ON-LINE COMMUNICATIONS AND THE COMPUTER

By

F.K. Morioka
and
R.M. Wainwright

Control Data Corporation
MDM Communications Division
3519 West Warner Avenue
Santa Ana, California 92704

Presented At The
ACM Symposium
On Problems In The Optimization Of
Data Communications Systems

Pine Mountain, Georgia

October 16, 1969

ON-LINE COMMUNICATIONS AND THE COMPUTER

INTRODUCTION

This paper contains information on the general characteristics of the Control Data Corporation, MDM Communications Division M1000 Communications Message Switching System.

It is the intent of this paper to discuss the analytic systems approach to the problem that led to the building of unique hardware and software to resolve the message handling functions. The M1000 hardware and software system features are discussed in the paper.

The M1000 Communications Switching System allows front-ending, stand-alone message switching and remote batch entry.

ON-LINE COMMUNICATIONS AND THE COMPUTER

PROBLEM ANALYSIS AND SYSTEM DEFINITIONS

The definition of the scope of a data communications requirement is one of the most difficult tasks a system designer is faced with. If the budget were an open check book, the effort would be simple. However, we all realize that this utopian event is rarely available. One must always fit the task to particular budgets.

Experience has shown that prior to any pricing activity one must carefully analyze the communications requirements and functionally define each element to arrive at a workable configuration. During this period, the following questions should be answered:

1. How to design and build a computerized communications system for the communications environment.

2. How to design the system with the requirement of open ended growth from both the hardware and software viewpoints.

3. How to design and build a system that is efficient within itself in order that in the final result the system is cost competitive.

4. Source data form and characteristics:
    a. Media (i.e., magnetic tape, cards, paper tape, etc.)
    b. Character code set
    c. Volume per unit time, both peak and average.

5. Traffic distribution:
    a. Routing (single, multiple, or group addressing required?)
    b. Sorting (allowable destinations as a function of data origin).

6. Response requirements:
    a. Turnaround time, question/answer time

b. History requirements, or how far back must one access file or perform an update.

7. Throughput:

a. Peak requirement, usually a short time interval

b. Typical average day or month.

8. Communications reliability:

a. Tolerance to human error

b. Tolerance to equipment failure or a function of availability desired.

(1) Seven-day week, 24-hour day, without down time

(2) Eight-hour day, without down time

(3) Tolerance other than full eight-hour day.

9. Expansion requirements:

a. Functions, such as the servicing of new terminals, message control, and message processing requirements.

b. Capacity, the ability to add the new terminals of software features while maintaining the necessary throughput and response times.

If one examines some of the above points in detail, the following characteristics should evolve:

1. Source data form and characteristics.

If a wide variety of terminal devices are used, code sets, data transmission speeds, device control, and error handling requirements will require:

a. Extensive control software in any computer to avoid system crashes with attendant down times.

b. Memory space to meet required device response times and processing transmission speed buffering.

c. Various grades of communications channels to meet the various data transmission speed requirements.

d. Secondary storage space to provide the necessary space for queuing or protecting peak character loads.

2. Traffic distribution.

In an on-line communications environment the distribution of traffic will vary by the hour or day. Particular terminals may generate peaks at different times and are generally affected by daily business patterns. The end design must be able to adapt to these conditions by reallocating space and time on-line.

Typical control techniques which may be applied are:

a. Secondary or alternate routing paths
b. Reallocation of channel time by controlling terminal time on/off on a given channel
c. Controlling deliveries by priorities.

All of these requirements are generally controlled by software. However, in some cases, a combination of hardware and software — if applied to the proper problem — can and will solve the problem.

3. Response time requirements.

The present pace of industry dictates ready, accurate answers, whether one is solving scientific problems or inquiring to a data base regarding a particular credit card. Response time or turnaround time becomes an important system criteria.

Ready response time will require:

a. High data transmission rates, hence higher cost communications channels.
b. Fast internal processing so that the file process time is significantly less than the transmission line time. As an example, if the communications channel data rate is 300 characters per second, a 30-character

inquiry only requires 0.1 second to transmit. If an answer or response is 300 characters in length, the return transmit time is 1 second. If the file processing time is 1 second, the total turnaround time is:

T_{IN} = 0.1 second
T_{FP} = 1.0 second
T_{OUT} = 1.0 second (Display time assumed to be
2.1 second overlapped with transmission time and equal to the transmission.)

If the data transmission rate were doubled to 600 characters per second (4800 bits per second with 8 bits per character), the total turnaround time would be reduced to 1.6 second. However, note that the time consumed is the file processing time, not the communication time.

Hence, single-station, fast response terminals are not efficient in terms of communications. The general practice on these types of terminals is to concentrate N terminals on one line to derive the necessary response time while maintaining low line costs.

If one were to set up a credit checking facility, file processing times and communications channel loading need careful attention to avoid high line costs. Large files will generally require more time to process. Hence, the designer must consider the file process time when balancing channel loads.

4. Throughput.

Response time and throughput are generally discussed under one heading — Response Time. A low throughput system will increase response time.

If one is conversing with a data base, file process time and channel loading are essential considerations.

If one is concerned with transmit time between point of origin and destination, queuing delays must be considered. The designer has the choice of queuing at a central point or causing outstation queuing during peak conditions. The second choice will allow the use of smaller secondary storage units. However, if sufficient input and output channel capability is provided, the system queues are minimized, although channel costs may be high.

To ensure that the system will not be identified as a message delay point, due to large queuing peak and average volumes, per unit time should be estimated under normal as well as abnormal conditions. The typical underestimation is the condition where a destination data sink fails. A large system queue builds up and bogs down the whole system. The destination sink is fixed. Data is dumped. The sink fails. The queue grows. Somebody makes a decision to airmail the data.

Experience has shown that the system throughput requirement is always underestimated. The basic reason is that as the system provides better service, applicational loads will rise unexpectedly.

5. Communications reliability.

The degree of service desired by any on-line system depends on the particular business demands. As an example, a brokerage firm may have short operating hours, but any system failure will seriously affect the buy and sell volume per day.

An airline is concerned with proper deliveries of reservations information, operations messages, maintenance data and schedules, on an around the clock basis.

ON-LINE COMMUNICATIONS AND THE COMPUTER

An automated credit bureau is concerned with ready reply during business hours.

A time-sharing facility must provide reliable service on demand.

The common carriers must provide round the clock service without causing undue delays while continuously expanding the network.

The necessary reliability from an on-line data communications system can be derived if some of the following considerations are included in the design analysis:

a. Is the software capable of handling a wide range of error capability?
   - (1) Human error - operators make mistakes. What action should the system take?
   - (2) Machine error - digital computers are not completely digital. Core memory, disc files, and magnetic tape storage units involve analog to digital signal conversion. An error rate does exist. The software must recognize this fact and include provisions for error protection. Also, provisions should be included for subsystem failure detection and recovery to avoid total failure due to subsystem failure or power loss.

b. Is the hardware designed for a communications environment?

   Intermittent outages can be avoided if the hardware and/or its power source is designed to withstand short power line interrupts.

   Communications interface equipment must be designed to be compatible with common carrier practices to avoid time consuming channel outages.

c. Is the communication controller software designed to readily accept new features?

Any addition of programmed features generally requires careful integration and listing.

The software should be able to readily accept a new requirement in such a manner that the total system is not disabled during the integration interval.

d. Is the system reliability compromised as the system grows?

The prediction of growth for most on-line systems has been seriously underestimated. Data processing needs always increase and unfortunately the need to communicate also increases.

As pointed out earlier, it is sometimes desirable to treat the needs independently so that each may grow without hampering the other's growth.

Whatever choice is made, serious thought should be directed towards what effect any expansion requirement will have on overall reliability. As an example, if N accesses to secondary storage are required to meet message protection needs for a given data or message rate it is quite obvious that the remaining accesses required for data processing will not be available when the data rate increases. The converse arguement is also true.

6. Expansion requirements.

The expansion of an on-line data communication system generally occurs in two areas:

a. Addition of new terminal types and corresponding communications channels.

b. Addition of special software such as forms control.

Both cannot be clearly separated functionally as forms handling will also affect the channel control procedures (such as calling all receive-only stations, confirming they are there, and then sending the message).

Any feature additions will affect throughput as more computer instructions per message will be required. Also, space must be allocated for the software. If more accesses are required to secondary space, the effect on throughput will be higher.

Initial planning must include provisions for features. Programming and hardware cost trade-offs must be made.

Although the consideration list is small, one will find that the answer derivation process must be accomplished in an orderly, well-specified manner to avoid poor data service.

With this list of considerations in mind, one can see what a system design must focus its attention towards if, in fact, the system to be designed is going to solve the problems facing the front-end, stand-alone message switching and remote batch entry users of today's growing marketplace.

WHAT THE MDM COMMUNICATIONS M1000 COMPUTER SWITCHING SYSTEM IS AND WHAT IT WAS DESIGNED AND BUILT TO DO

The M1000 System is a complete, self-contained, pre-programmed data communication system. No other systems, equipments or programs are required in order to perform the various communication tasks. Messages are received, edited, routed, translated to other codes as required, delivered or intercepted and stored for later retrieval or processing entirely within the confines of the M1000 System.

ON-LINE COMMUNICATIONS AND THE COMPUTER

Messages are switched within the system in a store-and-forward method for the purpose of stand-alone message switching. Message or job entries are collected, switched, sorted and queued within multi-job queue areas, within the M1000 System, in a store-and-forward method for the purpose of front-ending, or pre-processing for data record updating and inquiry/response applications related to a batch processing system. It should be realized that, within the M1000 System, not one application is done within the system, but rather a multiple of applications can be handled within the same hardware and software system. The three major applications that the M1000 System is designed to accommodate are listed as follows:

1. Stand-alone message switching for the purpose of user-to-user communications message switching.

2. Information exchanges/EDP front-ending for the purpose of collecting, switching, sorting and queuing inquiry/response applications and data record update information related to large-scale data batch process systems.

3. Remote, batch entry applications.

An M1000 System is comprised of one or more data communications centers called Exchanges. Each Exchange provides communication service for the lines and devices which it terminates. When a system consists of more than one Exchange, the Exchanges are interconnected by high-speed trunks providing full duplex inter-exchange communication at transmission rates of 9,600 bits/second in each direction.

An Exchange is a complex of one or more identical Exchange Units. In a Multi-Exchange Unit Exchange, each Exchange Unit terminates a portion of the lines, shares the total traffic load and provides redundancy for protection against equipment failures. Communication between Exchange Units is provided by a disc trunking process which optimizes intra-exchange communication while maintaining complete traffic protection.

An Exchange Unit consists of a special purpose communications processor, core and fast access disc memory, line termination equipment and auxiliary memory. The Exchange Unit is modular with respect to these components (except the processor which is a fixed component of each Exchange Unit), and the kinds and amounts employed are dependent on the application.

Two different system organizations can be derived using the M1000 System components, a centralized organization or a distributed organization. The selection of a particular system organization depends primarily upon (1) the geographical distribution of the communications network and related line costs, (2) bulk data transmission requirements, and (3) inter-communications requirements between such high-speed terminals as data processing centers.

The centralized data communications system consists of a single Exchange terminating all of the lines and terminal devices in the communication network. Capability is provided for data communication between all terminal devices in the network. Additionally, the Exchange may terminate several processing computers providing front-ending for the EDP computer systems.

A distributed system consists of two or more geographically separated Exchanges, each of which terminates a portion of the lines and devices in the total network. Each Exchange in the distributed system has the same general communications capability as previously described for the centralized system.

Additionally, high-speed, as well as low-speed devices can now be terminated in two or more locations and can inter-communicate over voice grade trunks at the high data rate.

When high-speed devices are terminated at two Exchanges, bulk data transfer can take place simultaneously with the transfer of messages to and from the low-speed lines. Low-speed line traffic usually has first priority on the inter-exchange trunks, and the bulk data transfer rate is reduced, as required, by the low-speed traffic volume. In some cases, the inter-exchange trunks will be lightly loaded with low-speed message traffic, and, therefore, most of their capacity is available for bulk data. In other cases, the trunk is reserved for higher priority traffic during peak hours and the lower priority traffic is accommodated during slack hours.

Each Exchange in an M1000 System will normally control the lines connected to it. It will issue polls and calls as required, receive and evaluate responses and monitor the lines for error conditions. Proper handling and routing of all messages is accomplished by the M1000 System using specific system parameters entered via the console. The M1000 can also terminate single-station, uncontrolled inputs.

Messages received from any line can be routed (switched) to any station(s) in the network if properly addressed. If any station or stations direct messages to just one fixed destination, the incoming messages from each such station can be directed via implied routing; i.e., the input station implies the message destination and the messages need not be addressed. Switched messages with improperly formatted addresses will be intercepted by the M1000 System and delivered to the intercept station designated for the input station generating the message. The texts of all messages received are delivered to the designated (or implied) outstation(s) without change except for code translation where required.

A record of all traffic through each Exchange Unit is maintained in the Exchange Unit's internal storage to the limit of its available capacity. Messages so stored may be automatically retrieved from storage by request from any of the outstations to which they were originally directed, or from the Exchange Control Console. If desired, this retrieval capability is expanded by the use of extended retrieval on disc packs and journal records on magnetic tapes.

ON-LINE COMMUNICATIONS AND THE COMPUTER

In some system applications a de-centralized switching system (central M1000 tying one or more remote M1000 Systems into one total, integrated M1000 System) has been required. The means by which the subsystems can be tied together, if desired, is a MDM Communications 9600-bit-per-second modem which utilizes a Tariff FCC #260 Type 3005 (Schedule 4, Type C4 Facility or its equivalent).

Inter-exchange traffic over the 9600-bit-per-second trunk is highly protected by error coding and by program/hardware interlocks and controls associated with the trunk. The data rate on the trunk is automatically adjusted, as required, to take advantage of the highest rate the trunk telephone line will support at any given time. Retransmission requests are issued automatically for blocks of data received in error. Although normally operating with one particular Exchange Unit at each of its terminals, the trunk will automatically switch to another Exchange Unit (if available) at either or both ends should its normal terminal Exchange Unit(s) fail.

An M1000 System may interface with a foreign communication system or data processing computer. In these cases all transfers of information are initiated by the foreign communication system or data processing computer.

When interfacing with a foreign communication system, the M1000 System will appear as an outstation in the foreign communication system network, responding as appropriate to polls and calls in accordance with the requirements of the foreign system outstation logic.

When interfacing with a data processing computer, the M1000 System will normally appear as a device on the I/O channel in the data processing computer system.

MDM Communications has developed a Local Terminal Interface Package (LTIP) which connects the IBM System 360 computer with the M1000. The LTIP-360 includes capabilities to connect up to eight S/360 I/O selector channel device ports to eight

M1000 Exchange Units each operating independently with overlapping of various channel communications. The interface package consists of (see Figure 1):

1. TIP Program.

   The program residing within the M1000 Exchange required to support the data processing computer interface. The TIP program passes or accepts data and control information to and from the D. P. computer system.

2. 360 Interconnect (360 IC).

   The 360 IC is a switching and holding device to logically connect an addressable 360 channel device port to a uniquely identified Exchange Unit. Connection is made on a first come, first serve basis.

3. M1000 Interface Control Unit (MICU).

   The MICU acts as the translator between the data processing computer and the M1000. The MICU performs signal analysis and provides device status to the D. P. computer.

4. M1000 Interface Program (MIP).

   MIP consists of a set of macros residing in the D. P. computer to control I/O operations with the M1000 System. The MIP passes or accepts data and control information to and from the M1000.

Messages across the computer interface are blocked to a maximum of 4096 eight-bit bytes. The M1000 assumes the responsibility of deblocking messages written by the D. P. computer; the user's D. P. program must perform the deblocking function on a read operation. All data transfers are at a minimum of 60,000 bytes per second to a maximum of 120,000 bytes per second and are parity checked by byte.

All operations across the computer interface are initiated by the data processing computer. On D. P. write operations, the M1000 deblocks the messages and passes them to the appropriate queues and processes. For each D. P. read, the M1000 passes a single message or block of messages to the D. P. computer.

The hardware and software of the M1000 System was designed to provide continuity of service in the presence of equipment and communication line failures. Techniques for coping with failures in each of the major portions of the system are discussed separately.

The connections provided for backup of one Exchange Unit by an alternate are shown in Figure 2. Connections to the lines and devices for which an Exchange Unit has primary responsibility are shown as vertical lines while the secondary connections are diagonal lines.

Each communication processor is connected to one bank of line modules for which it has prime responsibility and to another bank for which it has secondary responsibility. A two-phase clock provides the means for sharing access to a single bank of modules by the alternate and prime processor.

Each communication processor may have associated with it a transmitter/receiver for inter-exchange traffic on a voice bandwidth trunk. A transmitter/receiver switch, under control of the monitor unit, connects the outgoing trunk line either to the transmitter/receiver in the Exchange Unit having prime responsibility for the trunk, or to the transmitter/receiver in the adjacent Exchange Unit having secondary responsibility for the trunk.

Through the disc access controller, each Exchange Unit in the Exchange has access to each disc file in the Exchange. This provides for intra-exchange traffic trunking and for redundant recording of traffic for backup.

Through the disc pack access units, each Exchange Unit in the Exchange has access to each disc pack drive in the Exchange just as with each disc file.

Figure 3 illustrates the logical organization of a Multi-Exchange Unit Exchange. An understanding of this organization and of how Exchange Unit functions are redundantly distributed through an Exchange is basic to an understanding of the service continuity and traffic protection features of the M1000. Subsequent paragraphs will describe this logical organization and will relate it to the physical organization of the Exchange. It is of particular importance that the differences between the logical and physical organizations be clearly understood.

The set of processes normally associated with an Exchange Unit in conjunction with their associated tables and queues, comprise a control function. As shown in Figure 3, each control function has an associated logical disc system and an associated line group. Thus, for example, Control Function N is associated with Line Group N and with Disc System N. As indicated in the figure, only Control Function N can both transmit (or write) to and receive (or read) from Line Group N or Disc System N. However, Control Function N+1 can receive from the non-controlled lines of Line Group N and any control function can read from any disc via the disc trunk.

Figure 4 shows the relationship between the logical and physical organization of an Exchange Unit. Each physical Exchange Unit (the items arrayed vertically under the heading "Exchange Unit" on the figure) consists of three major elements as shown: a line group, a communication processor, and a physical disc unit. The logical disc system associated with a control function consists of two elements, one in each of two adjacent disc units. As shown in Figure 4, the disc system for Control Function N consists of a portion of Disc N and a portion of Disc N+1. These two portions are used to store identical information providing each control function with a redundant copy of the information it requires to service its line group. Since the two sets of information are identical, either Disc N or Disc N+1 can fail without interrupting the servicing of Line Group N. In a like manner, any disc in the Exchange can fail without interrupting the service of any line group.

As further illustrated in the figure, a particular control function may be activated in either of two physical Exchange Units or processors. Thus, for example, CFN may be active in either Processor N or in Processor N+1. Normally, it will be active only in Processor N as indicated by the solid enclosure in that processor in the figure. However, in the event Processor N is out of service for any reason, CFN will be activated in Processor N+1 as indicated by the dotted enclosure in the figure. (For certain purposes, to be discussed in a later paragraphs, CFN may be locked out of either or both processors under manual control.)

Each control function includes a set of tables and queues which provide all the information required by that control function to control and store the traffic of its associated line group. A control function is aware of only one set of tables and queues but duplicate sets are maintained under hardware control. These duplicate sets are maintained in each of the two elements of the control function's disc system previously described, the prime and copy areas. Control Function N, for example, services its lines from a single logical set of tables and queues. The physical set(s) actually employed are established under hardware control. Normally, anything logically written by CFN will be physically written, by hardware, to both the prime and copy areas of CFN and anything logically read by CFN will by physically supplied from the first of CFN's areas available. If physical Disc N should fail, the copy area in Disc Unit N+1 only will be used. If physical Disc N+1 should fail, the prime area in Disc Unit N only will be used. CFN operates identically in all cases, unaware of whether one disc or both discs are being used.

Exchange Unit failures can occur in three general ways: line group element failure, processor failure or disc failure. Failure of an element in a line group causes interruption of service to the line associated with that element only. Service to all other lines in the system continues uninterrupted. Failure of a processor or of a disc in an Exchange containing two or more Exchange Units produces no interruption of service due to the manner in which Exchange Unit functions are redundantly distributed as previously described. The new logical organizations to which this Exchange automatically reconfigures itself in the event of these failures are described below.

ON-LINE COMMUNICATIONS AND THE COMPUTER

Figure 5 shows the logical organization of an Exchange when a processor has failed (Processor N). At the time of failure, Control Function N, then active in Processor N, went out of service with the processor. However, CFN is subsequently activated in Processor N+1 and continues servicing its line group using its own disc system as indicated in Figure 5. As indicated in the figure, Control Functions N-1 and N+1 are still operating normally after the failure. Control Function N, however, while servicing its own line group is taking no action on the non-controlled lines of CFN-1. Thus, in this situation, although all lines are still being serviced, there is no longer any backup for Line Group N-1 until Processor N is returned to service. However, there is no service interruption without at least one additional processor failure. If Processor N-1 now fails, CFN-1 fails with it and cannot be activated in Processor N, already failed. Thus, Line Group N-1 is no longer serviced although service continues on Line Groups N and N+1. If Processor N+1 should fail instead of Processor N-1, Control Function N fails with it and cannot be activated in Processor N, already failed. Thus, Line Group N is no longer serviced. However, service of Line Group N+1 continues (as does service of Line Group N-1 in this case) as CFN+1 is activated and restarts in Processor N+2.

Figure 6 shows the logical organization of an Exchange when a disc has failed (Disc N). At the time of failure, the disc access control hardware discontinues use of Disc N. Disc N-1 only is now used for CFN-1 and Disc N+1 only is now used for CFN as indicated in the figure. The control functions, as previously noted, are unaware of this change in operation and take no action. As with processor failure, system operation continues without service interruption until at least one additional disc failure.

As noted above, system failure (absence of service to one or more line groups) cannot occur without multiple processor or disc failures. Implicit in this discussion was the fact that the multiple failures must be of adjacent processors or adjacent discs. In a many-Exchange Unit Exchange, every other processor and every other disc can fail without service interruption to any line group. The following discussion will further substantiate this by describing the effects of various combinations of processor and disc failures.

1. Processor N and Disc N Both Fail.

   As shown in Figure 7, CFN is activated in Processor N+1 and continues servicing Line Group N using that half of its logical disc in physical Disc Unit N+1. Line Groups N-1 and N+1 continue to be serviced as before, except there is no backup for Line Group N-1 and no copy for the CFN-1 prime area of Disc N-1.

2. Processor N and Disc N-1 Both Fail.

   Figure 8 illustrates this situation. All line groups continue to be serviced and either a prime or a copy area or both is available for all line groups.

3. Processor N-1 and Disc N Both Fail.

   Figure 9 illustrates this situation. Again, all line groups are serviced and all disc information is available.

As previously noted, service interruptions will result from the failure of two adjacent processors or two adjacent discs. Each line module in the system is equipped to automatically send "emergency stop" signals down outgoing lines whenever failure of two adjacent units is detected by the monitor units. Outstation transmitting equipments, equipped with emergency stop recognition devices, are inhibited from blindly sending further traffic during this catastrophic failure condition. When either Exchange Unit recovers from the failed condition, emergency stop is removed and service continuity is restored.

Input line or station failures are automatically detected and reported by the Exchange Unit at the Exchange Operator position. The Exchange Unit will continue to attempt to service the failed line or station (the failed condition is only reported once) thus restoring service continuity at such time as the outage is remedied.

The particular manner in which a failure manifests itself is dependent upon the particular kind of line and remote station equipment. Although provisions are made for terminating non-controlled lines (freewheeling), controlled lines should be used wherever practical since fault conditions on the line or within the remote station equipment are more readily determined.

M1000 SYSTEM HARDWARE

The M1000 Data Communication System consists of integrated sets of hardware elements designed to permit simplified implementation of data communications systems through standardization of electrical and mechanical interfaces. A wide range of communication systems can be constructed, maintained and expanded without causing redesign of the basic hardware set.

For a particular application, equipment and stored program modules are assembled to form an Exchange Unit that performs the basic data communications function. A basic Exchange Unit may be expanded by addition of memory units and line termination hardware. To provide additional capacity and continuity of service through load sharing, additional Exchange Units are added to form an Exchange which provides the data communication services at one site serving a geographical area. Exchanges, located in the different geographical areas that they service, may be interconnected by inter-exchange high-speed trunks for message service at 9600 bits per second over voice grade telephone circuits. The following paragraphs describe the equipment organization of an Exchange and Exchange Unit.

ON-LINE COMMUNICATIONS AND THE COMPUTER

1. Exchange Organization.

An M1000 Exchange consists of one or more Exchange Units located at the same geographic site. Figures 2 and 10 illustrate the interconnections at the Exchange level.

Basic communication between Exchange Units within an Exchange is by means of the fast access disc system. The process is accomplished by allowing any communications processor within an Exchange to read from any fast disc storage unit within an Exchange. Communication from Exchange Unit A to Exchange Unit B, for example, is effected by having Exchange Unit A write certain data on its fast disc storage unit and then having Exchange Unit B read that data from the disc of Exchange Unit A.

Service of the low-speed data handling devices attached to a system is by means of the line modules with appropriate signal level converters. These line modules are scanned for data by the line controller. In order to allow line service to continue when an Exchange Unit is down, a group of line modules are scanned by both the line controller of the Exchange Unit in which they are located and, on command, by the line controller in the adjacent Exchange Unit. Similar switching is also provided for the transmitter-receiver and the LTIP 360 units. This technique of allowing switch-over of various devices from one Exchange Unit to its neighbor allows system units to be backed up while at the same time provides for ease of expansion.

2. Exchange Unit.

The Exchange Unit is the basic self-contained subsystem from which M1000 data communications systems are configured. Each Exchange Unit can be expanded, as is shown in Figure 11. A special purpose communications processor controls all message switching and processing for the Exchange. Input and output functions are accomplished by means of the controllers shown. The line controller scans up to 128 line modules within the Exchange in addition to the monitor unit, keyboard entry/readout and monitor console. Although not shown in Figure 11, the line controller additionally scans up to 128 line modules in an adjacent Exchange Unit for backup.

All of the equipments of the Exchange Unit are modular and are described in the following paragraphs with the exception of the console.

a. Communications Processor.

The communications processor is an integrated circuit, fully parallel central processing unit which supplies the Exchange Unit with basic stored-program control and computational capabilities. It contains the facilities for addressing up to 65,536 twenty-four-bit words of core storage for fetching and storing instructions and data. It provides control for sequencing and executing instructions. It has capabilities for arithmetic and logical processing of data, and it initiates and controls communication between input-output devices and memory.

Integrated circuitry is used throughout the communications processor for reliability and low power consumption. A word length of 24 bits is used in the fully word parallel organization that provides fast arithmetic and logic operations. The ten hardware registers including three 16-bit index-base registers in conjunction with a special set of communications oriented instructions provide the processor with the special processing capability required for data communications.

Direct memory access is provided for up to twelve input-output device controllers, each operating at word transfer rates up to 200 KC simultaneously with processing. A four-level priority interrupt system is provided as part of the communications processor. Several devices can cause interrupts at the same priority level. When a device supplies an interrupt request, the processor completes execution of the present instruction and then services the interrupt by exchanging the contents of the program counter register with a memory cell associated with the priority level at which the interrupt occurs. This causes the program to branch to an interrupt subroutine, which tests the devices attached to that priority level to ascertain which device supplied the interrupt request. The program then executes the proper subroutine associated with that device.

The core storage units utilized are conventional magnetic core memories containing 4096 twenty-four-bit words of storage. System addressability is sufficient to address up to 16 core storage units for a maximum capacity of 65,536 words. The core storage units are highly reliable, conservatively designed, two-microsecond units, operated at 2.4 microseconds in order to further enhance operating margins.

b. Fast Access Disc System (FDS).

Each Exchange Unit is equipped with a fast access disc file for non-volatile storage of programs and messages. The disc controller can control up to eight disc units for each Exchange Unit.

ON-LINE COMMUNICATIONS AND THE COMPUTER

Communication between Exchange Units in an Exchange for trunking of message traffic and for traffic protection backup is provided through the disc access system. This system, implemented by the disc access controller, allows any Exchange Unit to transfer data to and from any disc file in the Exchange.

The operation of the disc access controller centers around a sequence of processor calls of discs, followed by disc calls of processors. The operation is as follows:

(1) A Communications Processor (CP) places a call for a certain data area on a specific disc file.

(2) An addressed file accepts the call by storing the identification of the requesting processor and the requested data area. If the disc storage unit is busy, the requesting processor is so notified, and it must try again.

(3) As the access period for the requested data area approaches, the disc file places a call to the requesting processor.

(4) The processor then writes to or reads from that data area on the disc file.

(5) The processor reads the data through the disc controller.

In addition to the communication link provided by the disc access controller which allows any processor to access any disc, a direct read-write communication is provided between a processor and its own disc storage unit. Also, in order to provide backup capability, another direct read-write communication is provided between a processor and the adjacent disc storage unit.

Expansion of an Exchange by addition of Exchange Units is an orderly and efficient process. Since Exchange Unit to Exchange Unit communication is accomplished through the disc access controller, expansion by addition of Exchange Units involves the attachment of the new Exchange Unit to the open-ended disc access controller trunk.

The disc files are of conventional fixed-head design. For most applications, a file with twelve million bits of storage and 17.3 microseconds average access time is used. All heads "float" above the disc platters once the disc is started to ensure maximum reliability.

The disc unit is organized into 512 data tracks and two timing tracks. Each track contains 64 sectors. Each sector is divided into a 429-bit data and parity area and a 47-bit guard area. Words may be stored on disc sequentially on one track in straight bit serial fashion, on two tracks as two bit "characters" in bit parallel/character serial fashion or on four tracks as four bit "characters" in bit parallel/character serial fashion. In each case, every twelve sequential data bits on a track are followed by an associated parity bit and every 32 such sets of 12 data bits and 1 parity bit are followed by an associated 13-bit vertical parity group. The selection of 1, 2 or 4 bit "characters" for fast disc storage is made on a system/application basis and provides for varying data transfer rates to and from disc.

c. Disc Pack Storage System (DPS).

The M1000 Disc Pack Storage (DPS) facility provides economical extension of the storage space available for queuing of traffic awaiting delivery and for storage of delivered traffic to provide fast access for subsequent retrievals. Additionally, DPS may be used for initial loading of programs into the Exchange Unit's fast access, fixed head disc storage system.

The equipment items used for DPS are a Disc Pack Controller, Disc Pack Access Unit, Disc Pack Drive Unit and the Disc Pack. As shown in Figure 10, one Disc Pack Controller and one Disc Pack Access Unit are associated with each Exchange Unit. Up to eight Disc Pack Drive Units can be accommodated on each access unit. A disc bus arrangement is provided such that any controller can access any drive through an access unit.

The Disc Pack used on the drive is a removable unit consisting of six 14-inch discs, mounted 1/2-inch apart on a central hub. Data is recorded on the inside ten disc surfaces by ten read/write heads mounted on a movable comb-like access mechanism. In operation, the access mechanism is moved horizontally to any one of 203 positions thus providing 203 data tracks on each disc surface or a total of 2030 tracks. When the drive is not in operation, the access mechanism is fully retracted to allow replacement of the disc pack.

d. Magnetic Tape System.

The M1000 Digital Magnetic Tape System provides the system with high-speed digital magnetic tape input-output capability. The magnetic tape system consists of digital magnetic tape units attached to the M1000 System through the magnetic tape controller. After initial setup of the magnetic tape controller by a processor input-output instruction, input-output by means of magnetic tape occurs simultaneously with other processing activity.

M1000 magnetic tape units are normally supplied in 37-1/2 inches-per-second versions operating at 800-bits-per-inch packing density, completely compatible with IBM 9-channel System 360 tape systems. Up to eight magnetic tape units may be attached to each Exchange Unit through the magnetic tape controller.

Tape controller operation is initiated by the communications processor. Execution by it of a properly coded input-output instruction, signals the tape controller to execute a particular command. The tape controller then executes the command under its own control, freeing the processor for other work. Actual input-output occurs through the memory access controller (portion of the communication processor) directly to core storage on a cycle-stealing basis, under control of the tape controller. When the tape controller determines that the input-output process is complete for a given setup command, it notifies the processor by means of an interrupt.

e. Low Speed Input-Output System.

The low-speed input-output system provides communication with lines having character rates up to 400 characters per second. The system is primarily composed of line modules (one for each full duplex line) that perform serial to parallel conversion on data characters and a line controller that scans the line modules and provides data transfer between them and core memory.

The line controller can scan and control up to 256 line modules (128 located in the same Exchange Unit with the line controller and 128 from the adjacent Exchange Unit). The line scan speed per line module is under program control which allows complete flexibility in terminating lines of different speeds from 0 to 400 characters per second. The scan program for each application is configured to optimize the line sampling rate for each group of lines operating at different speeds to minimize the number of memory accesses.

The basic sampling clock for the line controller is supplied by the scan initiate clock. Since line controllers scan line modules in other Exchange Units, the scan initiate clocks each backup their neighbor. Thus, should failure of one of the clocks occur, the adjacent clock supplies the sample timing.

The line modules sample serial data from the line at the line rate, assemble characters and transfer them, character at a time, to memory through the line controller. Characters from memory are transferred to the line module through the line controller where they are transferred to the line module through the line controller where they are serially placed on the line at the line rate. Additionally, the line modules perform supervisory code detection and generation for such line supervision as end of message, start of message, open line, connect answerback, etc.

The line module to line interface is normally logic level RS-232 compatible. For teletype lines solid state keyers are provided for either 20 or 60 ma. neutral current loop keying. Optional keyers for other line interface characteristics are also available.

3. Transmitter/Receiver.

The MDM 9600-bit-per-second modem, or Transmitter/Receiver (TXR) is a MDM Communications development which allows adaptive communications over voice grade telephone circuits.

Once initiated by program control, the transmitter accesses memory and encodes the digital information in a manner suitable for transmission over telephone channels. The generated analog signal exhibits a very narrow signal spectra which may contain digital information for 1200, 2400, 4800, 7200 or 9600 bits per second. The particular rate transmitted over any particular period is primarily controlled by the communications processor and its resident TXR control program. The rate control parameters were derived at the distant M1000 and transmitted to the local M1000 in real-time.

The receiver decodes the incoming analog signal and directly loads memory on a cycle-stealing basis.

The receiver performs bit and frame synchronization detections and the necessary serial to parallel conversion for word transferring. Its detection or sampling rate is controlled by the local M1000 Communications Processor (CP). The TXR is completely microminiaturized and requires no operator intervention. Its operation is completely controlled by the respective M1000 CP's.

4. Trunk Switch.

The trunk switch is a simple relay switch which allows one of the two TXR's to be connected to a given telephone circuit. Its action is controlled by the Exchange monitor hardware.

5. Line Termination Equipment.

The line termination equipment basically consists of cabinets, level converters, power supplies and jack fields which allow compatible termination of the telephone circuits to the M1000 line modules.

6. Monitor Unit.

The monitor unit is a portion of a special line module (the console line module). It implements switchover in the event of Exchange Unit failure. The monitor unit performs three basic functions:

a. It monitors its own Exchange Unit for a special output every character interval.

b. It communicates with the monitor units on both the Exchange Unit that it is providing backup to and the Exchange Unit that is backing it up.

c. In the event of Exchange Unit failure, both Exchange Units are informed via their monitor units and switchover takes place as previously described.

For maximum reliability, the monitor unit is an ultra-conservative, worst case design with a minimum of carefully derated components. Exchange Unit status and alarms are provided from this unit to the console.

M1000 SYSTEM SOFTWARE

The programming portion of the MDM Communications M1000 System is organized in modular packages. In general, these modules are standard programs independent of specific applications. Special modules may be customized for special requirements. Through the use of appropriate controls and interlocks within and between modules, the programming system provides for such features as prevention of message loss or duplication, prevention of contention between messages for service or storage, expeditious handling of priority messages, message intercepts, alarms, and the like.

ON-LINE COMMUNICATIONS AND THE COMPUTER

The program modules act in conjunction with each other, with various hardware modules and, on occasion, with human operators to form a set of closed processes.

1. Programming System Storage.

   An important aspect of the system is the use of tape, core and disc storage and the manner in which these forms of storage are utilized relative to the programming system.

   Basically, magnetic tape or disc pack storage is used for "off-line" program storage and for initial system set-up. Two primary classes of off-line program storage are significant:

   a. Initial Exchange Unit Load Storage.

      A collection of all the program elements a user may employ, but without the definition of specific system configuration parameters. It is used to prepare configured Exchange Unit load storage.

   b. Configured Exchange Unit Load Storage.

      The user's on-line programming system. It is used to load an Exchange Unit as necessary for on-line operation.

   Fast access disc and core are used for on-line storage. Additionally, the disc retains the entire programming system during operation so that normal recovery may be accomplished without recourse to off-line program storage. During operation certain programs are permanently retained in core storage (core resident) due to timing requirements. Other programs are moved to core storage for execution only as required (overlay). Certain areas in core storage are set aside for overlay use.

2. Relationships With Total System Organization.

As previously noted, all programs are related to various closed processes. In some cases, the program itself handles all of the functions required to form the process and "program" and "process" become synonymous. In general, however, a process is comprised of functions performed by one or more programs, hardware modules and human operators. A specific example of the distinction between the organization of programs and the functional closed process organization of the total system may be found in inter-exchange trunking. Each trunk program is a part of the programming system in its respective Exchange Unit while the trunk process involves the cooperative efforts of two of these programs and certain hardware elements in order to properly control traffic between the two associated Exchanges.

Information is transferred between processes by means of queues. One process will form an output queue of information which becomes an input queue for the next process which is to assume control of this information. Except for external queues of messages at outstations, queues are manipulated by the programming system and, in general, considered a part of its organization in the following descriptions. As major interfaces between programs or processes, they constitute the major links connecting these elements into an organized system.

Programs are designed to be totally independent of the many possible variations in types and numbers of lines and devices that may be encountered by the system. All such variations are specified by means of tables. Table entries define specific system parameters and are made and revised by the user in accordance with his specific requirements.

The set of programs and tables which are concerned with the control and servicing of the group of lines and devices which are primarily assigned to a particualr Exchange Unit constitute a control function. The set of programs and tables which are concerned with the service of the resident console constitute a resident function.

3. System Operation.

These paragraphs will be concerned with the scheduling and control of the programming system during operation.

For operational purposes, each of the processes of the control or resident functions is divided into a set of program loops. Each program loop is composed of a set of program modules. A program module is a basic block of logic. A program module is established on the basis of its size, its independence and similar considerations. A program loop is the basic system logical element which is scheduled for execution. Once initiated, it is cycled through to completion. It halts temporarily at the end of each program module and defines the next module to be called up from the fast disc system for execution. When the loop is completed, it is removed from the schedule until it is once again initiated. A loop may be initiated on the basis of clock time or via an initiation of "GO" signal from another loop.

A process, as previously defined, is the complete set of logic (implemented, in general, in programs, hardware and human operators) required to move a message from one queue to another.

A process may contain one or more program loops. A program loop may contain one or more program modules. TIP (Terminal Interface Package) programs are examples of program loops. They are parts of both the input and the output processes which also include, for example, the priority one interrupt loop. Most TIP programs are overlay type. The priority one interrupt loop is core resident. The switch process, which is synonymous with the switch loop, is an example of a one-loop process. As a further example of operation, the switch loop is initiated ("Loop Go") by a signal indicating that a message to be switched is in queue. This "GO" signal, in turn, was issued by another loop (the common enqueuer loop) which queues all incoming traffic.

Figure 12 illustrates a program loop. The following paragraphs describe the manner in which operation sequences through the loop.

A "Loop Go" signal is activated for a particular loop by another loop or on a timed basis. This signal places the loop on the operations control schedule. Operations Control will pass control to this loop on a priority basis as long as it is on this schedule. Program Module 1 will be called up from disc to an overlay area in core for execution at the first such step.

As each module completes its functions, it issues a "Program Complete" signal as indicated on the figure, releases its overlay area in core and indicates the next program module to be executed in this loop when it is next scheduled by Operations Control.

As indicated for Program Module 2, in Figure 12, a module may relinquish control before it is completed. This can occur if it is interrupted by a higher priority loop or if it must wait for an input, output or service operation it has requested. In this case, the module remains in its overlay area in core awaiting return of control so that it may continue to completion.

When the last module completes, the "Loop Complete" signal is issued. This releases the overlay area occupied by the last module and removes the loop from the operations control schedule. Operations Control will not again schedule this loop for execution until its "Loop Go" signal is again issued.

The preceding description applies primarily to loops comprised of overlay type modules. Some loops, as previously noted, are core resident and never relinquish their storage areas in core. However, the above description applies in all respects other than relinquishing of core storage areas.

Core storage is so organized that modules of two different program loops may be resident simultaneously. That is, there are two overlay areas. Therefore, a loop program module may not only pass control to some core resident program before completion, but to a program module of another loop in the other overlay area.

As one can see, a multitude of areas must be considered if one is going to design and build a computerized switching system to perform the functions of stand-alone message switching, EDP front-ending and remote batch entry applications, yet have the ability to allow open-ended growth — in both hardware and software — and terminate a wide range of past, present and future user terminal devices with no major change to the basic hardware and software systems.

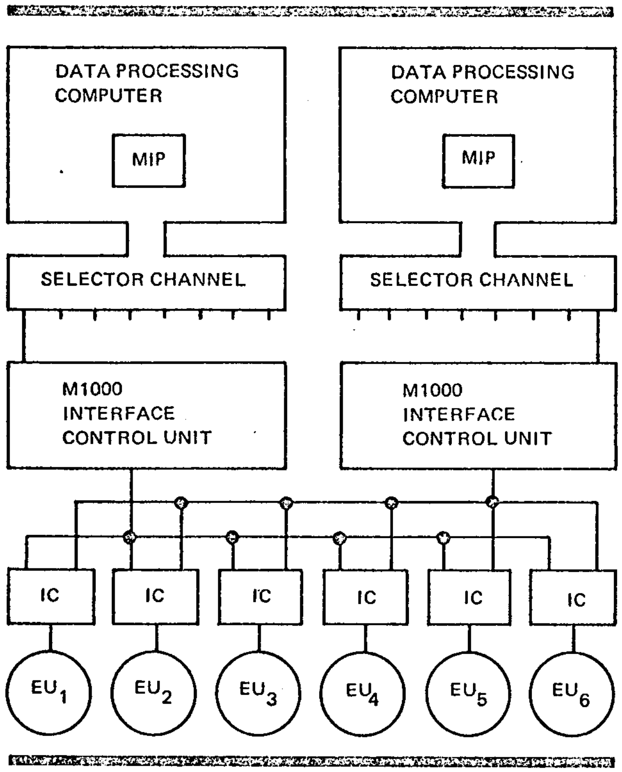

Figure 1 L TIP 360

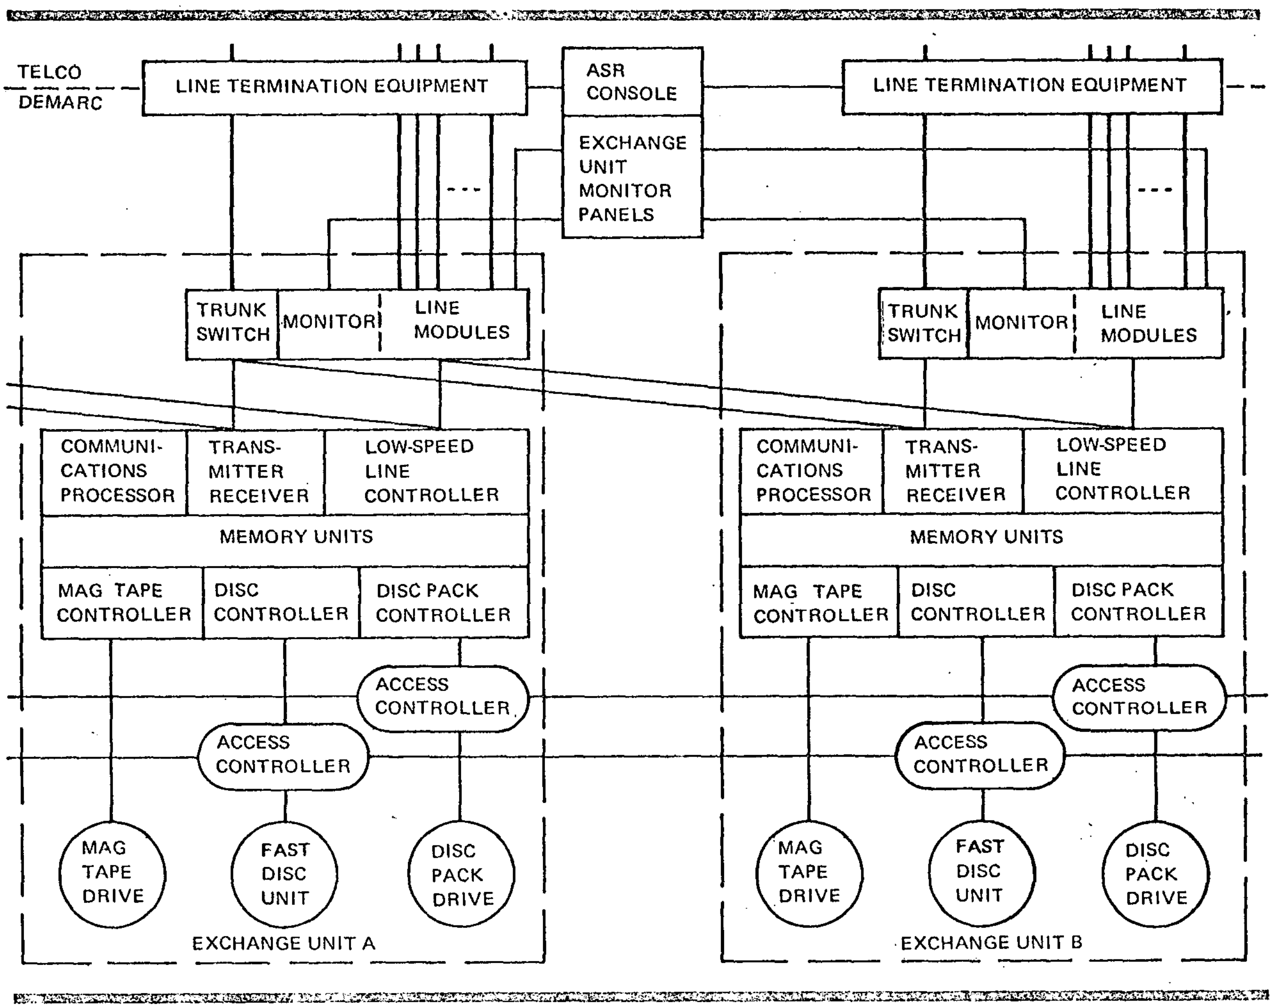

Figure 2 Exchange Unit Interconnections

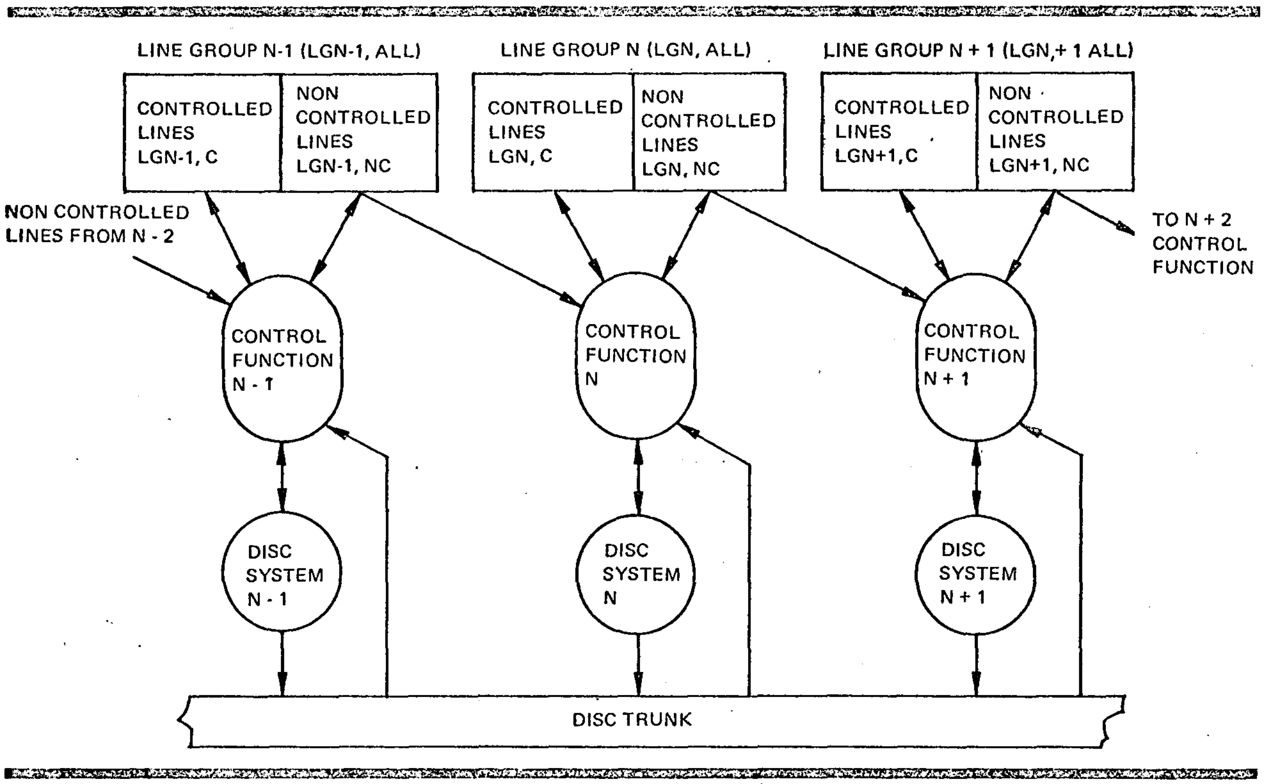

Figure 3 *Exchange Logical Organization*

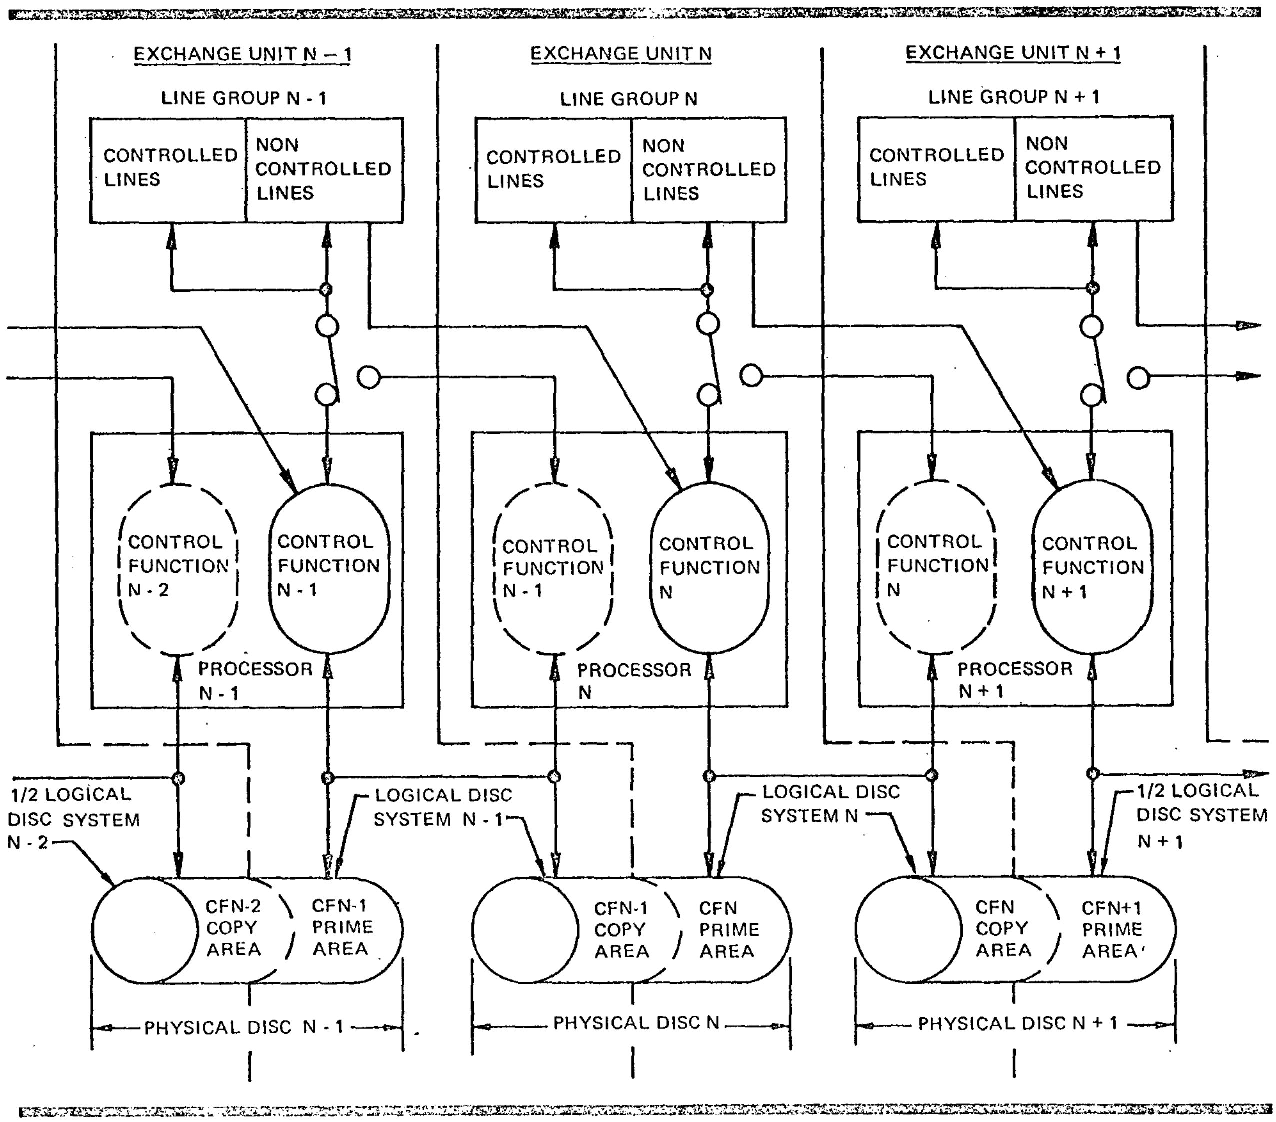

Figure 4 Relationship Between Exchange Logical and Physical Organizations

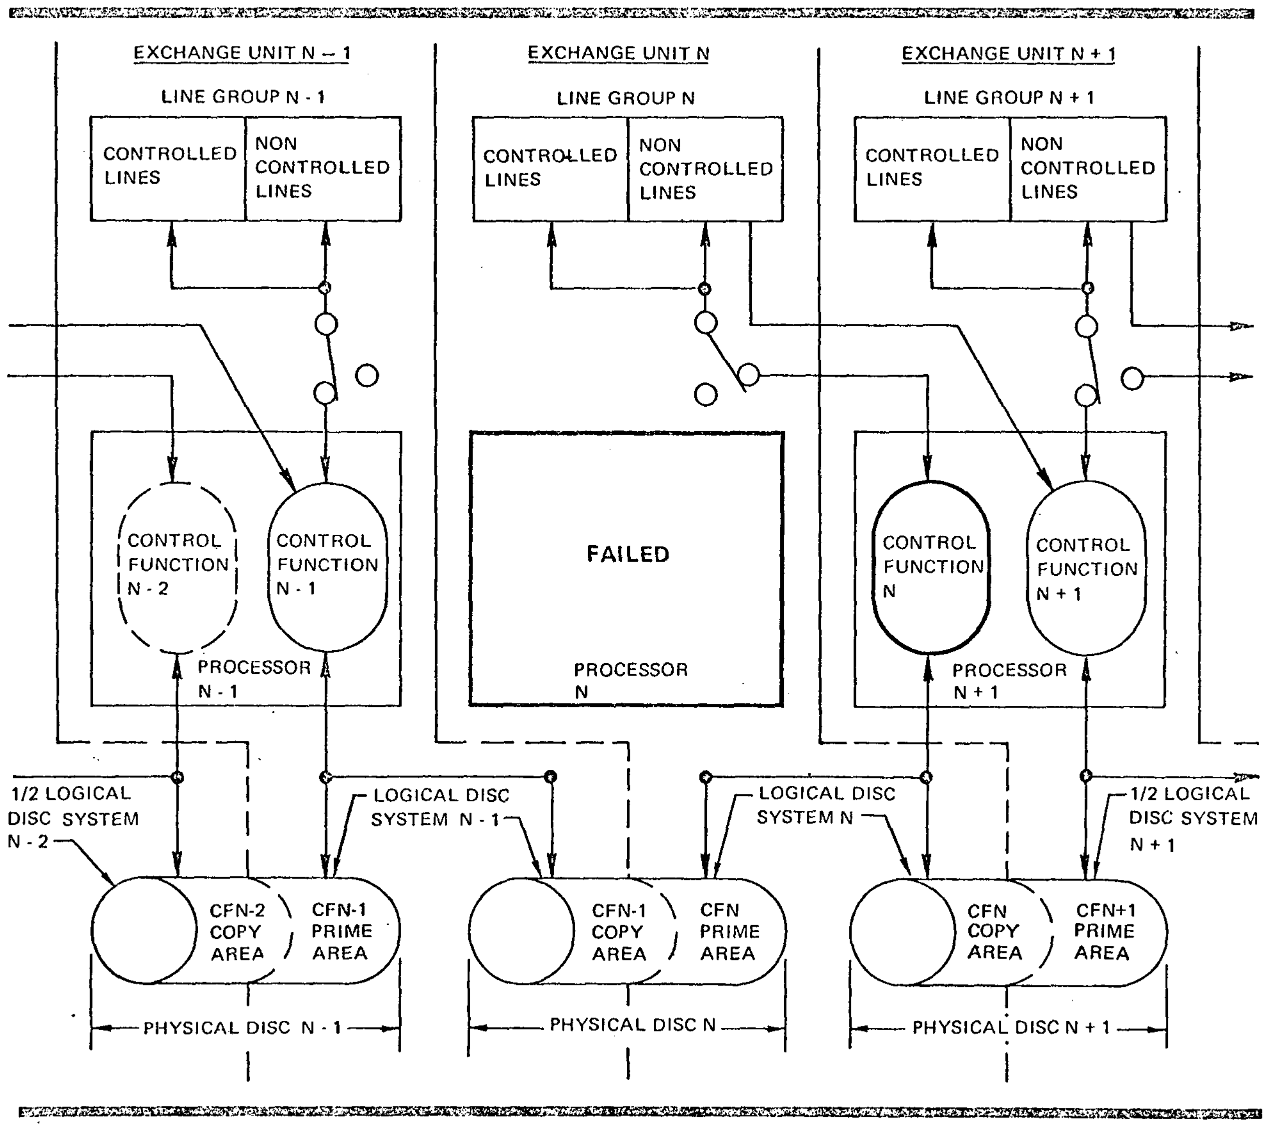

Figure 5 Exchange Organization, Processor N Failed

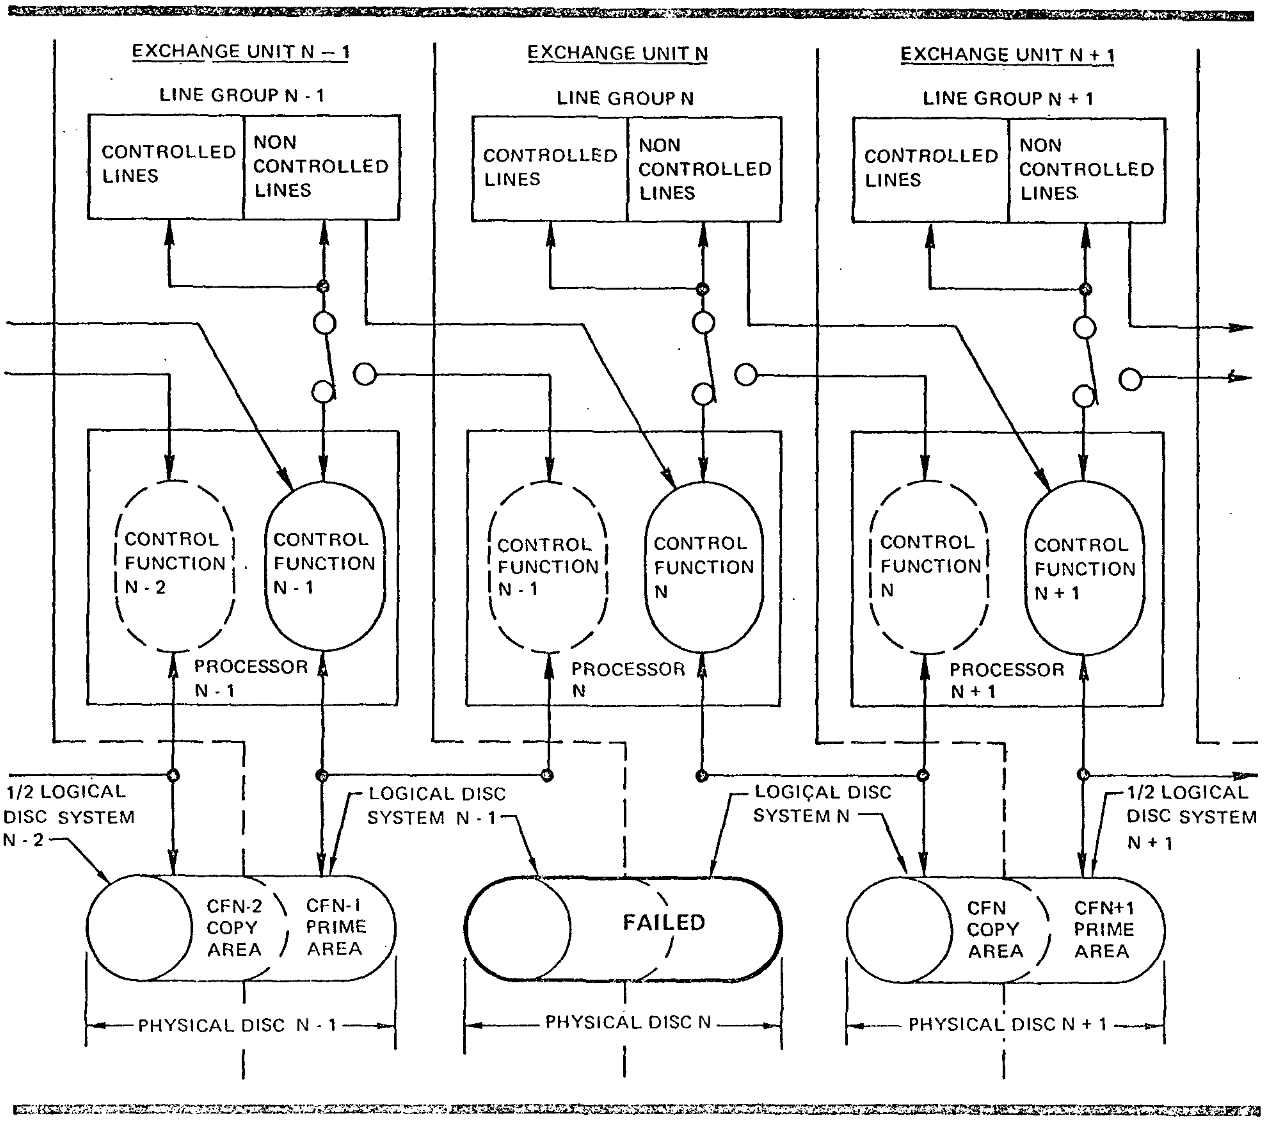

Figure 6 *Exchange Organization, Disc N Failed*

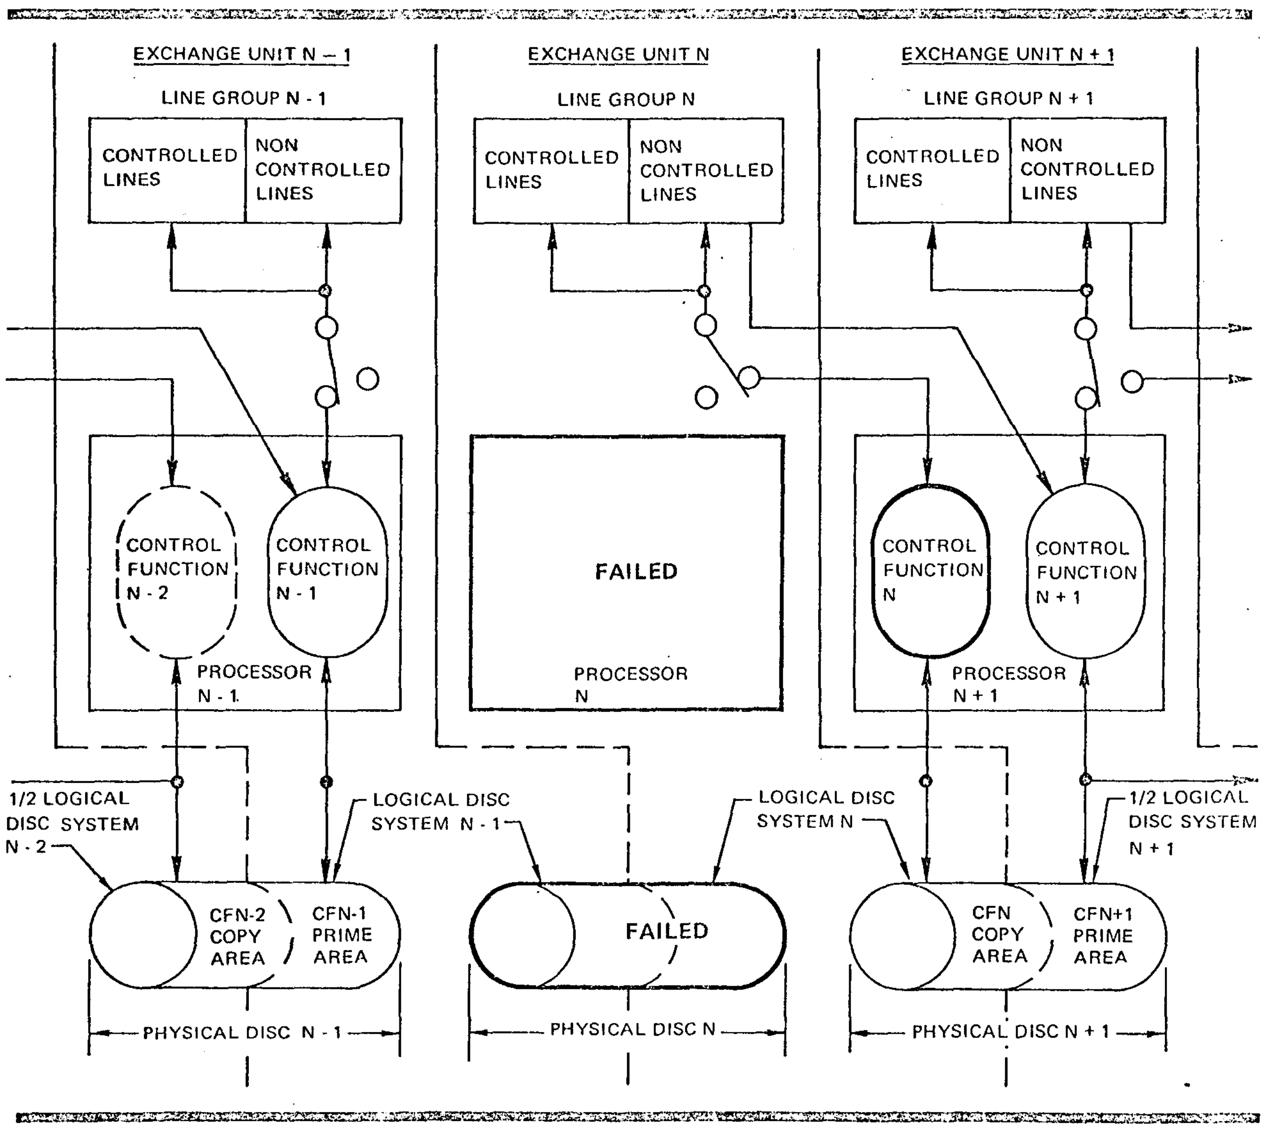

Figure 7 Exchange Organization, Processor N and Disc N Both Failed

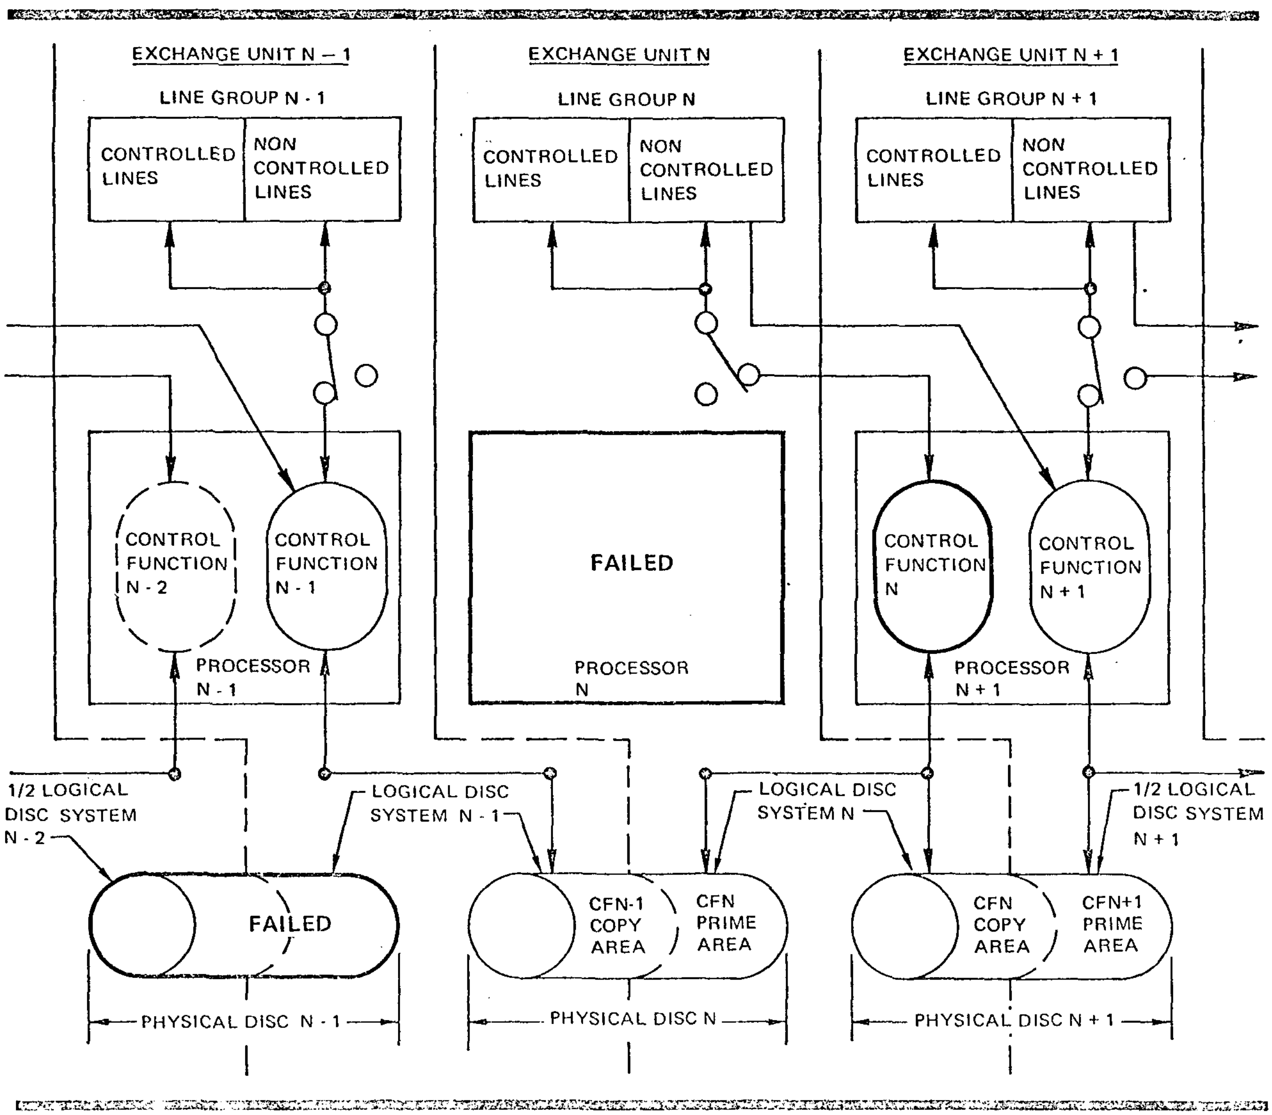

Figure 8 Exchange Organization, Processor N and Disc N-1 Both Failed

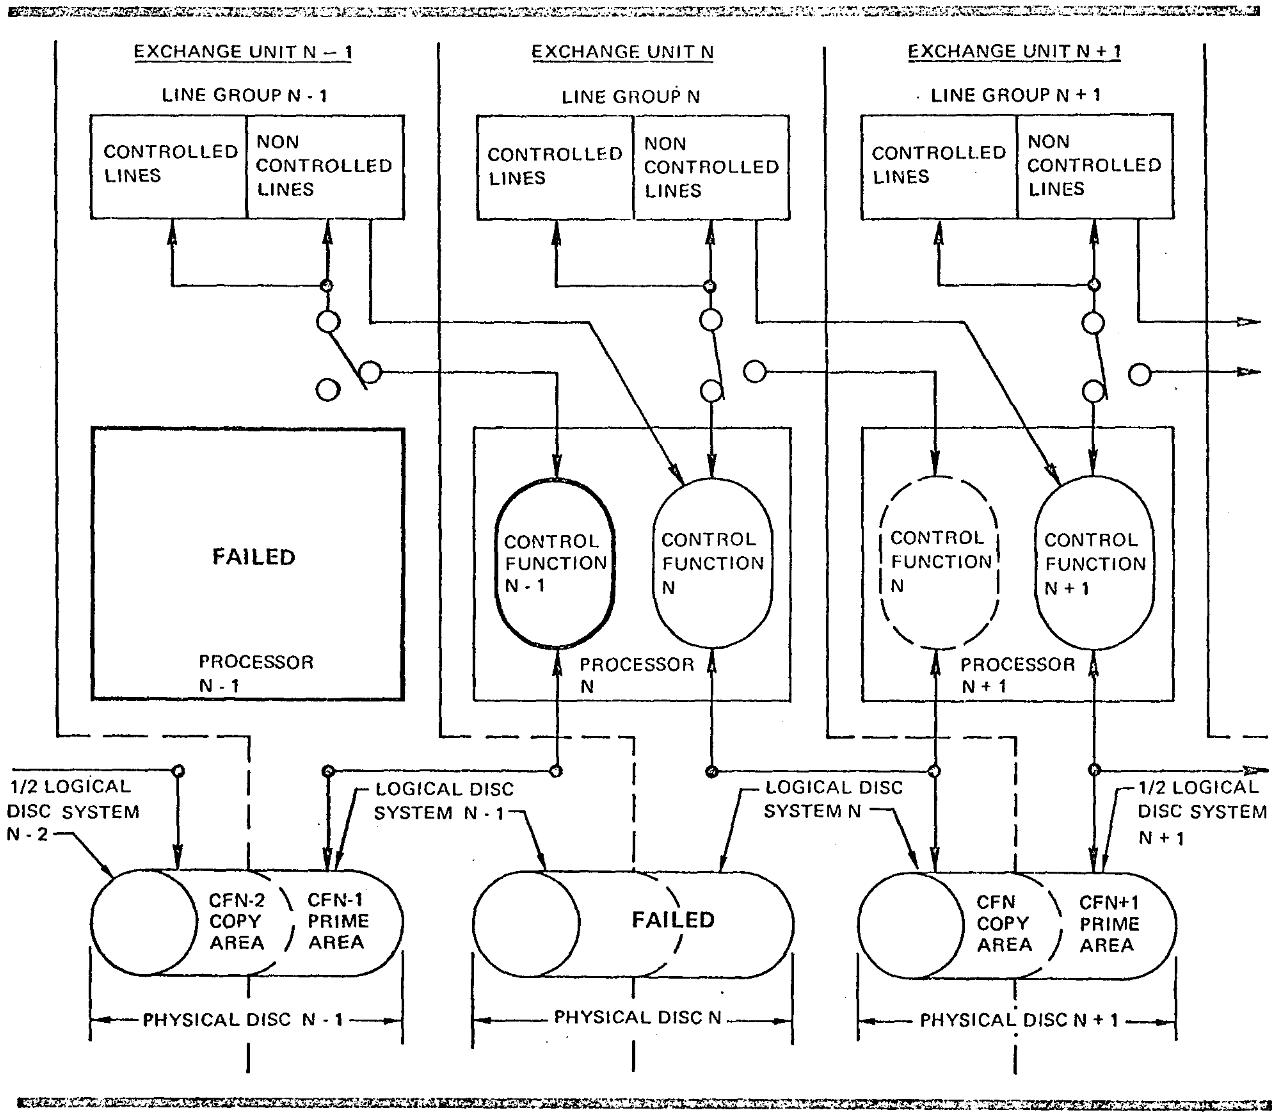

Figure 9 *Exchange Organization, Processor N-1 and Disc N Both Failed*

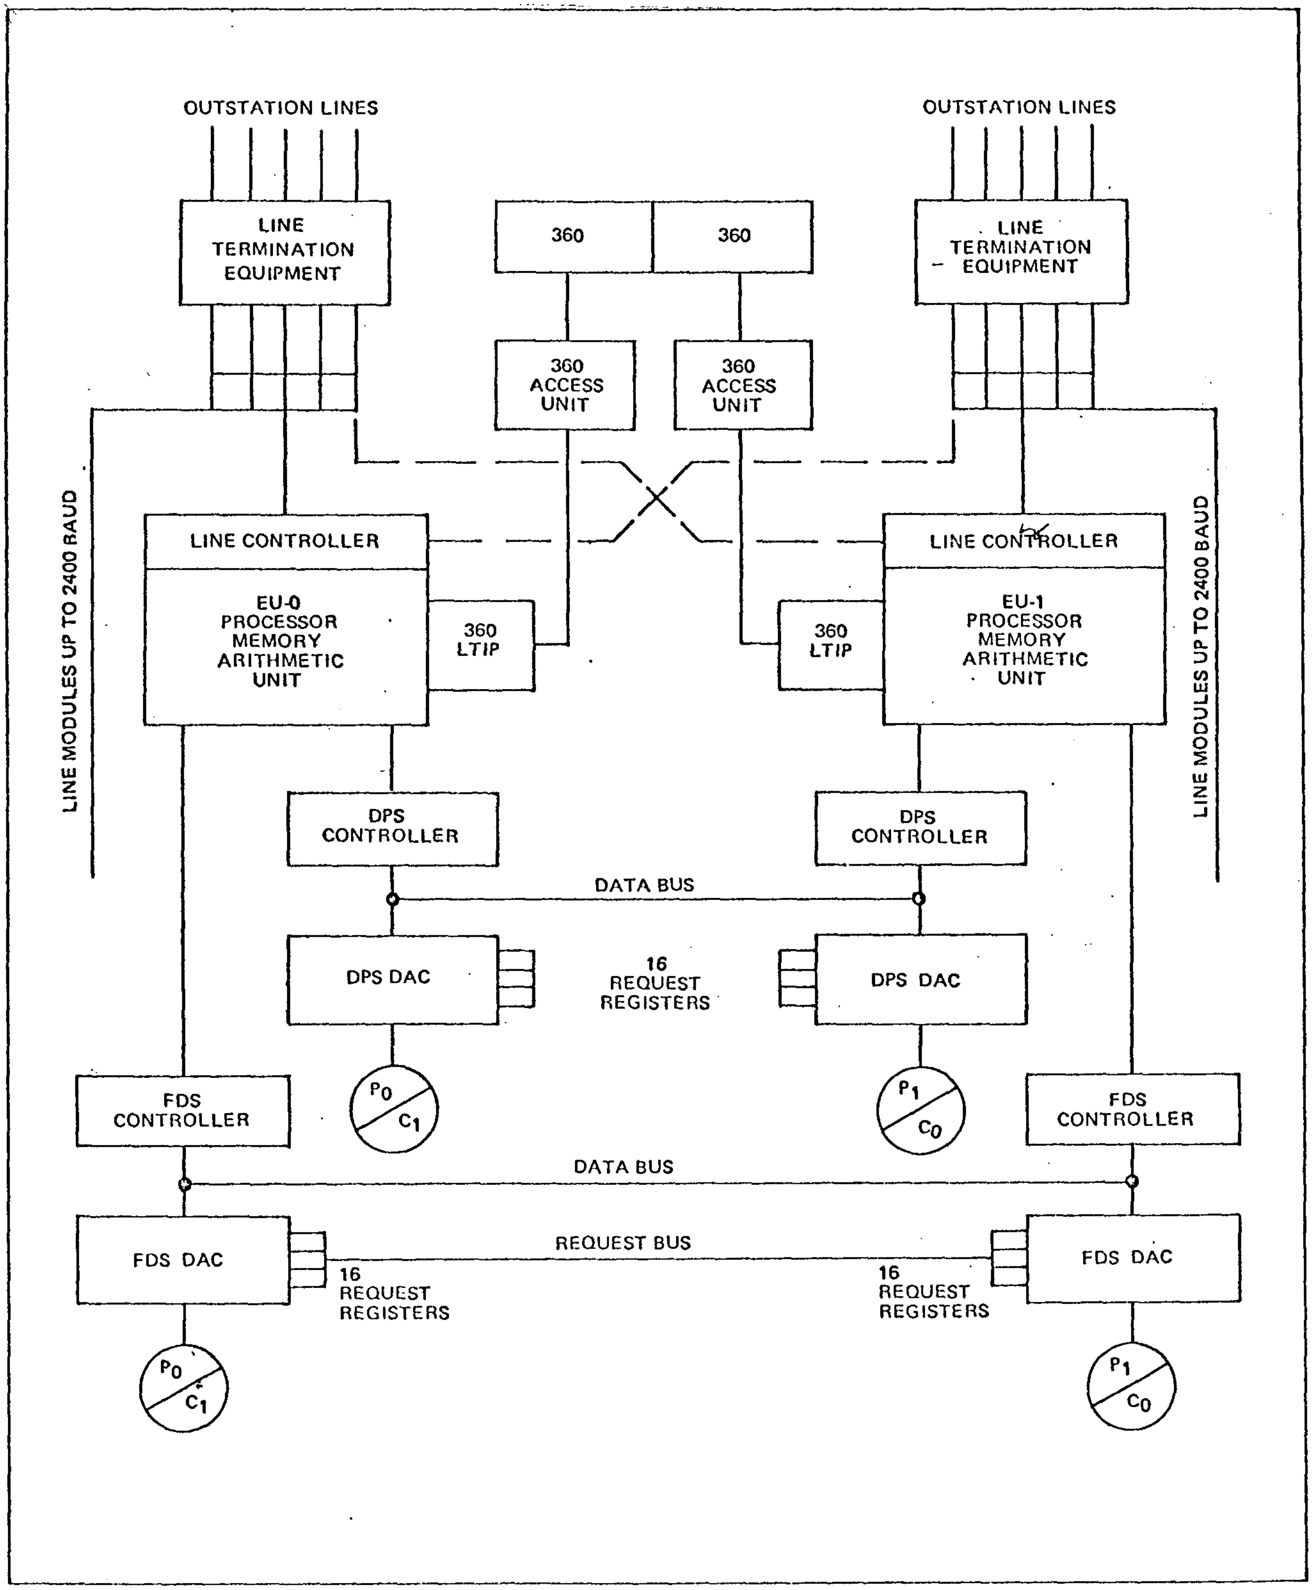

M1000 Dual System Concept

Figure 10

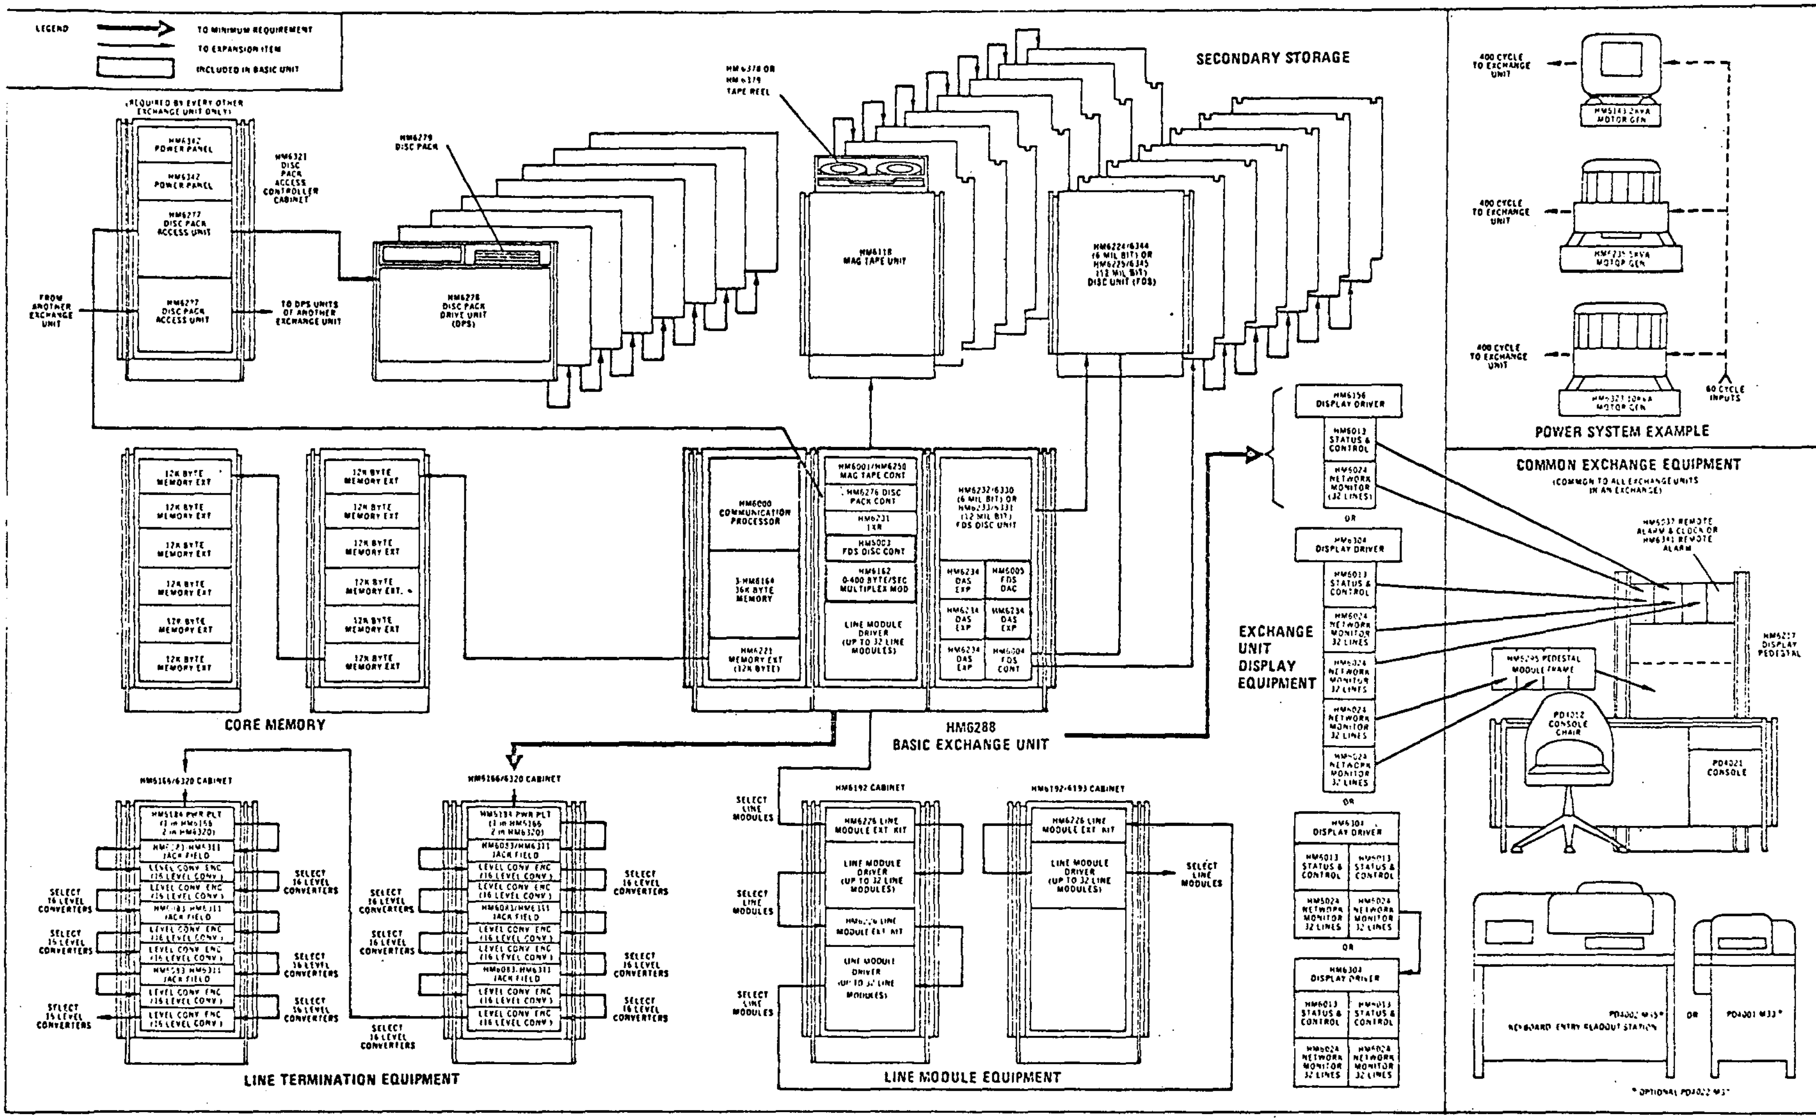

COMMUNICATIONS M1000 CONFIGURATOR

Figure 11

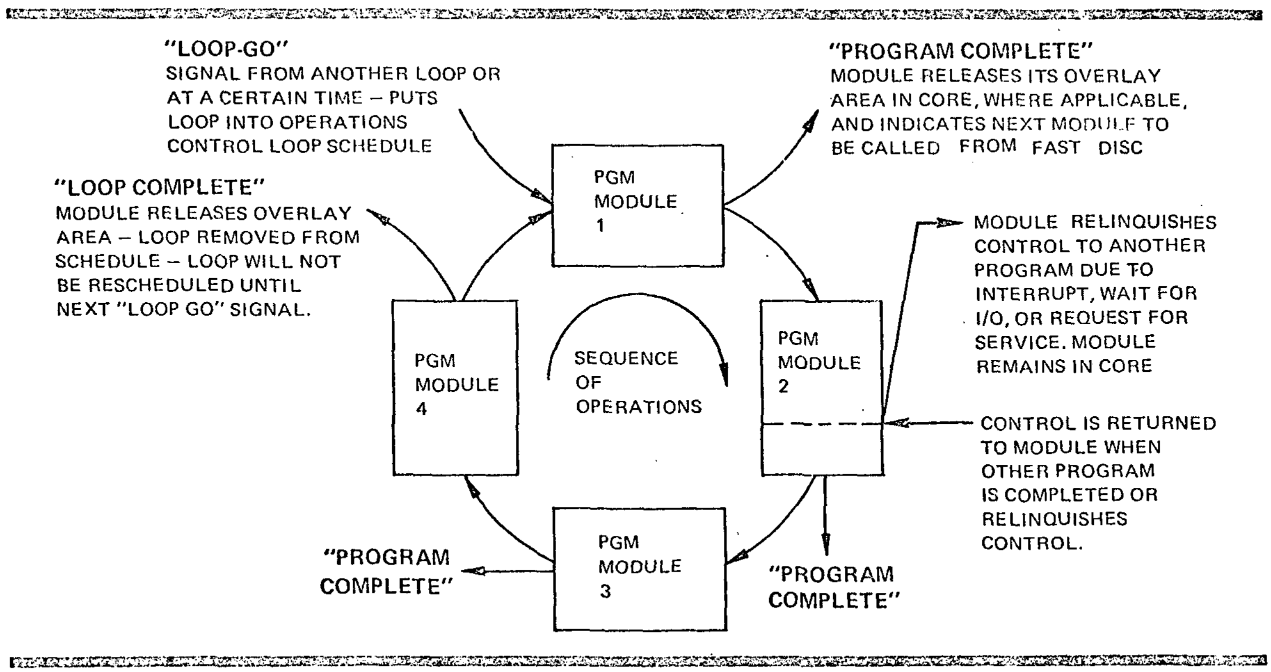

Figure 12 Program Loop

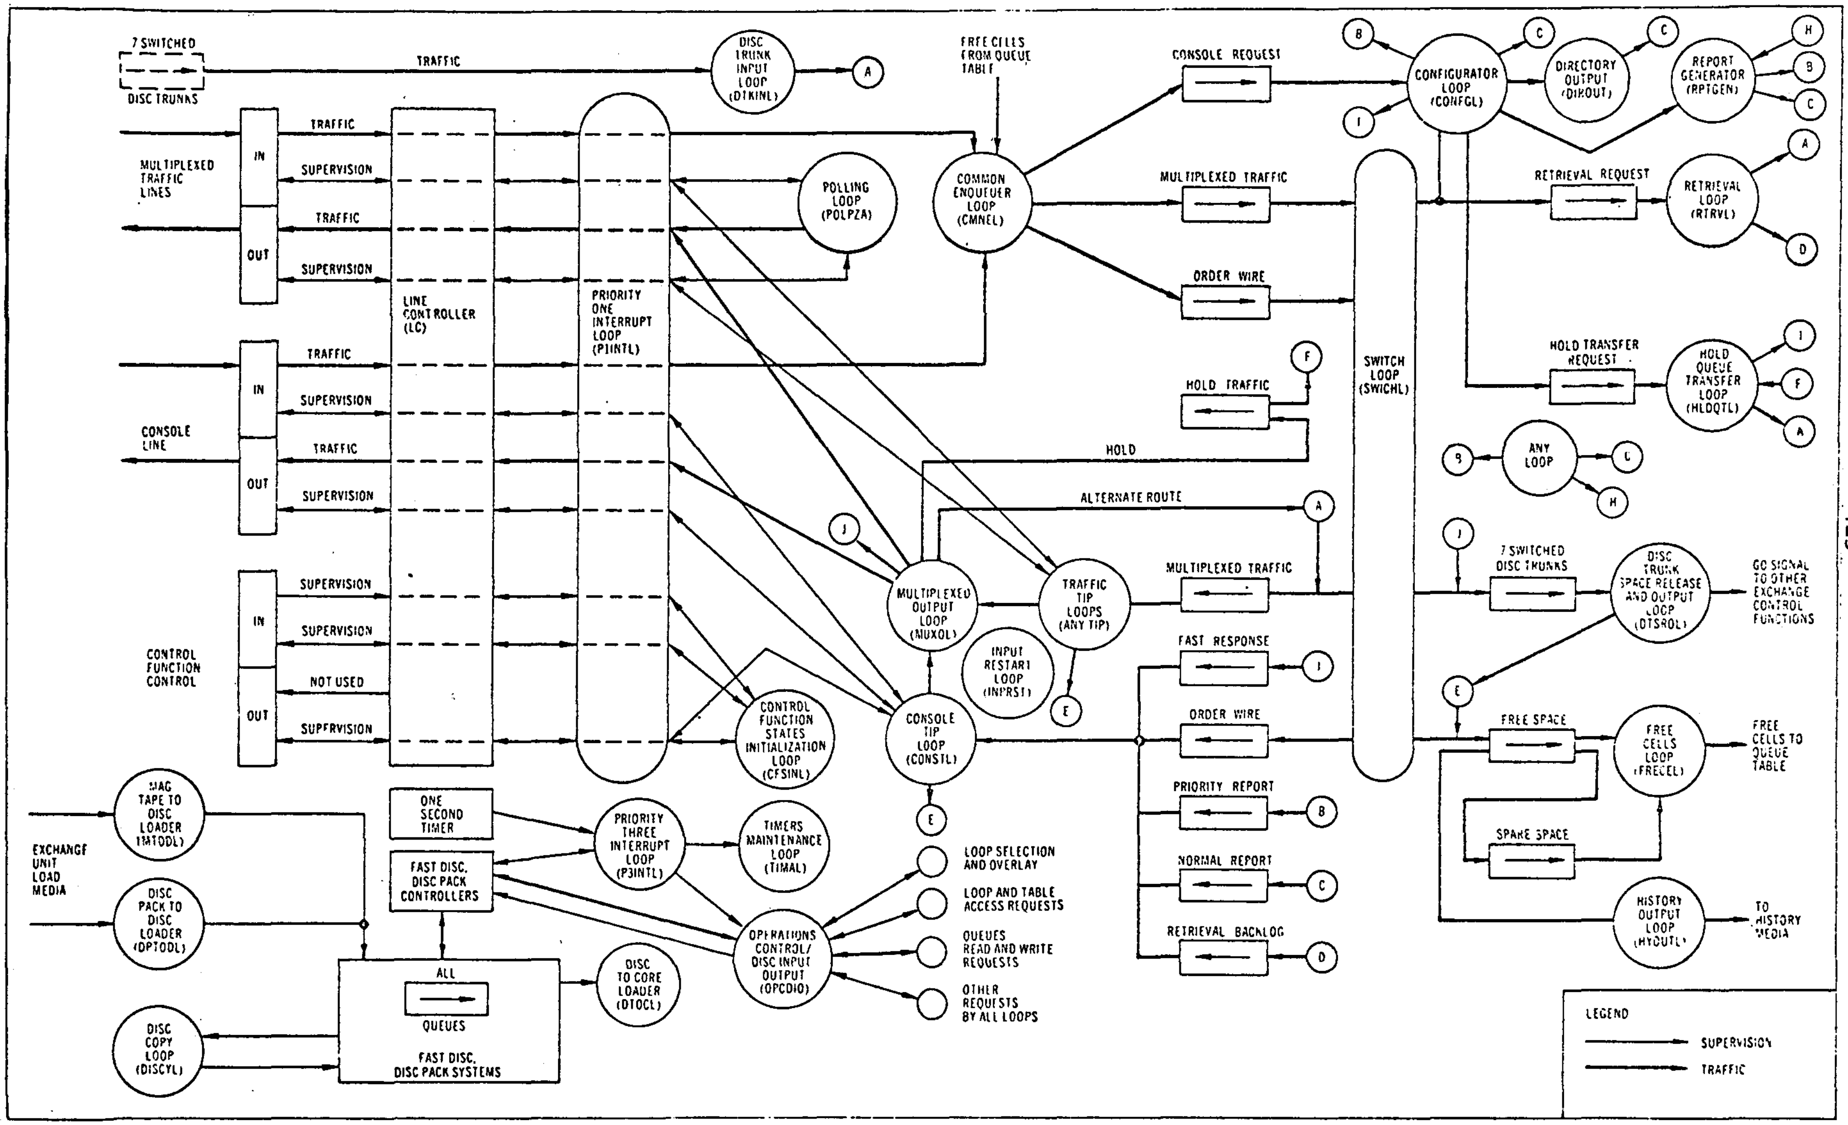

CONTROL FUNCTION LOOPS AND TRAFFIC AND SUPERVISION FLOW TO/FROM LINES