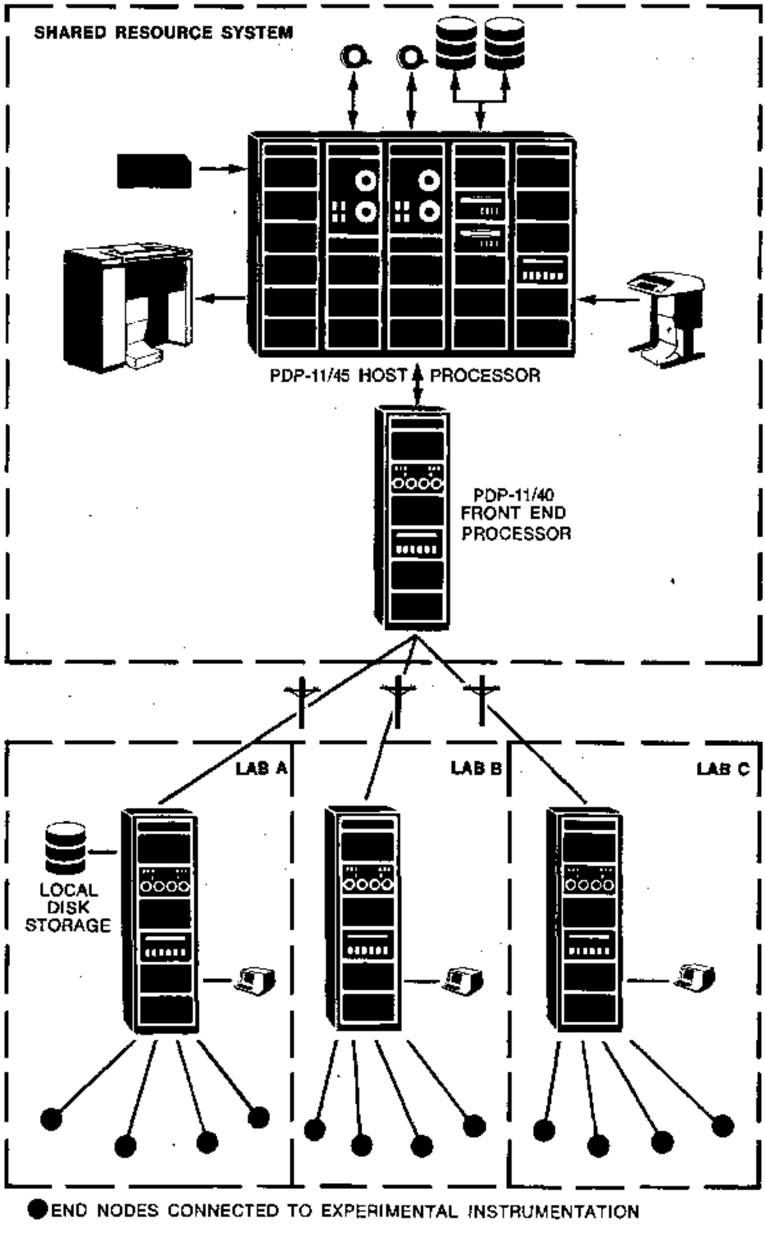

Figure 27—A University Laboratory Network

Peripheral devices connected to the host computer include magnetic tape drives and removable disk cartridge drives. The PDP-11/40 front-end accommodates four synchronous lines and has a 300 LPM Printer and a 300 CPM card reader attached to it.

From the console terminal, the experimenter can create programs on a local PDP-11/10 computer using the resources of the PDP-11/45 host computer. Programs stored in the host computer's disk file can be retrieved and loaded into the experimenter's local minicomputer and executed to conduct the experiment. While this is going on, the experimenter can be using the terminal to create new programs or to run other programs in the host computer to analyze data from a previous experiment. The foreground program conducting the experiment and gathering data can be sending data to a program in the central host system that analyzes the data almost as fast as it is created. The experimenter can have results printed on the central line printer or on a local terminal. Programs created by one experimenter can be requested by another experimenter for execution on his local PDP-11/10 system.

In summary, this type of network enables each of the outlying laboratories to share the following resources:

- Programs
- Program data created in one system and used by programs in another system.
- Peripheral devices at the host processing facility; e.g., line printer and magnetic tape drives.
- Data and programs stored in the host processor's files.
- Computing power; e.g., the host's processing power and main memory.

Other features inherent in this type of system include remote task loading, remote initiation of programs, and remote reporting of results.

A RESERVATION PROCESSING SYSTEM

A transportation system, serving a highly populated area, could install a PDP-11 minicomputer network for handling all passenger reservations. This type of network permits efficient allocation of seats, collects accounting information, prints tickets, and provides timely status information such as reservation lists and available seats.

The reservation system's data base, shown in Figure 28, is maintained at a central site consisting of two PDP-11/40 minicomputers, each with 64K words of core memory. Reservation files are stored on the central site disks, which have a storage capacity of approximately 230-million characters.

In the example shown in Figure 28, each of the PDP-11/40 central site processors is front-ended by another PDP-11/40 system with 24K words of memory. The front-ends handle communications-oriented tasks, such as terminal polling, error detection and correction, and the insertion and removal of control characters.

Inquiries from terminals are routed to the central data base via two additional PDP-11/40 systems that function as data concentrators. Each concentrator has approximately 16K words of memory. Messages are concentrated onto a single 2400 Baud leased line and transmitted to the front-ends. Other terminal inquiries are routed directly to the front-ends.

The front-ends are capable of temporarily storing messages prior to transferring them to the central system. The front-end also determines whether a message must be routed to another terminal or whether it is destined for the central computer.

Communications in the system can be between terminals, between terminals and the central system, between front-ends, and between concentrators. The network is designed so that the system will not be affected if a component fails. For instance, if a central site computer fails, the standby processor automatically takes over. The front-end processor is also backed up by a standby system.

Passenger reservations may be entered at asynchronous or synchronous buffered terminals. Prior to transmission, the reservation message is verified and edited by the ticket agent. It is then transmitted to one of the regional concentrators or directly to the front-end over either a 2400 Baud synchronous line or over a 75, 110, or 300 Baud asynchronous line. To assure message accuracy, Cyclic Redundancy Checks are performed on all synchronous lines.

This type of communications network provides reservation agents with extensive remote inquiry capabilities. For example, they can determine the availability of different types of seats such as window,

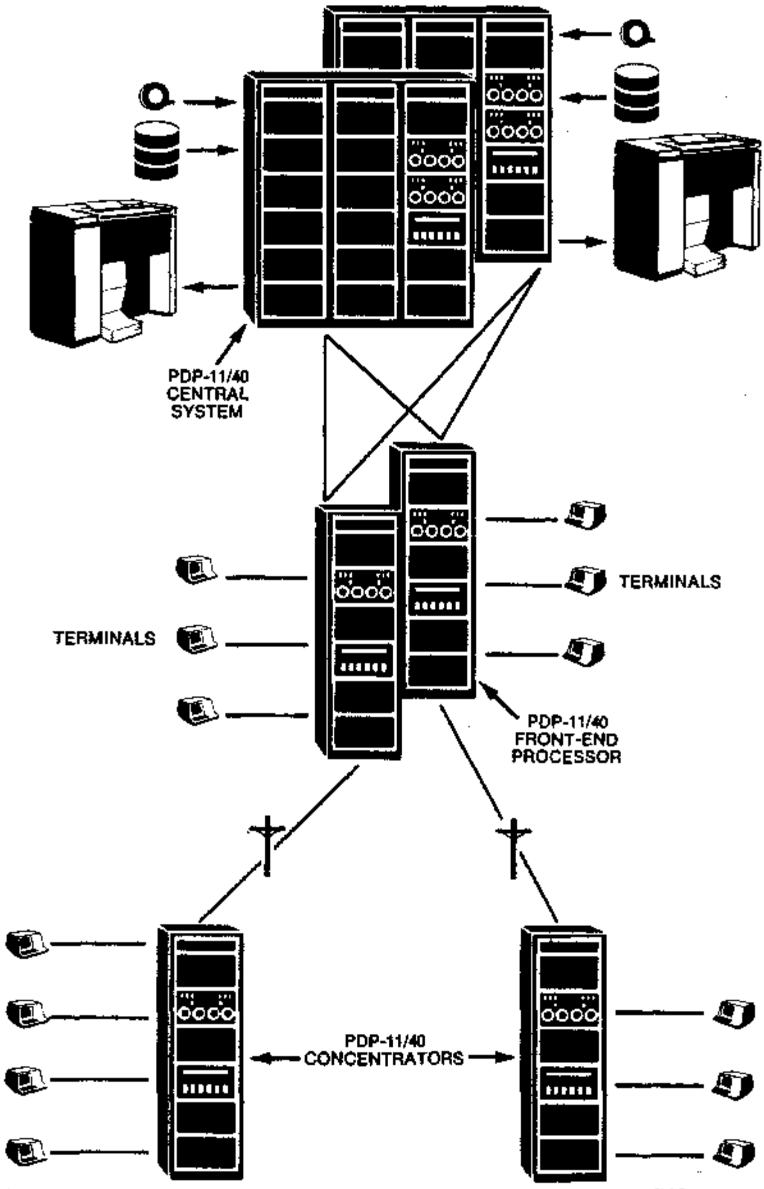

Figure 28—A Passenger Reservation Network

aisle, or smoking secton.

Traffic is controlled by a supervisor who can determine whether extra seats are needed. He may intervene at any time to make inquiries or changes.

A CREDIT CARD VERIFICATION SYSTEM

The use of credit cards to simplify consumer business transactions has been a part of the U. S. economy for many years. The validation, control and billing accuracy for the millions of cards in circulation, however, is anything but simple.

To facilitate the tremendous number of transactions that occur daily, any credit card service, with several hundred participating banks, thousands of merchant outlets, and millions of credit card holders, could use a minicomputer network to speed up its verification procedures.

When a credit card purchase is made, the merchant telephones the nearest message verification center where an operator enters the credit card number at a keyboard/display terminal. The message is sent via regional concentrators to a message switching center and is then routed to the credit card owner's local credit authorization center. The average time to process and respond to an inquiry is less than two minutes.

An example of such a system is illustrated in Figure 29. It consists of two centrally located PDP-11/45 systems. These two systems are interconnected so that they can back each other up to assure continuity of service. With this approach, a terminal will always receive a response, even if part of the system is shut down. The central computing facility can supplement credit authorization centers after hours by maintaining files about lost or stolen cards or fraudulent accounts.

Several PDP-11/10 systems functioning as data concentrators, may be strategically located to service specific regions. They can serve as extensions to the central facility for routing and traffic-control purposes. Each concentrator consists of two PDP-11/10 systems. One is a redundant system that can take over if the other system is not in service. The concentrators reduce line costs by concentrating data from different sources onto one high-speed line. The concentrators may be linked to the central system via four-wire, full-duplex leased telephone lines.

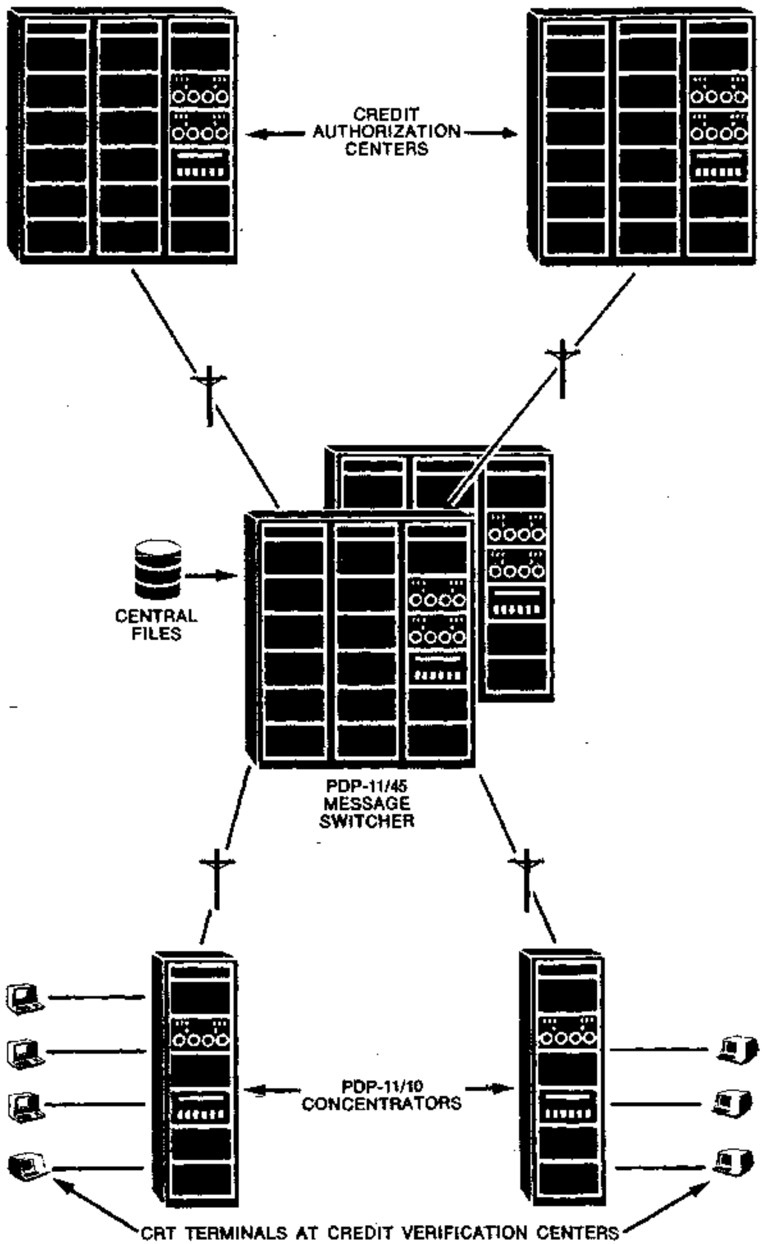

Figure 29—A Credit Card Verification Network

appendix a

data communications with the pdp-11

THE PDP-11 FAMILY

Architecture makes the PDP-11 the ideal computer for data communications networks. This architecture evolves around the UNIBUS—a high-speed, bi-directional, asynchronous bus to which all major system elements (processor, memory, mass storage, terminals, lines, etc.) interface. Since all system components communicate in similar fashion over the UNIBUS, the processor has the same easy access to peripherals as it has to memory, and devices can send, receive, and exchange data without processor intervention or intermediate buffering in memory. The UNIBUS design also allows easy and inexpensive use of Direct Memory Access (DMA) devices. And multiple single-line communications interfaces can be added to the PDP-11 without special multiplexing hardware.

Byte handling, the crux of data communications, is performed easily and efficiently; each eight-bit byte is directly addressable. Dynamic stacking along with subroutine call and interrupt processing permit reentrant coding and fully nested interrupts. This lets multiple devices share service routines, and nesting allows higher-priority routines to interrupt those having lower priority.

The extremely low time overhead of the PDP-11's communications I/O allows the CPU to be time-shared for other tasks while performing data communications functions. This ability, combined with the concept of a family of compatible and configurable machines, makes a PDP-11 the ideal choice for large computing networks requiring a hierarchy of subsystems. For example, PDP-

11's could be used in dedicated or semi-dedicated modes and communicate with a central, expanded PDP-11 which could also control a number of peripherals.

Multiprocessor configurations are also possible with the PDP-11. Bus switches are available to allow peripheral and CPU sharing between systems, and bus-to-bus interconnections permit data communication between separate operating systems.

Large multiprocessor PDP-11 based systems offer the user maximum utilization of shared peripherals, increased system throughput and computational power, and a redundant, failsoft configuration.

The PDP-11 family of computers ranges from the PDP-11/10, an extremely compact and low-priced minicomputer, to the PDP-11/50, an extremely fast and powerful MOS-memory medium-size computer. The architecture, instruction set, and operational characteristics of all PDP-11's are similar. This means that software is upward compatible, and PDP-11 peripherals can be used with all PDP-11 family processors. Thus smaller PDP-11's can be upgraded to larger PDP-11 systems without extensive replacement of hardware and software.

Features common to all PDP-11's include:

- 16-bit word length (two 8-bit bytes).
- Word or byte processing and addressing (very efficient handling of 8-bit characters).
- Modular design (provides extreme ease and flexilibity in configuring systems).
- UNIBUS architecture (single bus for all system components—simplifies interfacing and system expansion).
- Asynchronous operation (systems run at their highest possible speeds).
- Multiple general-purpose registers (simplifies programming).
- Stack processing (hardware sequential memory manipulation—eases handling of structured data, subroutines, and interrupts).
- Automatic priority processing (4-line, multi-level system is dynamically alterable).
- Vectored interrupts (fast interrupt response without device polling).

- Direct Memory Access
- Power fail and restart

An extensive range of peripherals and options is available for use with PDP-11 computers, including DIGITAL-built terminals, fixed and moving head disks, several types of magnetic tape storage media, and several types of memories. Specialized data communications hardware developed for the PDP-11 includes synchronous and asynchronous interfaces, error checking hardware, and various adaptors.

PDP-11 OPTIONS FOR DISTRIBUTED NETWORKS

A variety of standard PDP-11 options may be used to implement the communications channels of a computer network. The type of interprocessor link selected depends on the required data rate and the physical distance between the computers.

ASYNCHRONOUS SINGLE LINE INTERFACE (DL11)

The DL11 connects PDP-11 systems to a variety of serial asynchronous channels. Features include character-buffered receiver and transmitter, selectable data rates (between 50 and 9600 Baud), independent receive and transmit speeds, strap-selectable character size and stop code length, 20 mA or EIA output levels, and optional full data set control.

PROGRAMMABLE ASYNCHRONOUS DUAL LINE INTERFACE (DC11)

The DC11 interfaces local or remote terminals (via modems) to PDP-11 systems. Full or half-duplex operation at four programmable speeds. Split speed operation. Programmable character size: 5, 6, 7, or 8 bits. Automatic parity checking. Auto-answering capability. Interfaces to Bell 103, 202, or equivalent data sets. Reverse channel available for BELL 202 operation.

PROGRAMMABLE 16-LINE ASYNCHRONOUS MULTIPLEXER (DH11)

The DH11 multiplexer connects the PDP-11 with 16 asynchronous serial communications lines operating with individually programmable parameters. There are 14 standard speeds (0, 50, 75, 110, 134.5, 150, 200, 300, 600, 1,200, 1,800, 2,400, 4,800, or 9,600 Baud) plus two external inputs; four character sizes (5, 6, 7, or 8 data bits); three parity configurations generated and checked by the hardware (odd, even, or none); and three stop code lengths (1 and 1½ stop bits for 5 data bits, 1 and 2 stop bits for 6, 7, or 8 data bits).

16-LINE ASYNCHRONOUS MULTIPLEXER (DJ11)

The DJ11 is a multiplexed interface between 16 asynchronous serial data communications channels and the PDP-11 UNIBUS. It is a low cost, high performance unit whose character formats and operating speeds are jumper or switch-selectable in four-line groups. The customer may select from 12 standard speeds (50, 75, 110, 134.5, 150, 300, 600, 1,200, 1,800, 2,400, 4,800 or 9,600 Baud); four character sizes (5, 6, 7 or 8 data bits); three parity configurations generated and checked by the hardware (odd, even, or none); and three stop code lengths (1 and 1½ stop bits for 5 data bits, and 1 and 2 stop bits for 6, 7, or 8 data bits).

SYNCHRONOUS LINE INTERFACE (DU11)

The DU11 is a single-line, program-controlled, double buffered communications device designed to interface the PDP-11 processor to a serial synchronous line. The self-contained unit is fully programmable with respect to Sync character, character length (5 to 8 bits), and parity selection.

NPR SYNCHRONOUS LINE INTERFACE (DQ11)

The DQ11 is a high-speed, double-buffered, Non-Processor Request (NPR) communications device designed to interface the PDP-11 Processor to a serial synchronous communications channel. This interface allows the PDP-11 to be used for remote batch and remote concentrator applications. With the DQ11, the PDP-11 can also be used as a front-end asynchronous line controller to handle remote and local synchronous terminals. The DQ11 sets a new performance standard for the industry, with speeds up to 1.0 Megabaud.

PDP-11—To IBM 360/370 Channel Interface (DX11B)

The DX11B is a programmable interface between a PDP-11 UNIBUS and an IBM 360/370 multiplexer or selector channel. The DX11B hardware handles the detection and response to all channel-generated control signals. Features include Non-Processor Request (NPR) data transfers, data transfer rates up to 250,000 bytes per second (depending upon IBM model), byte multiplexed or burst mode operation, recognition of up to 128 IBM device addresses over the full range of 256 addresses, and built-in maintenance and protection capabilities. The DX11B can be programmed to emulate an IBM 2848, 2703, or 3705 control unit.

Automatic Calling Unit Interface (DN11)

The DN11 provides computer control of Bell 801A, 801C or equivalent Automatic Calling Units.

Signal Conditioning Interfaces (DF11)

The DF11 signal conditioning options permit most DIGITAL serial line interfaces to adapt to any of the common communications levels such as 20 mA Teletype or EIA. The options include both the electrical conversion circuitry and the physical electrical connectors. One DF11 model incorporates an asynchronous integral modem.

Communications Arithmetic Option (KG11-A)

The KG11-A computes Cyclic Redundancy Checks (CRC) and Logitudinal Redundancy Checks (LRC) for detecting errors in serially transmitted data.

The features and characteristics of the PDP-11 asynchronous and synchronous line interfaces are summarized in the following tables.

SUMMARY OF PDP-11 SYNCHRONOUS LINE INTERFACE CHARACTERISTICS

Serial Line Interface	DU11	DQ11
Maximum Number of Units Per System	16	16
Character Size	5, 6, 7 or 8*	Up to 16*
Sync Character Stripping	Yes	Yes
Baud Rate Range	Up to 9600	Up to 1 Megabaud
Sync Character Programming	Yes	Yes
Full- or Half-Duplex	Yes	Yes
Parity Generation (Transmit)	Yes	Yes
Parity Check (Receive)	Yes	Yes
Auto Answer	Yes	Yes
Full Data Set Control	Yes	Yes
Direct Memory Access (DMA)	No	Yes
Character Buffered	Double	Double

* Program Selectable

SUMMARY OF PDP-11 ASYNCHRONOUS LINE INTERFACE CHARACTERISTICS

Serial Line Interface	DL11	DC11	DJ11	DH11
Number of Lines Per Unit	1	2	16	16
Maximum Number of Units Per System	47	16	16	16
Programmable	No	Yes	No	Yes
Character Size 5, 6, 7, or 8	Yes [1]	Yes [2]	Yes [3]	Yes [2]
Stop Bits	1, 1.5 or 2 [1]	1 or 2 [2]	1, 1.5 or 2 [3]	1, 1.5 or 2 [2]
Baud Rate Range	40 to 9600 [1]	50 to 1800 [2]	50 to 9600 [1]	50 to 9600 [2]
Independent Receive and Transmit Speed	Yes	Yes	Yes	Yes
Full-or Half-Duplex	Yes	Yes	Yes	Yes
Break Detection	Yes	No	Yes	Yes
Reverse Break Generation	Yes	Yes	Yes	Yes
Parity (Transmit)	Yes	No	Yes	Yes
Parity Check (Receive)	Yes	Yes	Yes	Yes
Auto-Answer	Yes	Yes	No	Yes
Full Data Set Control	Optional	Yes	No	Optional
Direct Memory Access (DMA)	No	No	No	Transmit
Character Buffered	Double	Single	64 [4]	64 [4]

[1] Strap Selectable
[2] Program Selectable
[3] Switch Selectable
[4] Receive Only

DIGITAL TERMINALS

DIGITAL's product offerings in the terminal area include both typewriter-oriented terminals and CRT display terminals.

Typewriter-Oriented Terminals

The most recent addition to DIGITAL's typewriter-oriented terminal family is the LA36-DECwriter. This unit features 132 print positions, 30 character-per-second transfer rate, a 7-by-7 dot matrix print head for quiet, high-quality printing, multi-part paper handling, and a 96 ASCII character set which allows upper and lower-case characters to be printed.

Other DIGITAL product offerings in the typewriter-terminal class are the LT33 Teletype and the LA30 DECwriter. The LT33 may be used to type in or print out information, or to read in or punch out perforated paper tape. Printing, punching, and reading are performed at rates up to 10 characters per second.

The LA30 prints from a set of 64 characters at speeds up to 30 characters per second. Data entry is made from either a 97- or 128-character keyboard. The LA30 produces an original and one copy on a standard 9⅞-inch wide, tractor-driven continuous form.

CRT Display Terminals

DIGITAL's VT50 Video Terminal is an exceptionally low cost, high-performance terminal featuring a large-sized screen area capable of displaying twelve 80-character lines, and switch-selectable transfer rates up to 9600 Baud. A 20mA current loop interface is standard on the VT50. Depending upon user and system requirements, the host system designer can switch-select transmission rates from 75-to-9600 Baud. Even parity or mark (no parity) can also be selected. To permit telephone line communications, an acoustic coupler can easily be added.

Another CRT display terminal offered by DIGITAL is the VT05. This alphanumeric video display terminal is logically and electrically equivalent to the teletypewriter. It can be connected directly to modems, acoustic couplers, and other EIA/CCITT compatible devices. Or it may be connected directly to a computer via its 20mA current loop interface. A single switch on the rear of the VT05 allows the operator to select transfer rates of 110, 150, 300, 600, 1200, and 2400 Baud. Another switch changes the terminal from full to half-duplex operation.

PDP-11 DATA COMMUNICATIONS SOFTWARE

General

Digital Equipment Corporation has developed data communication software as an integral part of its hardware in order to minimize programming costs and complexity. To satisfy the unique requirements of a diversity of data communication system users, this software was designed with a building-block approach in mind. This approach provides each user with a solid foundation to which he can add his own unique applications software to make his system a functioning entity. Each software module is general purpose and easy to interface with.

PDP-11 data communications software ranges from modular communications executives, such as COMTEX, to sophisticated operating systems, such as RSX-11D and RSX-11M.

COMTEX (Communications-Oriented Multi-Terminal Executive)

Of DIGITAL's general-purpose data communications software, COMTEX is by far the simplest in structure and function. COMTEX is a modular, single-task oriented data communications executive with routines for servicing communication channel interfaces and terminals. It provides the framework and resources that the user can tailor to meet his specific data communications needs. The COMTEX system base contains routines for supporting DIGITAL's communications interfaces and such peripherals as the console terminal, card readers and line printers. Additional COMTEX routines may be added to the basic package depending on the user's needs.

COMTEX consists of:

- The System Control and Interface Program (SCIP)—an executive which manages the execution of various terminal and line-control related tasks.
- Interrupt Service Routines (ISR)—handlers which service the various types of communication line interfaces.
- Terminal Application Packages (TAP)—control modules which service various types of terminals and emulate specific functions.

COMTEX is not a turnkey software package, but a base which frees the user from having to develop his own communications-oriented software package. The COMTEX system is delivered to the customer in source form so that he can readily configure it to his particular application with minimal effort and, therefore, have a running system in the shortest possible period of time. By this time, too, he is sufficiently familiar with the inner workings of COMTEX to delete the general purpose service which he does not require and thus improve system throughput.

The user program receives data from terminals via communications lines that are controlled by COMTEX. The program manipulates the data, then transmits responses via COMTEX. Communication from the user program to COMTEX is by a set of easily understood commands.

The modular nature of COMTEX allows the user to easily replace and/or modify the terminal-dependent code. The protocol-dependent control modules (TAPs) are completely independent of a particular user application. TAPs perform such functions as special character detection, terminal control and code conversion. TAPs are table-oriented so that a single TAP can service multiple terminals of the same type.

The I/O device handlers performing communication line control functions (ISRs) are transparent to all functions not related to line control. They handle such functions as modem control and the mechanics of data input and transmission. One copy of an ISR can service multiple communication controllers of the same type.

All COMTEX internal operations are scheduled on a priority basis. Time-critical functions can be performed at high-priority levels; such tasks as code conversion can be performed at lower-priority levels. The priority scheduler is accessible to the user for priority scheduling of his tasks. The task scheduler can be driven by the system timer routines, allowing the user to schedule tasks on a time basis.

COMTEX operates on a PDP-11 with a minimum of 8K words of core memory. The following building blocks are available from DIGITAL and are supported by the COMTEX operating system.

I/O HANDLERS

DL11	—Single asynchronous line interface
DC11	—Programmable dual asynchronous line interface
DJ11	—16-line asynchronous multiplexer
DH11/DM11-BB	—Programmable 16-line asynchronous multiplexer with modem control
DM11	—Single-speed 16-line asynchronous multiplexer
DU11	—Synchronous line interface
DP11	—Synchronous line interface
KG11-A	—Communications Arithmetic Unit (only in 2780 package)
DX11-B	—360/370 Channel interface (as part of Front-End System Base package)

Control Modules (TAPs)

Interactive Teletype
IBM 2741

DN11 Autodial Interface
2848 Emulator (as part of the Front-End System Base package)
Bisync Communications (only as part of the 2780 and HASP packages)
Console Terminal

I/O Device Handler/Control Module
CR11 Card reader
LP11 Line printers

Applications Program
2780 Emulator
HASP Workstation Emulator (binary package only)

DOS/COMTEX

DOS/COMTEX combines the communications capabilities of COMTEX with DIGITAL's versatile DOS/BATCH-11 single-task disk operating system. It can serve as either a powerful program development tool or a run-time executive. COMTEX is handled by DOS as a user program and is called into core memory when needed.

DOS file-handling services provide considerable flexibility. Files can be accessed randomly or sequentially. DOS also provides a useful directory structure that enables quick location of files, file protection, and the ability to append one file to another.

The services available with DOS include date and time of day, number conversion routines, and console command functions such as starting or stopping a program and transferring data from a device to a user buffer.

As a program development tool, DOS offers a full complement of utilities such as the EDIT-11 text editor, MACRO-11 for translation of assembly language source statements into relocatable object files, and LINK-11 for combining modules into a binary load module.

Since these utilities are for program development, they are not available to terminal users in a communications environment. DOS also has several useful off-line capabilities (not available to terminal users) such as FORTRAN and BATCH, which can be used for running local FORTRAN jobs or for editing files created from terminals.

DOS/COMTEX requires a PDP-11 with a minimum of 16K words of core memory and mass storage. It can expand to 28K and multiple disk drives. DOS supports a variety of DIGITAL I/O peripherals which augment the capabilities of a communications system. Among these are:

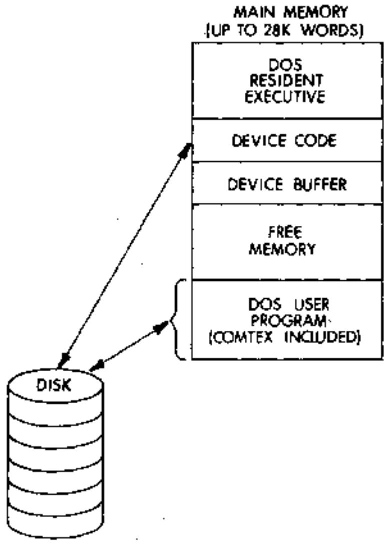

Figure A-1—DOS/COMTEX

RK11 —DECpack 1.2M word removable disk cartridge system

RP11 —20M word removable disk pack system

TA11 —DECcassette system

TC11 —DECtape system

TM11 —Magnetic tape system

LP11 —Line Printer

CR11 —Card Reader

RF11 —Fixed Head Disk

RSX-11M

DIGITAL's RSX-11M operating system provides multi-task capabilities without sacrificing high-speed real-time response. It has the ability to control a large number of communications devices and to simultaneously perform such tasks as on-line data reduction, message store and forward, spooling of data to I/O devices, and communications circuit load leveling.

The natural synchronization requirements of programs—coupled with the disparity between the time required for I/O transfers or communications control, and the time required for processing—can result in idle time for system resources. Multiprogramming is a method of enhancing system efficiency by building a queue of demands for resources. Leveling of demand is achieved by maintaining in main memory two or more tasks waiting for a system resource. The concurrent tasks are then multiplexed, the needs of one task being serviced during the "dead time" intervals of the other task.

Multiprogramming is accomplished by dividing available memory into a number of named, fixed-sized partitions. Tasks are built to execute out of a specific partition, and all partitions in the system can operate in parallel. RSX-11M has been designed to keep many tasks operating simultaneously without degrading critical response time by minimizing system overhead. Communications throughput is maximized by keeping all I/O handlers core resident. Control modules and applications programs are supported as tasks by the operating system.

RSX-11M consists of a core resident executive and provides the management facilities to allocate system resources to user and system tasks. It also resolves the conflicts that may arise between several tasks competing for the same resources. It provides such services as:

- Task scheduling on a priority basis—to greatly facilitate priority message service in a message switching application.
- Time-dependent task scheduling—for message store and forward applications and periodic interrogation of batch devices.
- Task I/O requesting—enables leveling of line loads by selection of alternative routing and message transmission to alternate stations if the addressed terminal is unavailable.
- Intertask communications—permit the easy implementation of complex editing and error detection; provide the capability of run-time responses from the communications control processor's data base; and enable the user to easily log communications activity.

- Multiprogramming services—enables the processing and file management of the transaction data collected via the terminals to be done in real-time and/or batch mode by the terminal control processor. This provides faster results and reduces the volume of data transmitted to a host computer. (In fact, the link to the host may now be dial-up and/or eliminated). In addition, multiprogramming facilitates the spooling of data for more efficient communications.

The RSX-11M executive also provides service routines such as:

- Power fail/system restart
- Memory partition management
- Hardware memory management option support

The RSX-11M executive is written in a modular fashion. This permits modules to be assembled at system generation time using conditional assembly techniques.

I/O handlers are linked into the system at system generation time. These memory-resident routines result in an overall savings in memory allocation and system response time.

The I/O handlers provide funtions such as:

- Interrupt servicing
- Address checking
- I/O queue management

Both user and system tasks (such as the FORTRAN compiler and MACRO assembler) are treated identically under RSX-11M. Tasks are referenced by name and have an execution priority associated with them. They execute in a specific partition or sub-partition. All tasks can be made of a single unit or can be built with overlays. They always occupy contiguous memory locations when they execute, with the exception of sharable libraries and sharable data areas, which need not be contiguous with the task itself.

The low system overhead, coupled with core resident I/O handlers linked at system generation time, enable RSX-11M to operate on a 16K word PDP-11 with mass storage. RSX-11M will expand to handle up to 64K words of core and additional disks as required.

RSX-11M supports the following communications devices:

I/O Handlers

DL11	—Single asynchronous line interface (excluding modem control support)
DJ11	—16-line asynchronous multiplexer
DH11/DM11-BB	—Programmable 16-line asynchronous multiplexer with modem control
DU11	—Synchronous line interface (up to 4800 bps)
DP11	—Synchronous line interface (up to 4800 bps)

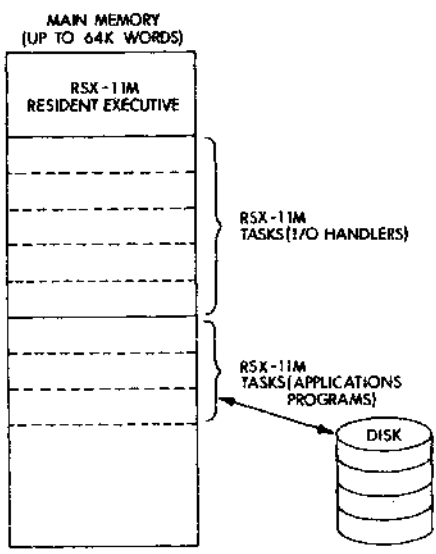

Figure A-2—RSX-11M

Control Modules
ASCII Teletype (full duplex, 8-bit only)
Console terminal

In addition, RSX-11M supports a large number of DIGITAL I/O peripherals, which can add even more capability to a communications system. Among these are:

RK11	—DECpack 1.2M word removable disk cartridge system
RP11	—20M word removable disk pack system
RJSO3/04	—Fixed Head Disk systems
TC11	—DECtape system
TA11	—DECcassette system
LP11	—Line Printers
LS11	—Line Printer
CR11	—Card Reader

RSX-11D

RSX-11D is an advanced real-time system with comprehensive multitask capabilities. It is fully upward compatible from, but more powerful than, RSX-11M. The operating system software is capable of reallocating system resources dynamically in response to changing system demands. Memory utilization is optimized by eliminating the need to preallocate memory. Permanent memory areas are allocated only for very high priority tasks. All remaining memory is shared dynamically among tasks, including batch, which reside on disk. When needed, these tasks load into any available memory area. They do not have to wait for preassigned parts of memory to become free. As soon as a task completes, its memory becomes available to any other task. If the first task loaded into such a space does not fill it, the remainder is used for another task. If the memory space is full when an urgent task requires loading, low priority tasks may be checkpointed (automatically interrupted and moved back to disk temporarily) to make room for the new task.

RSX-11D consists of an executive which is core resident and provides the management facilities for dynamic memory allocation, task relocation and loading, queueing I/O requests, and file management. Volume protection for data storage is also provided.

I/O handlers are treated as tasks to permit on-line loading and unloading as needed.

Both user and system tasks are treated similarly under RSX-11D. Several tasks may share the same memory partition and yet be completely protected from the other. Although every task is hardware protected from every other task, more than one task can still share common data and common subprograms. RSX-11D systems may utilize any number of global Common areas, with each task having access only to the ones it needs. The same Common may be made available to one set of tasks for both reading and writing while it is hardware-write protected to another set of tasks. All tasks also have access to a library of commonly used subroutines on a hardware-protected, read-only basis.

RSX-11D is a base upon which the user can build a complex multitask communications system. I/O handlers, control modules and applications programs are all treated as tasks and are easily integrated to form a complete, working system. RSX-11D requires a 48K PDP-11 with mass storage, with support for expansion up to 124K core and additional disks as required.

The following communications building blocks are available and supported by RSX-11D:

I/O Handlers

DL11-A, -B	—Single asynchronous line interface
DJ11	—16-line asynchronous line interface
DH11/DM11-BB	—16-line programmable asynchronous line multiplexer with modem control

Control Modules

ASCII Teletype (full duplex, 8-bit only)
Console terminal

Applications program

2780 Emulator

In addition, RSX-11D supports a large repertoire of DIGITAL I/O peripherals which may be useful in communications applications. Among these are:

RF11	—256K word fixed head disk
RK11	—DECpack 1.2M word removable disk cartridge system
TA11	—DECcassette system

TC11 —DECtape system

TM11 —Magnetic tape

LP11 —Line Printers

CR11 —Card Reader

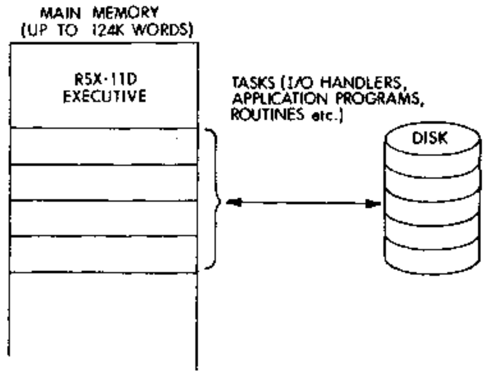

Figure A-3—RSX-11D

RSTS/E for Interactive Timesharing

RSTS/E (Resource Sharing Timesharing System/Extended) is the PDP-11's timesharing system, designed for use wherever large numbers (16 to 32) of interactive terminals are the fundamental requirement. This requirement may be in a school or scientific research department, where users are constantly developing new programs as well as running existing ones. Or it may be in a business area (such as on-line order entry where the mix of dedicated applications programs may change very little but the on-line data base is constantly being updated). The power of the RSTS/E system makes it uniquely appropriate for either type of application—or both simultaneously.

For the scientific user, RSTS/E offers the computational power of the BASIC-PLUS language. BASIC-PLUS is one of the most powerful BASIC languages ever developed, with more than forty basic commands, 35 built-in functions, and three different data types: integer, string, and floating point (single precision or double precision). Each of these data types may be used in arrays which, because of RSTS/E's virtual data storage capability, are not limited in size by the size of main memory. RSTS/E programs have access not only to disk files, but to magnetic tapes, card readers, and line printers as well. A single program may have as many as 12 files open simultaneously.

For the business-oriented user, RSTS/E offers the data management power of the RSTS/E system. Individual files as large as 33 million characters may be accessed by sequential, random, or indexed access methods. Many users may be accessing and even updating the same file simultaneously, on a fully protected basis. (The record locking facility of RSTS/E assures that two users will never end up updating the same record simultaneously since this can result in an erroneous entry.) Files that do not need to be accessed randomly may be stored on magnetic tape instead. Both disk packs and tape reels may be mounted and dismounted on-line to assure maximum utilization of individual drives. Information from these files can be precisely formatted for display at individual terminals or routed to a high-speed line printer for printing. RSTS/E's output formatting controls include comma insertion, floating dollar sign, trailing minus sign, and asterisk protect. A single system may have as many as eight line printers, all operating under control of a spooler system.

RT-11 F/B Operating System

RT-11 is a high performance Foreground/Background operating system that combines powerful computing hardware with user-oriented software. RT-11 is designed for the single user involved in program development and/or real-time applications. It provides fast, simple, on-line access to the full power of any DIGITAL PDP-11 processor with at least 8K words of memory plus mass storage (16K words of memory for F/B operation).

RT-11's interactive nature and quick response offer give-and-take flexibility plus rapid turnaround. Its queued, real-time I/O and fast throughput take full advantage of the outstanding real-time features of the PDP-11—features like hardware stack processing, multi-level, multi-line priority system, and vectored interrupts.

The RT-11 operating system actually provides two monitors: the single-job monitor and the F/B monitor. The F/B monitor allows two programs to operate: a foreground program and a background program. The real-time function is accomplished in the foreground which has priority on system resources. Functions which do not have critical response time requirements (e.g., program development) are accomplished in the background, which operates whenever the foreground is not busy. Within their priorities, both foreground and background are complete RT-11 systems with access to all system functions. Although they operate independently, foreground and background can communicate through disk files and/or job communication areas in memory.

If F/B operation is not required, the single-job monitor—which requires less memory and overhead—can be utilized. Should requirements change, upgrading is easily accomplished since programs are completely interchangeable between the single-job monitor and the F/B monitor.

appendix b

common carrier offerings

Common carriers offer many facilities and services which can be used to configure a data communications network. Just a few years back the term "common carrier" was synonymous with American Telephone and Telegraph Corporation (AT&T) or Western Union (WU). Recently, however, the FCC passed a ruling which allows other companies to enter the communications industry as a common carrier. Among these new companies are Datran and Microwave Communications, Inc. (MCI). Datran proposes to transfer data on its switched network. MCI will use microwave links to connect cities.

FACILITIES

To determine what facility or facilities will be used in a data communications network, one must consider the requirements of the system. The facilities or types of lines available should be considered as pipes which carry the data to and from the various devices (computers and terminals) comprising the network. The pipe must be sized to accommodate the amount of data the terminals and computers can send or receive. It would be incorrect to connect a 300 character-per-second terminal to a line which can only handle 15 characters per second. The system would be unusable. It would be equally incorrect to attach a 15 character-per-second terminal to a line which is capable of handling 300 characters per second. The system would be workable but not economical.

With leased lines, the tariffs are broken down into four basic categories: teletype, sub-voice, voice, and wide-band services. Lines are usually categorized by their maximum data transfer speeds in bits per second.

Telegraph grade lines will operate at speeds up to 75 Baud. This is equivalent to 100 words per minute with teletype equipment. This grade line is the slowest type available and operates on a Direct Current (DC) carrier. Because of the DC carrier and slow speed, digital-to-analog signal conversion is not required. Lines which are capable of handling higher-speed transmission require that only analog signals be transmitted. Therefore, with the exception of teletype speeds, digital data must be converted to analog signals at the transmitting stations and back to digital data at the receiving station.

Sub-voice grade lines are capable of transmitting data for the popular 15 character-per-second terminals. The facility operates at speeds up to 180 Baud. Any higher speed transmission must be handled by voice-grade lines.

Voice-grade lines are capable of handing voice transmission without distortion. These lines are capable of transmitting data at speeds up to 10,000 Baud, although they are normally used in systems with transmission speed requirements of 2400 Baud.

The highest grade facility is the wide-band lines. These lines can handle the equivalent of 60-to-120 voice grade lines.

With each of these facilities, a different rate is charged. These rates are based on a per-mile-per-month tariff.

CONDITIONING

Conditioning applies to voice grade lines. The better conditioned a line is, the less distortion is found and errors are less. Conditioning ranges from C1 to the best, C4. The common carrier will often recommend no conditioning for lines transmitting at 1200 Baud, C1 conditioning is recommended for 2400 Baud, C2 for 4800 Baud, and C4 for speeds above 4800 Baud.

SERVICES

There are two basic types of services: switched and leased. Switched service (sometimes called dial-up) employs the use of the public switched network. Leased services are based on dedicated lines whose facilities were dscribed earlier.

The tariffs for data transmission using switched (dial-up) services fall under AT&T Data-Phone and Western Union broad band. The advantage of using switched services is that you pay for what you use (billing is at normal phone rates). This service will transmit at voice grade data rates. If a terminal does not have much traffic, this service would be economical. The disadvantage in using this service is that it is subject to all the problems encountered in the public switched network. There can be long connect times, disconnects, and busy signals. For a system requiring short response times, this service is not advisable.

Leased or private lines are dedicated to the user. The advantage is that the terminal or computer is always physically connected to the line. Very short response times are met with this service. This service is billed on a 24-hour-a-day, 7-day-a-week basis.

Some of the more widely used leased services are summarized below:

WATS (Wide Area Telephone Service)

- Basically Dial Facility
- Country Split into Zones
- Full Service
- Measured Service
- Different from Dial (Zones)
- WATS In Service
- WATS Out Service

TELPAK

- Bundles of Voice Grade Leased Lines
- 12 or 24 Intrastate
- 60 or 120 Interstate
- Shared Services

TELEX (Teletypewriter Exchange)

- 66 words per minute
- Certain Cities

TWX (Teletypewriter Exchange Service)

- 60 wpm, 150 Baud Operation

Data Phone 50

- Dial-up 50 KB Service
- Between Specific Cities
- Equivalent to 12 Voice Frequency Channels

DATACOM—Western Union

- Multiplexer between two or three cities
- Derive Multiple Low Speed and TTY Lines
- Joint Users (Same Industry)

appendix C
glossary

This appendix contains definitions of terms unique to data communications networks. Definitions used are common to the industry where such definitions exist. Other definitions were specially written for this glossary.

ACCESS CONTROL
The tasks imposed on a network or any of its components, performed by hardware, software, and administrative controls, to control usage of the system. Included are monitoring of system operation, insuring of data integrity, user identification, recording system access and changes, and methods for granting users access.

ACK 0, ACK 1
(AFFIRMATIVE ACKNOWLEDGMENT)
These replies (DLE sequences in Binary Synchronous Communications) indicate that the previous transmission block was accepted by the receiver and that it is ready to accept the next block of the transmission. Use of ACK 0 and ACK 1 alternately provides sequential checking control for a series of replies. ACK 0 is also an affirmative (ready to receive) reply to a station selection (multipoint), or to an initialization sequence (line bid) in point-to-point operation.

ACOUSTIC COUPLER
A device that converts electrical signals into audio signals, enabling data to be transmitted over the public telephone network via a conventional telephone handset.

ALTERNATE ROUTE
A secondary path used to reach a destination if the primary path is unavailable.

ASYNCHRONOUS TRANSMISSION

Transmission in which time intervals between transmitted characters may be of unequal length. Transmission is controlled by start and stop elements at the beginning and end of each character. Also called Start-Stop transmission.

ASCII

American Standard Code for Information Interchange. This is a seven-bit-plus-parity code established by the American National Standards Institute (formerly American Standards Association) to achieve compatibility between data services. Also called USASCII.

AMPLITUDE MODULATION (AM)

A method of transmission whereby the amplitude of the carrier wave is modified in accordance with the amplitude of the signal wave.

AUDIO FREQUENCIES

Frequencies which can be heard by the human ear (usually between 15 cycles and 20,000 cycles per second).

AUTOMATIC CALLING UNIT (ACU)

A dialing device supplied by the communications common carrier. This device permits a business machine to automatically dial calls over the communications network.

BACKGROUND PROCESSING

The automatic execution of a low-priority computer program when higher priority programs are not using the system resources.

BANDWIDTH

The range of frequencies assigned to a channel or system. The difference expressed in Hertz between the highest and lowest frequencies of a band.

BASEBAND SIGNALLING

Transmission of a signal at its original frequencies, i.e., unmodulated.

BATCH PROCESSING

A technique of data processing in which jobs are collected and grouped before processing. Data thus are normally processed in a deferred mode.

BAUD

A unit of signalling speed equal to the number of discrete conditions or signal events per second. In asynchronous transmission, the unit of signalling speed corresponding to one unit interval per second; that is, if the duration of the unit interval is 20 milliseconds, the signalling speed is 50 Baud. Baud is the same as "bits per second" only if each signal event represents exactly one bit.

BAUDOT CODE

A code for the transmission of data in which five bits represent one character. It is named for Emile Baudot, a pioneer in printing telegraphy. The name is usually applied to the code used in many teleprinter systems, which was first used by Murray, a contemporary of Baudot.

BINARY DIGIT (BIT)

In the binary notation either of the characters 0 or 1. "Bit" is the commonly used abbreviation for Binary Digit.

BINARY SYNCHRONOUS COMMUNICATIONS (BSC)

A uniform discipline, using a defined set of control characters and control character sequences, for synchronized transmission of binary coded data between stations in a data communications system. Also called BISYNC.

BIT

Abbreviation for BINARY DIGIT.

BIT TRANSFER RATE

The number of bits transferred per unit time, usually expressed in Bits Per Second (BPS).

BLOCK

A group of digits transmitted as a unit, over which a coding procedure is usually applied for synchronization or error control purposes. See also: Packet.

BLOCK CHARACTER CHECK (BCC)

The result of a transmission verification algorithm accumulated over a transmission block, and normally appended at the end: e.g., CRC, LRC.

BROADBAND
See Wideband.

BROADBAND EXCHANGE (BEX)
Public switched communication system of Western Union featuring various bandwidth full-duplex connections.

BUFFER
A storage device used to compensate for a difference in the rate of data flow when transmitting data from one device to another.

BYTE
A binary element string operated upon as a unit and usually shorter than a computer word, e.g., six-bit, eight-bit, or nine-bit bytes.

CARRIER
A continuous frequency capable of being modulated or impressed with a signal.

CARRIER SYSTEM
A means of obtaining a number of channels over a single path by modulating each channel upon a different "carrier" frequency and demodulating at the receiving point to restore the signals to their original form.

CATHODE-RAY TUBE (CRT)
A television-like picture tube used in visual display terminals.

CCITT
Comité Consultatif Internationale de Telegraphie et Telephonie. An international consultative committee that sets international communications usage standards.

CENTRALIZED (COMPUTER) NETWORK
A computer network configuration in which a central node provides computing power, control, or other services. Compare: Decentralized Network.

CHANNEL
That part of a communications system that connects a message

source to a message sink. See also: Information Transfer Channel.

CHANNEL CAPACITY

A term which expresses the maximum bit rate that can be handled by the channel.

CIRCUIT

In communications the complete electrical path providing one or two-way communication between two points comprising associated go and return channels. Compare: Channel.

CIRCUIT SWITCHING

A method of communications, where an electrical connection between calling and called stations is established on demand for exclusive use of the circuit until the connection is released. See also: Packet Switching, Store and Forward, Message Switching.

CODE

1) A set of unambiguous rules specifying the way in which data may be represented, e.g., the set of correspondences in the Standard Code for Information Interchange. 2) In data communications, a system of rules and conventions according to which the signals representing data can be formed, transmitted, received and processed. 3) In data processing, to represent data or a computer program in a symbolic form that can be accepted by a data processor.

COMMON CARRIER

In data communications, a public utility company that is recognized by an appropriate regulatory agency as having a vested interest and responsibility in furnishing communication services to the general public, e.g., Western Union, The Bell System. See also: Specialized Common Carrier, Value Added Service.

COMMUNICATIONS COMPUTER

A computer that acts as the interface between another computer or terminal and a network, or a computer controlling data flow in a network. See also: Front-End Computer, Switching Computer, Concentrator.

COMMUNICATION CONTROL CHARACTER

A functional character intended to control or facilitate transmission over data networks. There are ten control characters specified in ASCII which form the basis for character-oriented communications control procedures. See also: Control Character.

COMPUTER NETWORK

An interconnection of assemblies of computer systems, terminals and communications facilities.

CONCENTRATOR

A communications device that provides communications capability between many low speed, usually asynchronous channels and one or more high speed, usually synchronous channels. Usually different speeds, codes, and protocols can be accommodated on the low-speed side. The low-speed channels usually operate in contention requiring buffering. The concentrator may have the capability to be polled by a computer, and may in turn poll terminals.

CONDITIONING

The addition of equipment to leased voice-grade lines to provide specified minimum values of line characteristics required for data transmission, e.g., equalization and echo suppression.

CONNECT TIME

A measure of system usage by a user, usually the time interval during which the user terminal was on line during a session. See also: CPU Time.

CONSOLE

1) A part of a computer used for communication between operator or maintenance engineer and the computer. 2) Part of a terminal providing user input and output capability.

CONTENTION

A condition on a communications channel when two or more stations try to transmit at the same time.

CONTROL CHARACTER

1) A character whose occurrence in a particular context initiates, modifies, or stops a control function. 2) In the ASCII code, any of the 32 characters in the first two columns of the standard code table. See also: Communications Control Character.

CONTROL PROCEDURE

The means used to control the orderly communication of information between stations on a data link. Syn: Line Discipline. See also: Protocol.

CONTROL STATION

The station on a network which supervises the network control procedures such as polling, selecting and recovery. It is also responsible for establishing order on the line in the event of contention, or any other abnormal situation, arising between any stations on the network. Compare: Tributary Station.

CONVERSATIONAL

Pertaining to a mode of processing that involves step-by-step interaction between the user at a terminal by means of keyboard and display and a computer. See also: Interactive.

CPU TIME

Central Processing Unit Time; a measure of system usage by a user, based on the total amount of computer processing time used. See also: Connect Time.

CROSS TALK

The unwanted transfer of energy from one circuit, called the disturbing circuit, to another circuit, called the disturbed circuit.

CYCLIC REDUNDANCY CHECK (CRC)

An error detection scheme in which the check character is generated by taking the remainder after dividing all the serialized bits in a block of data by a predetermined binary number.

DATA ACCESS ARRANGEMENT (DAA)

Data communication equipment furnished by a common carrier, permitting attachment of privately owned data terminal and data communication equipment to the common carrier network.

DATA ACQUISITION

The retrieval of data from remote sites initiated by a central computer system; e.g., retrieving data during off-hours processing from a previously mounted magnetic tape at an unattended terminal; or taking periodic readings from an unattended real-time station.

DATA BASE

1) The entire collection of information available to a computer system. 2) A structured collection of information as an entity or collection of related files treated as an entity.

DATA COMPRESSION

A technique whereby a repetitive string of data (usually on a byte basis) is transmitted as a count plus a string value.

DATA CONCENTRATION

Collection of data at an intermediate point from several low and medium-speed lines for retransmission across high-speed lines.

DATA COLLECTION

The act of bringing data from one or more points to a central point.

DATA COMMUNICATION

The interchange of data messages from one point to another over communications channels. See also: Data Transmission.

DATA COMMUNICATION EQUIPMENT

The equipment that provides the functions required to establish, maintain, and terminate a connection, the signal conversion, and coding required for communication between data terminal equipment and data circuit. The data communication equipment may or may not be an integral part of a computer; e.g., a modem. See also: Terminal Installation, Data Link.

DATA INTEGRITY

A performance measure based on the rate of undetected errors.

DATA LINK

An assembly of terminal installations and the interconnecting circuits operating according to a particular method that permits information to be exchanged between terminal installations. NOTE: The method of operation is defined by particular transmission codes, transmission modes, direction, and control.

DATA ORIGINATION

The earliest stage at which the source material is first put into machine-readable form or directly into electrical signals.

DATA-PHONE SET/SERVICE

The data sets or modems manufactured and supplied by the Bell System for use in the transmission of data over the regular telephone network. The service mark, of the Bell System for the transmission of data over the regular telephone network.

DATA-PHONE 50

A public switched communications service of the Bell System featuring high-speed data communications at 50 kbps.

DATA-PHONE DIGITAL SERVICE (DDS)

A communications service of the Bell System in which data is transmitted in digital rather than analog form, thus eliminating the need for modems.

DATA SET

1) A modem. 2) A collection of data records, with a logical relation of one to another. See also: Data Phone, Modem.

DATA SHARING

The ability of users or computer processes at several nodes to access data at a single node.

DATA TERMINAL EQUIPMENT (DTE)

1) The equipment comprising the data source, the data sink, or both. 2) Equipment usually comprising the following functional units: control logic, buffer store, and one or more input or output devices or computers. It may also contain error control, synchronization, and station identification capability. See also: Data Communications Equipment, Data Link, Terminal Installation.

DATA TRANSMISSION

The sending of data from one place for reception elsewhere. Compare: Data Communication.

DDCMP

Digital Data Communications Message Protocol. A uniform discipline for the transmission of data between stations in a point-to-point or multipoint data communications system. The method of physical data transfer used may be parallel, serial synchronous, or serial asynchronous.

DECENTRALIZED (COMPUTER) NETWORK

A computer network, where some of the network control functions are distributed over several network nodes. Compare: Centralized Network.

DELAY DISTORTION

Distortion resulting from non-uniform speed of transmission of the various frequency components of a signal through a transmission medium.

DELIMITER

A character that separates and organizes elements of data.

DEMODULATION

The process of retrieving an original signal from a modulated carrier wave. This technique is used in data sets to make communication signals compatible with computer signals.

DIAL-UP LINE

A communications circuit that is established by a switched circuit connection.

DIRECT DISTANCE DIALING (DDD)

A telephone exchange service which enables a user to directly dial telephones outside his local area without operator assistance.

DIRECT MEMORY ACCESS (DMA)

A facility that permits I/O transfers directly into or out of memory without passing through the processor's general registers; either performed independently of the processor or on a cycle-stealing basis.

DISTRIBUTED NETWORK

A network configuration in which all node pairs are connected either directly, or through redundant paths through intermediate nodes. Compare: Fully Connected Network.

DLE (Data Link Escape)

A control character used exclusively to provide supplementary line-control signals (control character sequences or DLE sequences). These are two-character sequences where the first character is DLE. The second character varies according to the function desired and the code used.

DUPLEX

Simultaneous two-way independent transmission in both directions. Also referred to as full-duplex.

EBCDIC

Extended Binary Coded-Decimal Interchange Code. An 8-bit character code used primarily in IBM equipment. The code provides for 256 different bit patterns.

ECHO

A portion of the transmitted signal returned from the distant point to the source with sufficient magnitude and delay so as to cause interference.

ECHO CHECK

A method of checking the accuracy of transmission of data in which the received data are returned to the sending end for comparison with the original data.

ECHO SUPPRESSOR

A device used to suppress the effects of an echo.

ELECTRONIC INDUSTRIES ASSOCIATION (EIA)

A standards organization specializing in the electrical and functional characteristics of interface equipment.

ELECTRONIC SWITCHING SYSTEM (ESS)

The common carrier communications switching system which uses solid state devices and other computer-type equipment and principles.

ENQ (Enquiry)

Used as a request for response to obtain identification and/or an indication of station status. In Binary Synchronous (BSC) transmission, ENQ is transmitted as part of an initialization sequence (line bid) in point-to-point operation, and as the final character of a selection or polling sequence in multipoint operation.

EOT (End of Transmission)

Indicates the end of a transmission, which may include one or more messages, and resets all stations on the line to control mode (unless it erroneously occurs within a transmission block).

EQUALIZATION

Compensation for the increase of attenuation with frequency. Its purpose is to produce a flat frequency response.

ETX (End of Text)

Indicates the end of a message. If multiple transmission blocks are contained in a message in BSC systems, ETX terminates the last block of the message. (ETB is used to terminate preceding blocks). The block check character is sent immediately following ETX. ETX requires a reply indicating the receiving station's status.

FREQUENCY DIVISION MULTIPLEXING (FDM)

Dividing the available transmission frequency range into narrower bands each of which is used for a separate channel.

FREQUENCY MODULATION (FM)

A method of transmission whereby the frequency of the carrier wave is changed to correspond to changes in the signal wave.

FREQUENCY SHIFT KEYING (FSK)

Also called frequency shift signalling. A method of frequency modulation in which frequency is made to vary at significant instants by smooth as well as abrupt transitions.

FRONT END PROCESSOR

A communications computer associated with a host computer. It may perform line control, message handling, code conversion, error control and applications functions such as control and operation of special-purpose terminals. See also: Communications Computer.

FOREGROUND PROCESSING

High-priority processing, usually resulting from real-time entries, given precedence by means of interrupts, over lower priority "background" processing.

FOREIGN EXCHANGE LINE

A line offered by a common carrier in which a termination in one central office is assigned a number belonging to a remote central office.

FORWARD CHANNEL

A data transmission channel in which the direction of transmission coincides with that in which information is being transferred. Compare: Reverse Channel.

FORWARD SUPERVISION

Use of supervisory sequences sent from the primary to a secondary station or node.

FRAME

See: BLOCK

FULL DUPLEX

See: DUPLEX

FULLY CONNECTED NETWORK

A network in which each node is directly connected with every other node.

HALF DUPLEX

A circuit designed for transmission in either direction but not both directions simultaneously.

HARDWARE

Physical equipment, as opposed to a computer program or method of use, e.g., mechanical, electrical, magnetic, or electronic devices.

HASP

Houston Automatic Spooling Program. An IBM 360/370 OS software front-end which performs job spooling and controls communications between local and remote processors and Remote Job Entry (RJE) stations.

HEADER

The control information prefixed in a message text, e.g., source or destination code, priority, or message type. Syn: Heading, Leader.

HERTZ

A unit of frequency equal to one cycle per second. Cycles are referred to as Hertz in honor of the experimenter Heinrich Hertz. Abbreviated Hz.

HETEROGENEOUS (COMPUTER) NETWORK

A network of dissimilar host computers, such as those of different manufacturers. Compare: Homogeneous Network.

HIERARCHICAL (COMPUTER) NETWORK

A computer network, in which processing and control functions are performed at several levels by computers specially suited for the functions performed, e.g., in factory or laboratory automation.

HOMOGENEOUS (COMPUTER) NETWORK

A network of similar host computers such as those of one model of one manufacturer. Compare: Heterogeneous (Computer) Network.

HOST COMPUTER

A computer attached to a network providing primarily services such as computation, data base access or special programs or programming languages.

HOST INTERFACE

The interface between a communications processor and a host computer.

IDENTIFICATION

1) The process of providing personal, equipment, or organizational characteristics or codes to gain access to computer programs, processes, files, or data. 2) The process of determining personal, equipment, or organizational characteristics or codes to permit access to computer programs, processes, files or data.

INFORMATION BIT

A bit which is generated by the data source and which is not used for error control by the data transmission system. Compare: Overhead Bit.

INFORMATION PATH

The functional route by which information is transferred in a one-way direction from a single data source to a single data sink.

INFORMATION (TRANSFER) CHANNEL

1) The functional connection between the source and the sink data terminal equipments. It includes the circuit and the asso-

ciated data communications equipments. 2) The assembly of data communications and circuits including a reverse channel if it exists.

INTERACTIVE

Pertaining to exchange of information and control between a user and a computer process, or between computer processes. See also: Conversational.

INTERCHANGE POINT

A location where interface signals are transmitted between equipments by means of electrical interconnections. See also: Interface.

INTERFACE

1) A shared boundary defined by common physical interconnection characteristics, signal characteristics, and meanings of interchanged signals. 2) A device or equipment making possible interoperation between two systems, e.g., a hardware component or a common storage register. 3) A shared logical boundary between two software components.

ITB (Intermediate Text Block)

In Binary Synchronous Communications, a control character used to terminate an intermediate block of characters. The block check character is sent immediately following ITB, but no line turnaround occurs. The response following ETB or ETX also applies to all of the ITB checks immediately preceding the block terminated by ETB or ETX.

LEASED LINE

A line reserved for the exclusive use of a leasing customer without interexchange switching arrangements. Also called Private Line.

LINE

1) The portion of a circuit external to the apparatus consisting of the conductors connecting a telegraph or telephone set to the exchange or connecting two exchanges. 2) The group of conductors on the same overhead route in the same cable.

LINK

1) Any specified relationship between two nodes in a network.

2) A communications path between two nodes. 3) a Data Link. See also: Line, Circuit, Virtual Circuit.

LINK REDUNDANCY LEVEL

The ratio of actual number of links to the minimum number of links required to connect all nodes of a network. See also: Fully Connected Network.

LOAD SHARING

The distribution of a given load among several computers on a network.

LOCAL EXCHANGE

An exchange in which subscribers' lines terminate. Also called "End Office."

LOGIN

A user access procedure to a system involving identification, access control and exchange of network information between user and system. Syn: Logon.

LOGOUT

A user exit procedure from a system often providing usage statistics to the user. Syn: Logoff.

LONGITUDINAL REDUNDANCY CHECK (LRC)

An error checking technique based on an accumulated exclusive OR of transmitted characters. An LRC character is accumulated at both the sending and receiving stations during the transmission of a block. This accumulation is called the Block Check Character (BCC), and is transmitted as the last character in the block. The transmitted BCC is compared with the accumulated BCC character at the receiving station for an equal condition. An equal comparison indicates a good transmission of the previous block.

MARK

Presence of a signal. In telegraphy, mark represents the closed condition or current flowing. Equivalent to a binary one condition.

MASTER STATION,

See Primary Station.

MESSAGE SWITCHING

A method of handling messages over communications networks. The entire message is transmitted to an intermediate point (i.e., a switching computer), stored for a period of time, perhaps very short, and then transmitted again towards its destination. The destination of each message is indicated by an address integral to the message. Compare: Circuit Switching.

MODEM

Modulator-Demodulator. A device that modulates signals transmitted over communications circuits. Syn: Data Set.

MULTILEAVING

A technique which allows simultaneous bidirectional communications traffic; e.g., output from a previous remote batch job may be received while a new job is being transmitted.

MULTIPLEXING

A division of a transmission facility into two or more channels. See also: Frequency Division Multiplexing and Time Division Multiplexing.

MULTIPLEXER

A device used for multiplexing. It may or may not be a stored program computer.

MULTI-POINT LINE

A single communications line to which more than one terminal is attached. Use of this type of line normally requires some kind of polling mechanism, addressing each terminal with a unique ID. Also called: "Multi-Drop."

NEGATIVE ACKNOWLEDGMENT (NAK)

Indicates that the previous transmission block was in error and the receiver is ready to accept a retransmission of the erroneous block. NAK is also the "not ready" reply to a station selection (multipoint) or to an initialization sequence (line bid) in point-to-point operation.

NARROWBAND CHANNELS

Sub-voice grade channels characterized by a speed range of 100 to 200 bits per second.

NETWORK

1) An interconnected or interrelated group of nodes. 2) In connection with a disciplinary or problem oriented qualifier, the combination of material, documentation, and human resources that are united by design to achieve certain objectives, e.g., a social science network, a science information network. See Also: Computer Network.

NETWORK CONTROL PROGRAM

That module of an operating system in a host computer, which establishes and breaks logical connections, communicating with the network on one side, and with user processes within the host computer, on the other side.

NETWORK OPERATIONS CENTER

A specialized network installation that assists in reliable network operations. Typical activities are monitoring of network status, supervision and coordination of network maintenance, accumulation of accounting and usage data, and user support.

NETWORK REDUNDANCY

The property of a network to have additional links beyond the minimum number necessary to connect all nodes. See also: Link Redundancy Level.

NETWORK SECURITY

The totality of measures taken to protect a network from an unauthorized access, accidental or willful interference with normal operations, or destruction. This includes protection of physical facilities, software, and personnel security. See also: Privacy.

NETWORK TOPOLOGY

The geometric arrangement of links and nodes of a network.

NODE

1) An end point of any branch of a network, or a junction common to two or more branches of a network.

NOISE

Undesirable disturbances in a comunications system. Noise can generate errors in transmission.

NON-PROCESSOR REQUEST (NPR)

High priority data transfers to the PDP-11 Processor. These are direct memory access type transfers, and are honored by the processor between bus cycles of an instruction execution. NPR data transfers can be made between any two peripheral devices without the supervision of the processor. Normally, NPR transfers are between a mass storage device, such as a disk and core memory. An NPR device has very fast access to the bus and can transfer at high data rates once it has control. The processor state is not affected by the transfer; therefore, the processor can relinquish control while an instruction is in progress.

NON-TRANSPARENT MODE

Transmission of characters in a defined character format, e.g., ASCII or EBCDIC, in which all defined control characters and control character sequences are recognized and treated as such.

NON-SWITCHED LINE

A communications link which is permanently installed between two points.

NULL MODEM

A device which interfaces between a local peripheral that normally requires a modem, and the computer near it that expects to drive a modem to interface to that device; an imitation modem in both directions.

OFF-LINE

Pertaining to equipment or devices not under control of the central processing unit.

ONE-WAY ONLY OPERATION

A mode of operation of a data link in which data are transmitted in a preassigned direction over one channel. Syn: Simplex Operation.

ON-LINE

1) Pertaining to equipment or devices under control of the central processing unit. 2) Pertaining to a user's ability to interact with a computer.

OPERATING SYSTEM

Software that controls the execution of computer programs and that may provide scheduling, debugging, input and output control, accounting, storage assignment, data management, and related service. Sometimes called Supervisor, Executive, Monitor, Master Control Program depending on the computer manufacturer.

OVERHEAD BIT

A bit other than an information bit, e.g., check bit, framing bit.

PACKET

A group of bits including data and control elements which is switched and transmitted as a composite whole. The data and control elements and possibly error control information are arranged in a specified format.

PACKET SWITCHING

A data transmission process, utilizing addressed packets, whereby a channel is occupied only for the duration of transmission of the packet. NOTE: In certain data communication networks the data may be formatted into a packet or divided and then formatted into a number of packets (either by the data terminal equipment or by equipment within the network) for transmission and multiplexing purposes. See also: Circuit Switching, Message Switching, Store and Forward.

PARALLEL TRANSMISSION

Method of data transfer in which all bits of a character or byte are transmitted simultaneously either over separate communication lines or on different carrier frequencies on the same communication line.

PARITY CHECK

Addition of non-information bits to data, making the number of ones in each grouping of bits either always odd for odd parity or always even for even parity. This permits single error detection in each group.

PASSWORD

A word or string of characters that is recognizable by automatic means and that permits a user access to protected storage, files, or input or output devices.

PHASE MODULATION (PM)

A method of transmission whereby the angle of phase of the carrier wave is varied in accordance with the signal.

POINT-TO-POINT CONNECTION

1) A network configuration in which a connection is established between two, and only two, terminal installations. The connection may include switching facilities. 2) A circuit connecting two points without the use of any intermediate terminal or computer. Compare: Multi-point Connection.

POLLING

The process of inviting another station or node to transmit data. Compare: Selecting.

PRIMARY STATION

1) The station which at any given instant has the right to select and to transmit information to a secondary station, and the responsibility to insure information transfer. There should be only one primary station on a data link at one time. 2) A station which has control of a data link at a given instant. The assignment of primary status to a given station is temporary and is governed by standardized control procedures. Primary status is normally conferred upon a station so that it may transmit a message, but a station need not have a message to be nominated primary station.

PRIVACY

The right of an individual to the control of information about himself. Compare: Network Security.

PROCESS

1) A systematic sequence of operations to produce a specified result. 2) A set of related procedures and data undergoing execution and manipulation by one or more computer processing units.

PROGRAM SHARING

The ability for several users or computers to utilize a program at another node.

PROTOCOL

A formal set of conventions governing the format and relative timing of message exchange between two communicating processes. See also: Control Procedure.

PULSE CODE MODULATION (PCM)
Modulation of a pulse train in accordance with a code.

RANDOM ACCESS
Pertaining to a storage device where data or blocks of data can be read in any particular order (e.g., disk or DECtape). In random access devices you do not have to read from the beginning to find what you want as you do with paper tape and industry compatible magnetic tape.

REAL TIME SYSTEM
A system performing computation during the actual time the related physical process transpires, so that the results of the computation can be used in guiding the process.

REDUNDANCY
In a protocol the portion of the total characters or bits that can be eliminated without any loss of information.

REGIONAL (COMPUTER) NETWORK
1) A computer network whose nodes provide access to a defined geographical area. 2) A network whose nodes provide access to a specified class of users.

REGULATORY AGENCY
In data communications, an agency controlling common and specialized carrier tariffs, e.g., the Federal Communications Commission and the State Public Utility Commissions.

REMOTE JOB ENTRY
1) Submission of jobs through an input device that has access to a computer through a communications link. 2) The mode of operation that allows input of a batch job by a card reader at a remote site and receipt of the output via a line printer or card punch at a remote site. Abbr: RJE.

REMOTE STATION
(Multipoint) Synonymous with tributary station. (Point-to-point switched network) a station that can be called by the central station, or can call the central station if it has a message to send.

RESOURCE

Any means available to network users, such as computational power, programs, data files, storage capacity, or a combination of these.

RESOURCE SHARING

The joint use of resources available on a network by a number of dispersed users.

RESPONSE TIME

The elapsed time between the generation of the last character of a message at a terminal and the receipt of the first character of the reply. It includes terminal delay, network delay, and service node delay.

REVERSE CHANNEL

A channel used for transmission of supervisory or error-control signals. The direction of flow of these signals is in the direction opposite to that in which information is being transferred. The bandwidth of this channel is usually less than that of the Forward Channel, i.e., the information channel.

REVERSE INTERRUPT (RVI)

In Binary Synchronous Communications, a control character sequence (DLE sequence) sent by a receiving station instead of ACK1 or ACK0 to request premature termination of the transmission in progress.

RING NETWORK

A computer network where each computer is connected to adjacent computers.

SDLC

Synchronous Data Link Control. A uniform discipline for the transfer of data between stations in a point-to-point, multipoint, or loop arrangement, using synchronous data transmission techniques.

SECONDARY STATION

A station that has been selected to receive a transmission from the primary station. The assignment of secondary status is temporary, under control of the primary station, and continues for the duration of a transmission. Compare: Primary Station.

SELECTING
A process of inviting another station or node to receive data. Compare: polling.

SERIAL TRANSMISSION
A method of transmission in which each bit of information is sent sequentially on a single channel rather than simultaneously as in parallel transmission.

SIGNAL ELEMENT
Each of the parts of a digital signal, distinguished from others by its duration, position and sense, or by some of these features only. In start-stop operation a signal element has a minimum duration of one unit interval. If several unit intervals of the same sense run together, a signal element of duration of more than one unit element may be formed. Signal elements may be start elements, information elements or stop elements.

SIGNAL-TO-NOISE RATIO (SNR)
Relative power of the signal to the noise in a channel, usually measured in decibels.

SILO
A 16-bit wide by 64-word deep first-in/first-out hardware buffer used with the PDP-11's DJ11 and DH11 Asynchronous Multiplexers.

SIMPLEX MODE
Operation of a channel in one direction only with no capability of reversing.

SINK
1) The point of usage data in a network. 2) A data terminal installation that receives and processes data from a connected channel.

SLAVE

A remote system or terminal whose functions are controlled by a central "master" system. It is similar in concept to a Host system in that it responds to remotely generated requests, but unlike a Host system is usually capable of performing a limited range of operations.

SOFTWARE

A set of computer programs, procedures, rules and associated documentation concerned with the operation of network computers, e.g., compilers, monitors, editors, utility programs. Compare: Hardware.

SOURCE

1) The point of entry of data in a network. 2) A data terminal installation that enters data into a connected channel. Data entry may be under operator or machine control.

SPECIALIZED COMMON CARRIER

A company that provides private line communications services, e.g., voice, teleprinter, data, facsimile transmission. See also: Common Carrier, Value Added Service.

SPOOLING

The technique by which output to slow devices is placed into queues on mass storage devices to await transmission. This allows more efficient use of the system since programs using low-speed devices can run to completion quickly and make room for others.

STAR NETWORK

A computer network with peripheral nodes all connected to one or more computers at a centrally-located facility. See also: Centralized Network.

START ELEMENT

In start-stop transmission, the first element in each character, which serves to prepare the receiving equipment for the reception and registration of the character.

START OF HEADER (SOH)

A communication control character used at the beginning of a sequence of characters which constitute a machine-sensible address or routing information. Such a sequence is referred to as the *heading*.

START OF TEXT (STX)
A communication control character which precedes a sequence of characters that is to be treated as an entity and entirely transmitted through to the ultimate destination. Such a sequence is referred to as *text*. STX may be used to terminate a sequence of characters (heading) started by SOH.

START-STOP TRANSMISSION
Asynchronous transmission in which a group of code elements corresponding to a character signal is preceded by a start element and is followed by a stop element.

STATION
That independently-controllable configuration of data terminal equipment from or to which messages are transmitted on a data link. It includes those elements which serve as sources or sinks for the messages, as well as those elements which control the message flow on the link, by means of data communication control procedures. See also: Terminal Installation.

STOP ELEMENT
In start-stop transmission, the last element in each character, to which is assigned a minimum duration, during which the receiving equipment is returned to its rest condition in preparation for the reception of the next character.

SUPERVISORY PROGRAMS
Computer programs that have the primary function of scheduling, allocating, and controlling system resources rather than processing data to produce results.

SUPERVISORY SEQUENCE
In data communication, a sequence of communication control characters, and possibly other characters, that perform a defined control function.

SWITCHED LINE

A communications link for which the physical path may vary with each usage, e.g., the dial-up telephone network.

SYNCHRONOUS IDLE (SYN)

Character used as a time fill in the absence of any data or control character to maintain synchronization. The sequence of two continuous SYN's is used to establish synchronization (character phase) following each line turnaround.

SYNCHRONOUS TRANSMISSION

Transmission in which the data characters and bits are transmitted at a fixed rate with the transmitter and receiver synchronized. This eliminates the need for start-stop elements, thus providing greater efficiency. Compare: Asynchronous Transmission.

TARIFF

1) A published rate for services provided by a common or specialized carrier. 2) The means by which regulatory agencies approve such services. The tariff is a part of a contract between customer and carrier.

TELETYPE

Trademark of Teletype Corporation. Usually refers to one of their series of teleprinters.

TELETYPEWRITER EXCHANGE SERVICE (TWX)

A public teletypewriter exchange (switched) service in the United States and Canada formerly belonging to AT&T Company which is now owned by the Western Union Telegraph Company. Both Baudot and ASCII-coded machines are used.

TELEX SERVICE

A Western Union world-wide teletypewriter exchange service that uses the public telegraph network. Baudot equipment is used.

TELPAK

The name given to the pricing arrangement by AT&T in which many voice-grade telephone lines are leased as a group between two points.

TEMPORARY TEXT DELAY (TTD)

In Binary Synchronous Communications, a control character sequence (STX . . . ENQ) sent by a transmitting station to either indicate a delay in transmission or to initiate an abort of the transmission in progress.

TERMINAL

A device or computer which may be connected to a local or remote Host system, and for which the Host system provides computational and data access services. Two common types of terminals are timesharing (typically interactive keyboard terminals) and remote batch, such as the IBM 2780.

TERMINAL INSTALLATION

1) The totality of equipment at a user's installation including data terminal equipment, data communication equipment, and necessary support facilities. See also: Terminal, Station. 2) A set composed of data terminal, a signal converter, and possibly intermediate equipment; this set may be connected to a data processing machine or may be part of it.

TEXT

1) A sequence of characters forming part of a transmission which is sent from the data source to the data sink, and contains the information to be conveyed. It may be preceded by a header and followed by an "End of Text" signal. 2) In ASCII and communications, a sequence of characters, treated as an entity if preceded by a "Start of Text" and followed by an "End of Text" control character.

TIE LINE

A private line communications channel of the type provided by communications common carriers for linking two or more points together.

TIME-DIVISION MULTIPLEXING

A system of multiplexing in which channels are established by connecting terminals one at a time at regular intervals by means of an automatic distribution.

TIME-SHARING

A method of operation in which a computer facility is shared by several users for different purposes at (apparently) the same time. Although the computer actually services each user in sequence, the high speed of the computer makes it appear that the users are all handled simultaneously.

TORN-TAPE SWITCHING CENTER

A location where operators tear off incoming printed and punched paper tape and transfer it manually to the proper outgoing circuit.

TOUCH-TONE DEVICE

AT&T device for pushbutton dialing. The signaling form is multiple tones.

TRANSPARENT MODE

Transmission of binary data with the recognition of most control characters suppressed. In Binary Synchronous Communications, entry to and exit from the transparent mode is indicated by a sequence beginning with a special Data Link Escape (DLE) character.

TRIBUTARY STATION

A station, other than the control station, on a centralized multipoint data communications system, which can communicate only with the control station when polled or selected by the control station.

TRUNK

A single circuit between two points, both of which are switching centers or individual distribution points. Compare: Local Line.

TURNAROUND TIME

1) The elapsed time between submission of a job to a computing center and the return of results. 2) In communications the actual time required to reverse the direction of transmission from sender to receiver or vice versa when using a two-way alternate circuit. Time is required by line propagation effects, modem timing and computer reaction.

TWO WAY ALTERNATE OPERATION

A mode of operation of a data link in which data may be transmitted in both directions, one way at a time. Syn: Half-duplex Operation (US).

TWO WAY SIMULTANEOUS OPERATION

A mode of operation of a data link in which data may be transmitted simultaneously in both directions over two channels. Note: One of the channels is equipped for transmission in one direction while the other is equipped for transmission in the opposite direction. Syn: Full Duplex, Duplex.

UNATTENDED OPERATION

The automatic features of a station's operation which permit the transmission and reception of messages on an unattended basis.

UNIBUS

The single, asynchronous, high-speed bus structure shared by the PDP-11 processor, its memory, and all of its peripherals.

UNIT ELEMENT

A signal element of one unit element duration.

UNIT INTERVAL

A unit interval is the duration of the shortest nominal signal element. It is the longest interval of time such that the nominal durations of the signal elements in a synchronous system or the start and information elements in a start-stop system are whole multiples of this interval. The duration of the unit interval (in seconds) is the reciprocal of the telegraph speed expressed in Baud.

USASCII

See ASCII

USER EXIT

An external reference to which a user program may be linked, e.g., a user procedure for processing RJE batch job output.

VALUE ADDED SERVICE

A communication service utilizing communications common carrier networks for transmission and providing added data services with separate additional equipment. Such added service features may be store and forward message switching, terminal interfacing, and host interfacing.

VERTICAL REDUNDANCY CHECK (VRC)
A check or parity bit added to each character in a message such that the number of bits in each character, including the parity bit, is odd (odd parity) or even (even parity).

VIRTUAL CIRCUIT
A connection between a source and a sink in a network that may be realized by different circuit configurations during transmission of a message. Syn: Logical Circuit.

VOICE-GRADE CHANNEL
A channel used for speech transmission usually with an audio frequency range of 300-to-3400 Hertz. It is also used for transmission of analog and digital data. Up to 10,000 bits per second can be transmitted on a voice-grade channel.

WACK (WAIT BEFORE TRANSMITTING POSITIVE ACKNOWLEDGMENT)
In Binary Synchronous Communications, this DLE sequence is sent by a receiving station to indicate that it is temporarily not ready to receive.

WATS
Wide Area Telephone Service. A service provided by telephone companies in the United States that permits a customer to make calls to or from telephones in specific zones for a flat monthly charge. The monthly charges are based on size of the zone instead of number of calls. WATS may be used on a measured-time or full-time basis.

WIDEBAND
Communications channel having a bandwidth greater than a voice-grade channel characterized by data transmission speed of 10,000-to-500,000 bits per second.

digital

DIGITAL EQUIPMENT CORPORATION, Maynard, Massachusetts, Telephone: (617) 897-5111 • ARIZONA, Phoenix • CALIFORNIA, Sunnyvale, Santa Ana, Los Angeles, Oakland, San Diego and San Francisco (Mountain View) • COLORADO, Englewood • CONNECTICUT, Meriden, Fairfield • DISTRICT OF COLUMBIA, Washington (Lanham, Md.) • FLORIDA, Orlando • GEORGIA, Atlanta • ILLINOIS, Northbrook • INDIANA, Indianapolis • LOUISIANA, Metairie • MASSACHUSETTS, Marlborough and Waltham • MICHIGAN, Ann Arbor and Detroit (Southfield) • MINNESOTA, Minneapolis • MISSOURI, Kansas City and Maryland Heights • NEW JERSEY, Fairfield, Metuchen and Princeton • NEW MEXICO, Albuquerque • NEW YORK, Huntington Station, Manhattan, New York, Syracuse and Rochester • NORTH CAROLINA, Durham/Chapel Hill • OHIO, Cleveland, Dayton and Euclid • OKLAHOMA, Tulsa • OREGON, Portland • PENNSYLVANIA, Bluebell and Pittsburgh • TENNESSEE, Knoxville • TEXAS, Dallas and Houston • UTAH, Salt Lake City • WASHINGTON, Bellevue • WISCONSIN, Milwaukee • ARGENTINA, Buenos Aires • AUSTRALIA, Adelaide, Brisbane, Canberra, Melbourne, Perth and Sydney • AUSTRIA, Vienna • BELGIUM, Brussels • BRAZIL, Rio de Janeiro, Sao Paulo and Porto Alegre • CANADA, Calgary, Montreal, Ottawa, Toronto and Vancouver • CHILE, Santiago • DENMARK, Copenhagen • FINLAND, Helsinki • FRANCE, Grenoble and Paris • GERMANY, Berlin, Cologne, Hannover, Frankfurt, Munich and Stuttgart • INDIA, Bombay • ISRAEL, Tel Aviv • ITALY, Milan and Turin • JAPAN, Tokyo • MEXICO, Mexico City • NETHERLANDS, The Hague • NEW ZEALAND, Auckland • NORWAY, Oslo • PHILIPPINES, Manila • PUERTO RICO, Santurce • SPAIN, Barcelona and Madrid • SWEDEN, Stockholm • SWITZERLAND, Geneva and Zurich • UNITED KINGDOM, Birmingham, Bristol, Edinburgh, London, Manchester and Reading • VENEZUELA, Caracas

6802 • 11274 • 4582
ES • 09 • 45

Printed in U.S.A.