Computer Information Center

Computation Department

LTSS Livermore Time-Sharing System

Part I: OCTOPUS

Chapter Ø: INTRODUCTION TO THE OCTOPUS NETWORK

John G. Fletcher

December, 1970 Edition - 1

Lawrence Radiation Laboratory

University of California/Livermore

LIVERMORE TIME SHARING SYSTEM

PART I: OCTOPUS

CHAPTER Ø: INTRODUCTION TO THE OCTOPUS NETWORK

EDITION 1

by: JOHN G. FLETCHER

Systems Development Section
Computation Department

published by the

Computer Information Center
Computation Department
December, 1970

CONTENTS

		Page
Preface		1
Section 1:	Computers: The Network Components	3
	1.1 The Structure of a Computer	3
	Processor	3
	Memory	3
	Secondary Storage	4
	I/O (Input-Output) Devices	4
	Network Interfaces	5
	1.2 Types of Computers	5
	Worker Computers	5
	Elephant Computer	6
	Concentrators	6
	Terminal Computers	6
Section 2:	The Linkage of the Network	7
	2.1 The File Transport Network	9
	2.2 The Message Network	10
	Teletypewriter Subnetwork	10
	Data Acquisition Subnetwork	14
	Remote I/O Subnetwork	14
	2.3 Protocol of Transmissions	14
	Status Bits	16

C O N T E N T S

		Page
Section 3:	Message Standards	19
	3.1 Headings	19
	3.2 Character Set	23
	3.3 Control Characters	27
	Format Effectors	27
	Device Controls	28
	Information Separators	28
	Deletion Sentinels	28
	Code Extenders	29
	Communication Controls	29
Section 4:	The User	31
References		33
Index		35

I L L U S T R A T I O N S

Figure		Page
A	File Transport Network	8
B	The Teletypewriter Subnetwork	11
C	Data Acquisition Subnetwork	13
D	Remote I/O Subnetwork	15
E	Octopus Message Heading	20
F	Computer Numbers (in octal)	22
G	ASCII Code (in octal)	24

P R E F A C E

The *Octopus network* is the major information processing facility of the Lawrence Radiation Laboratory at Livermore; in fact, it is one of the largest (if not the largest) information processing facilities in the world. The network consists of a number of *computers* joined together by information-carrying *channels* (e.g., cables).

This publication describes the structure of the network and the major standards and conventions employed.

Section 1: COMPUTERS: THE NETWORK COMPONENTS

1.1 The Structure of a Computer

A computer is an information processing device capable of performing effectively, but with reduced capabilities, when disconnected from the network. A computer usually consists of several identifiable parts of varying function, as now summarized:

1. A *processor* is a machine which manipulates or transforms information. It contains units capable of performing arithmetic, logical, and decision making functions.

Typically, a processor consists of a number of *registers* which can hold the information upon which the processor is operating. Usually, one such register is the *program counter* or *instruction counter*, which selects a *word* from the *memory* of the computer (see 2. below). The selected word is viewed as an *instruction* to the processor as to what manipulations it should carry out next. An instruction may direct the processor to move a word between a register and the memory, to alter the content of a register on the basis of the content of the same or other registers, to operate an *I/O device* (see 4. below), to depart from the normal sequential incrementing of the program counter, etc.

Some computers (e.g., PDP-8's) have only a single processor. Others (e.g., PDP-10's) have two or more nearly identical processors. Still others (e.g., 7600's) have several processors of two or more kinds. A processor is usually constructed of high-speed, solid-state, electronic elements.

2. The *memory* of a computer is a device which stores (or remembers) a large amount of information in a form which can be rapidly and randomly accessed by the processor or processors.

Usually, memories are organized into words. A *word* is a collection of a fixed number of bits; the fixed number, the word *size*,

Section 1: COMPUTERS: THE NETWORK COMPONENTS

depends upon the computer and typically ranges from about 12 bits (e.g., PDP-8) to 64 bits (e.g., STAR). Each word in the memory has an *address*, a number which identifies the word. (The content of a program counter is, for example, an address.) Memories typically range from 2^{12} words (e.g., PDP-8) to 2^{19} words (e.g., 7600).

Memories are commonly constructed of tiny magnetic elements called "ferrite cores."

3. The *secondary storage* of a computer consists of a number of more-or-less independent information storing devices which are normally of greater capacity than the computer memory, but are slower to access. Information is moved between these devices and the memory as the computer's need for the information waxes and wanes.

The typical secondary storage devices are disks and drums, which keep information in a medium rotating at high speed; such devices typically hold 10^8 to 10^{10} bits. One secondary storage device in the network (the Photostore) holds over 10^{12} bits.

The distinction between secondary storage devices and I/O devices (see 4. below) is not clear cut, and often such devices as magnetic tape transports are used for secondary storage.

Information in secondary storage is usually divided into logical units called *files*.

4. I/O (input-output) devices are more-or-less independent devices which enable information to be moved between the computer memory and the outside (human) world.

Typical I/O devices are card readers, card punches, line printers, magnetic tape transports, paper tape reader-punches, teletypewriters, and consoles (switches and lights on the front of the processors).

5. A distinction is made here between I/O devices and the *network interfaces* which connect a computer to the channels leading to other computers in the network. An interface may connect to only one channel, or it may be a selector interface which connects to several channels.

1.2 Types of Computers

The computers of the Octopus network can be classed into four types, depending upon their function, as is now summarized:

1. *Worker computers* are very fast, very large computers. Their function is to execute (process) programs (sequences of instruction and data words) for the users of the network; these programs include those written by the user as well as generally available programs, such as compilers and assemblers. The other computers in the network exist to enhance the capabilities of the workers and to make the use of them more convenient.

Each worker computer is named by a single letter, its *designator*. Currently, the network includes the following workers (designators in parentheses): two CDC 6600's (L, M), two CDC 7600's (R, S), and one CDC STAR (T). (A third 6600 (G) is not part of the network, but operates autonomously; so, too, do an XDS Sigma-7, a DEC PDP-1, and an IBM 7094.)

Each worker computer is *time-shared*; that is, it alternates execution among the programs of many users, to some extent giving each user the illusion of having a computer to himself; in this way, the "thinking time" and other delays in the activity of a single user do not result in idle time for the computer. The scheduling of this alternating execution is controlled by special programs which constitute the *operating system* of the computer. The user controls

Section 1: COMPUTERS: THE NETWORK COMPONENTS

the activity of his programs in a worker computer by exchanging messages with that computer from an *interactive terminal*, an I/O device (e.g., teletypewriter) somewhere in the network.

The bulk of this manual is concerned with describing the user facilities available on the worker computers. Each section usually concerns some aspect of the computer or some service it might perform; where there are differences among the worker computers, these are treated in different sections of this LTSS Chapter.

2. The *Elephant computer* is a dual-processor DEC PDP-10. Its major function is to manage the large central storage facility of the network, the *Elephant file maintenance system* discussed in LTSS Chapter 4.[1]

The PDP-10 also performs for the network other service functions which require a fairly large computer, such as controlling some high-performance displays. (See Section 2.1).
Further, it provides an alternate or "back-up" routing for messages between the workers and interactive terminals (see section 2.2). Finally, the PDP-10 can be used to execute certain user programs, just like a worker computer. Its designator is P.

The facilities of the Elephant computer are discussed in the same LTSS Chapters as the corresponding worker facilities.

3. *Concentrators* are small computers which move messages and other information between worker computers, the Elephant computer, and various I/O devices. Typical concentrators are DEC PDP-8's and DEC PDP-11's. Concentrators are not available for the execution of user-written programs. They are described briefly in connection with discussion of the Octopus facilities that use them.

4. A *terminal computer* is a small computer attached to a channel from a concentrator in lieu of an I/O device. It usually controls several actual I/O devices. For the most part, a terminal computer differs from a concentrator in that it has a more peripheral and subordinate role in the network.

Section 2: THE LINKAGE OF THE NETWORK

The Octopus network has two major functions, both of which enhance the power of the worker computers. First, it provides access to the large central storage facility maintained by the Elephant computer. Second, it provides interconnections so that the workers and the Elephant computer may converse with one another and with interactive terminals and other remote I/O devices.

The volume of information and the transfer rate appropriate to the first of the functions is much greater than those appropriate to the second. Consequently, the Octopus network really consists of two networks, the *file transport network* and the *message network*, as now described.

Section 2: THE LINKAGE OF THE NETWORK

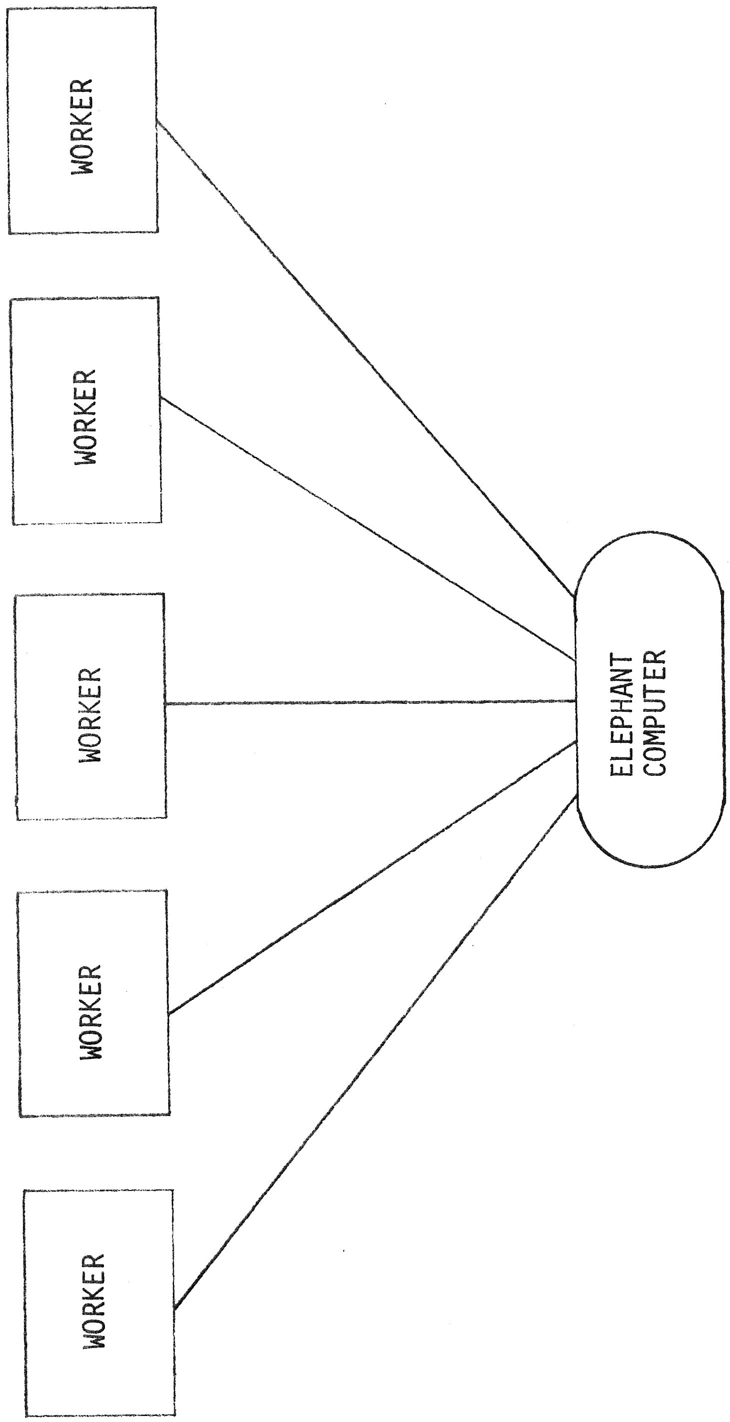

Figure A: File Transport Network

Section 2: THE LINKAGE OF THE NETWORK

2.1 The File Transport Network

Each worker computer is connected directly to the Elephant computer by a high-speed channel with a transfer rate of about 12 megahertz. (See Figure A.)

These channels are used to transport files between the workers and the Elephant computer. Requests for transports come to the Elephant computer over the message network. (See LTSS Chapter 4.) All transports are scheduled by the Elephant computer in such a way as to optimize the use of Elephant storage media.

A transport is initiated by a brief message from the Elephant computer to the scheduled worker computer over the file transport channel. This message conveys the essential information the worker requires, such as the identity of the file by name and *user number* (see Section 4), first word address and count, and direction of transport. The worker is then expected to respond in less than a second and begin reading or writing the data. At the end of the transfer, the worker sends a brief message to the Elephant computer indicating whether it detected any error.

Based on this message and its own observations, the Elephant computer decides whether the transport was successful, whether it should be rescheduled, or whether it should be abandoned. In the first and last of these cases, a suitable reply is sent through the message network to the computer from which the original request was made.

If calculations are made, taking into account such things as the buffer sizes, effects of rotational delay in secondary storage, the initial delay in the response of the worker computer to the transport-initiating message, and the amount of information in a

Section 2: THE LINKAGE OF THE NETWORK

single transport (called a *burst*, with a maximum of a little under four million bits), it is found that an average transport rate of about 3 megahertz can be maintained. This is adequate for present needs, in view of the characteristics of Elephant storage. (See LTSS Chapter 4.)

Information received by the Elephant computer is stored in the Elephant filing system, is forwarded to another worker computer, or is displayed on a display device operated by the Elephant computer.

The file transport network is used for one other purpose, as an alternate or "backup" path for certain messages which (because of a failure) cannot take their normal routing through the message network. Details are given in the next section.

2.2 The Message Network

The message network is far more complex than the file transport network.

Each link in this network joins a concentrator with a worker, the Elephant computer, another concentrator, an I/O device, or a terminal computer. The number of links and concentrators is subject to fairly frequent change.

At the present time, the situation can best be described by dividing the message network into three subnetworks, as follows:

1. The *teletypewriter subnetwork* connects all worker computers and the Elephant computer to about ~~300~~ interactive teletypewriter terminals. (See Figure B.) ~300

? 2. Data Acquisition Network 9/72

3. RIOT

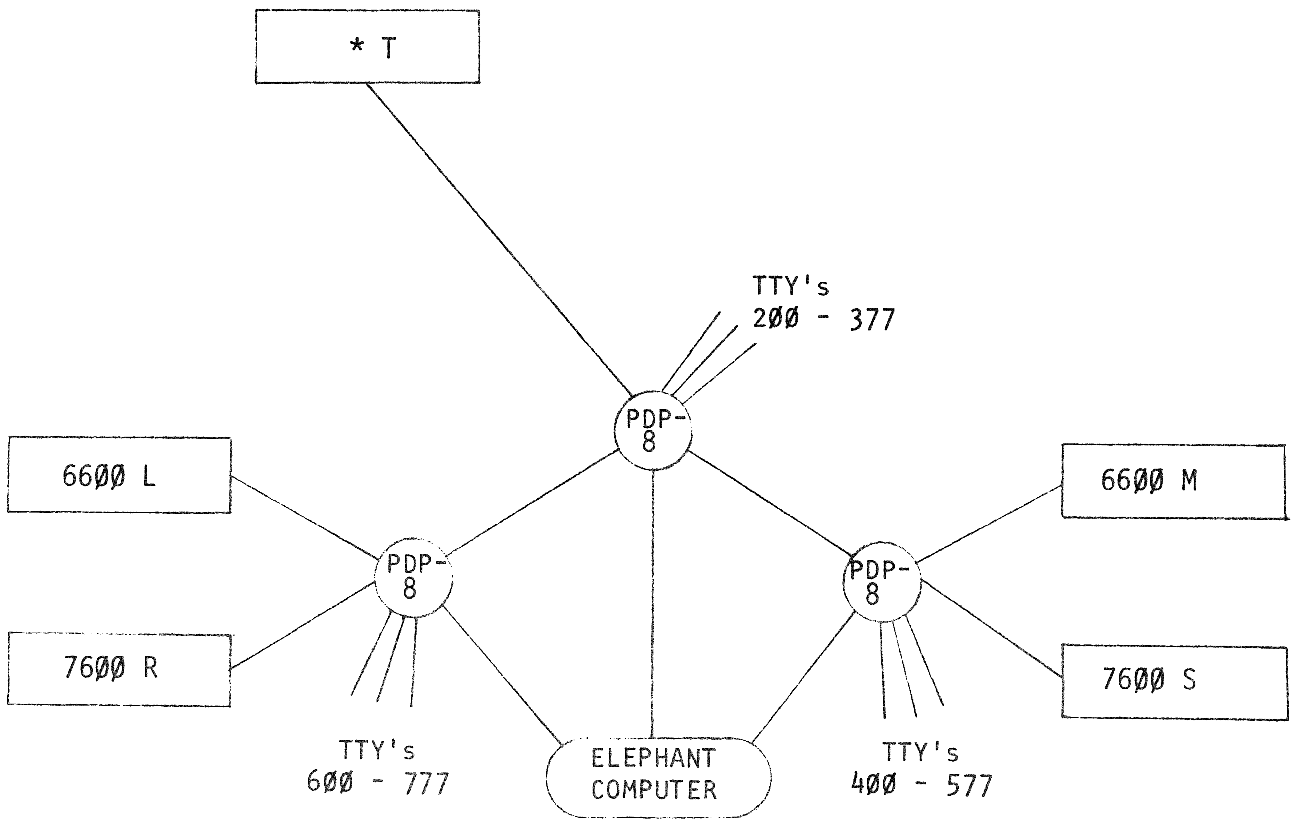

Figure B: The Teletypewriter Subnetwork

Section 2: THE LINKAGE OF THE NETWORK

The concentrators are DEC PDP-8 computers called *tentacles*; each tentacle is connected to one or more workers, one or more other tentacles, the Elephant computer, and up to 128 teletypewriters. The teletypewriters attached to each tentacle are assigned consecutive octal numbers beginning at some multiple of (octal) 200.

A brief examination of Figure B will show that any teletypewriter can communicate with any worker computer or the Elephant computer, the message passing through one, two, or all three tentacles. However, if the channel connecting a worker to a tentacle (or the tentacle itself) should fail, the worker is cut off from the subnetwork. Therefore, whenever a tentacle finds that it cannot forward a message for a worker along its normal path, it sends it to the Elephant computer which then may use the file transport network to try to forward the message; similarly messages from a worker can alternatively be sent over the file transport network and then be forwarded to the proper tentacle by the Elephant computer.

It is possible for nonteletypewriter messages to pass through a tentacle (or the alternate path through the file transport network), but if they do, they are limited to the maximum size imposed on teletypewriter messages (1212 bits of text). Teletypewriter use is discussed in LTSS Chapter 3.[2]

Section 2: THE LINKAGE OF THE NETWORK

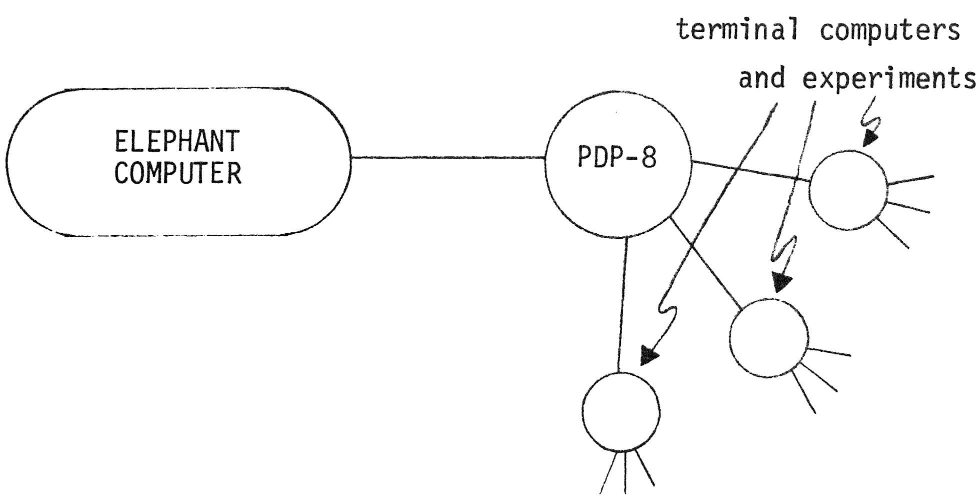

Figure C: Data Acquisition Subnetwork

2. The *data acquisition subnetwork* consists of a single concentrator (a PDP-8), attached to the Elephant computer, and to up to sixteen terminal computers. (See Figure C.) The latter are small computers (programmed and maintained by users) which usually collect data from semiautomated experimental arrangements. This data is then sent to the concentrator and on to the Elephant computer, where it is usually entered into the Elephant filing system; however, forwarding of messages to the workers through the teletypewriter subnetwork is also possible.

Data may also travel through the data acquisition subnetwork in the reverse direction. Data acquisition will be discussed in a later LTSS chapter.

3. The *remote I/O subnetwork* consists of a single concentrator (a multiprocessor DEC PDP-11), attached to all workers and the Elephant computer and to a number of terminal computers (PDP-8's). (See Figure D.)

Each terminal computer typically manages such I/O devices as a card reader and a line printer. This subnetwork is used for the remote input or output of files; that is, information is transported from a card reader (or other input device) to a worker computer or Elephant file or is transported from the file to a line printer (or other output device).

Nonfile information may also be conveyed. Remote I/O operation will be discussed in a later LTSS chapter.

2.3 Protocol of Transmissions

Although it is of little concern to the user, the protocol of transmissions in the message network is now reviewed. Only transmissions among workers, the Elephant computer, and concentrators are considered. (Transmissions between concentrators and terminal computers conform to varying standards.)

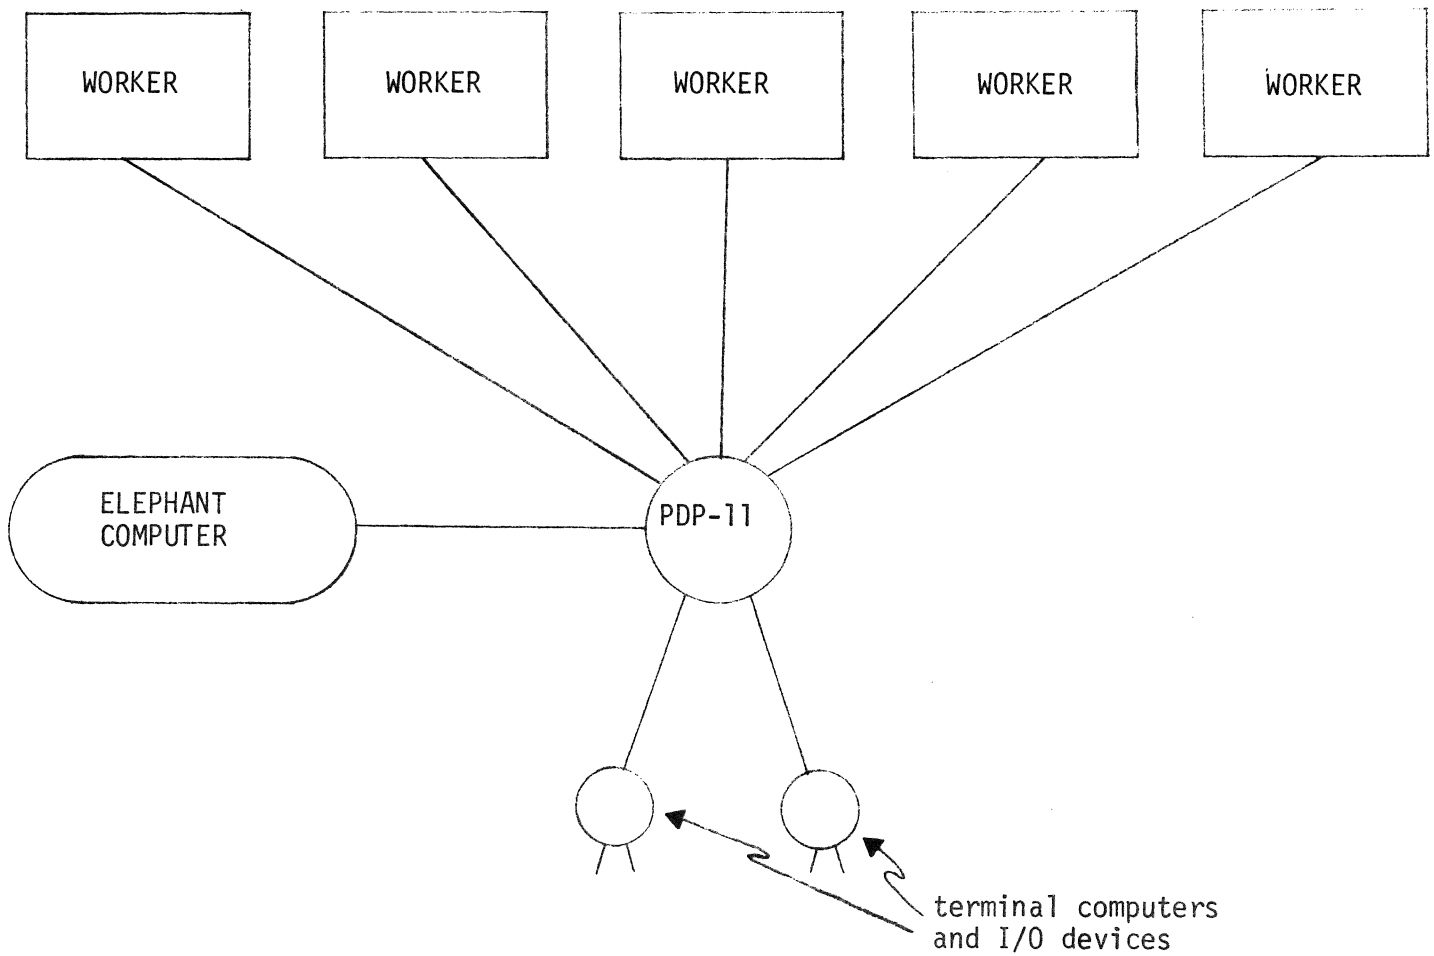

Figure D: Remote I/O Subnetwork

Section 2: THE LINKAGE OF THE NETWORK

For the present purpose, attention can be limited to two computers and the channel joining them. Further, the point of view of one of the two computers is taken; this computer is therefore called "self," as opposed to the other computer.

The network interface has four status bits known as *send requested, send permitted, receive requested,* and *receive permitted.* Only the first and last of these are under the direct control of (may be turned on and off by) self; the other two are controlled by the other computer and may only be sensed by self.

Send requested is set by self if and only if it has one or more messages it wishes to send to the other computer.

Receive permitted is set by self if and only if it is willing to receive messages sent by the other computer.

Send permitted is a copy of the other computer's receive permitted, and receive requested is a copy of the other computer's send requested.

Thus, whenever both the requested and the permitted bit for the same direction of transmission are set, self should issue the necessary instructions to begin the transmission; the other computer presumably will issue the appropriate instructions for its end of the channel at about the same time.

At the conclusion of the transmission, both computers check the error status bits in the interface, which should be set at both ends of the channel or not at all. Among the possible error status bits ~~are~~ is one indicating that one computer or the other has issued a transmission-aborting instruction, and one indicating that a clock in the interface has timed out because one computer or the other unduly delayed in proceeding with the transmission.

Section 2: THE LINKAGE OF THE NETWORK

Concentrators leave their receive permitted bits set most of the time and preferentially engage in a receive transmission if both send and receive are possible. This is ideal when dealing with a worker or the Elephant computer (as is explained below) but raises a problem when dealing with another concentrator, for both concentrators may set their send requested (in addition to their receive permitted) bits simultaneously, after which they would both begin a receive operation. To prevent this, the interface is designed so as to inhibit the setting of receive requested if send requested is already set.

Concentrators are expected to wait; workers and the Elephant computer are not. A worker computer can leave its send requested and receive permitted bits off most of the time (even when it has a message to send). Whenever it notices that a concentrator has caused send permitted or receive requested to be set, it can then (if it wishes) set send requested or receive permitted and expect that the concentrator will begin the transmission within a few hundred microseconds.

Section 3: MESSAGE STANDARDS

?

Since the user and/or his program is often the sender or receiver of information conveyed by the message network, some discussion of the standards and conventions used in that network is appropriate. This is the concern of the present section.

3.1 Headings

Messages in transit usually consist of two parts: *heading* and *text*. The text is the message proper, the information that the sender, origin, or *source* is trying to convey to the receiver, destination, or *sink*. However, this information is preceded by the heading, which contains information essential to the proper delivery of the message, such as a description of the sink. The heading is interpreted by each network node through which the message passes in order that it can be decided what is next to be done with the message. The Octopus message network uses a standard heading format, which is now described, although headings are always generated by operating systems, not by the user.

The basic heading consists of 48 bits divided into fields, as indicated in Figure E; the significance of each field is now discussed (proceeding from left to right).

The *sink* field (16 bits) indicates the intended receiver of the message; similarly the *source* field (16 bits) indicates the sender. Each of these fields is divided into two subfields: the *computer* field (8 bits) and the *device* field (8 bits).

SINK		SOURCE		ID	BYE	ERROR	BINARY	TYPE	CONTROL	CHAIN
COMPUTER	DEVICE	COMPUTER	DEVICE							
8	8	8	8	1	1	1	1	4	7	1

Figure E: Octopus Message Heading

Section 3: MESSAGE STANDARDS

Each computer in the network, except for terminal computers, is assigned a computer number; in the case of worker computers and the Elephant computer, this number is the obvious translation of the letter designator. Current assignments are given in Figure F.

The device number has varied meaning: in the case of a worker or the Elephant computer, it indicates a part of the operating system that is interested in the message. In the case of a concentrator, it usually indicates an I/O device, such as a particular teletypewriter or experiment; details are given in LTSS Chapter 3 and the Chapters (not yet published) on "Data Acquisition" and "Remote I/O Terminals."

The *ID* bit indicates that the message is the first of a sequence of messages, such as a log in the message from a teletypewriter. Similarly, the *bye* bit indicates that the message is the last of a sequence of messages, such as a log out message to a teletypewriter.

The *error* bit indicates that the message is being returned to its original source as undeliverable or unacceptable. The network node setting this bit also interchanges the source and sink fields so as to redirect the message whence it came; it may also set the *bye* bit if it knows that the intended sink can never be reached. Intermediate nodes of a message transfer never otherwise alter a message. The intended sink may set control bits (see below) in connection with the error bit to indicate why the message is unacceptable. If, when attempting to set the error bit, a network node finds it already set, it simply discards the message, since presumably neither its source nor its sink can be reached; this prevents messages from moving in endless loops.

The *binary* bit indicates the message text is not to be regarded as consisting of characters (as described in the next section).

Section 3: MESSAGE STANDARDS

ØØØ	Unused
ØØ1 - Ø32	Worker (or Elephant) computers A - Z
Ø33 - Ø37	Unused
Ø4Ø	Tentacle, teletypewriters ØØØ - 177
Ø41	Tentacle, teletypewriters 2ØØ - 377
Ø42	Tentacle, teletypewriters 4ØØ - 577
Ø43	Tentacle, teletypewriters 6ØØ - 777
Ø44 - Ø47	Unused
Ø5Ø	Data collection concentrator
Ø51 - Ø57	Unused
Ø6Ø	Remote I/O concentrator
Ø61 - 377	Unused

Figure F: Computer Numbers (in octal)

Section 3: MESSAGE STANDARDS

The *type* field (4 bits) indicates the general nature of the message. In particular, type Ø is for messages between interactive terminals and the programs which they control.

The *control* field (7 bits) has various interpretations, depending upon the content of the other fields, particularly the type field. If and only if the *chain* bit is set, an additional 48 bits of heading follows, consisting of 47 control bits and another chain bit; in this way the control field may be extended to any length. An extended control field often is used to hold such items as a user number and a *security level* (see Section 4).

3.2 Character Set

Messages often include character strings; if the binary bit is not set, the text consists solely of such strings. The Octopus standard character set is the American (National) Standard Code for Information Interchange (ASCII). ASCII is a 7-bit code comprising 128 characters. The Octopus standard treats ASCII in 8-bit bytes, the high order bit being Ø; this provides for future expansion to a 256-character set. Characters in messages are packed without regard to word boundaries; however, various computers in the network may internally store the characters in an alternate format more suitable to their word size.

The 64-character subset from (octal) 4Ø to (octal) 137, except that the upper case letters are to be regarded as replaced by the corresponding lower case letters, contains the most widely used characters; it is called the standard 6-bit subset. For some purposes (e.g., Elephant file names) only this subset is used and is stored as 6-bit bytes; the 6-bit representation is obtained by subtracting (octal) 4Ø from the ASCII representation, after the codes for the upper case letters have been interchanged. This represents a very recent decision which will ease the transition to keyboards with lower case letters.

Last digit / First digits	Ø	1	2	3	4	5	6	7
ØØ	NUL	SOH	STX	ETX	EOT	ENQ	ACK	BEL
Ø1	BS	HT	LF	VT	FF	CR	SO	SI
Ø2	DLE	DC1	DC2	DC3	DC4	NAK	SYN	ETB
Ø3	CAN	EM	SUB	ESC	FS	GS	RS	US
Ø4		!	"	#	$	%	&	´
Ø5	(	)	*	+	,	-	.	/
Ø6	Ø	1	2	3	4	5	6	7
Ø7	8	9	:	;	<	=	>	?
1Ø	@	A	B	C	D	E	F	G
11	H	I	J	K	L	M	N	O
12	P	Q	R	S	T	U	V	W
13	X	Y	Z	[	\	]	^	_
14	`	a	b	c	d	e	f	g
15	h	i	j	k	ℓ	m	n	o
16	p	q	r	s	t	u	v	w
17	x	y	z	{	¦	}	∼	DEL

Figure G: ASCII Code (in octal)

Section 3: MESSAGE STANDARDS

Sometimes the following standard escape convention is used to represent any character by 6-bit bytes: the 6-bit code (octal) 4Ø is recognized as an escape, rather than as "@"; any character is then represented with three 6-bit bytes obtained by adding (octal) 4ØØ2ØØ to its ASCII code, except that the codes of the upper and lower case letters are interchanged.

The ASCII code is summarized in Figure G. It consists of two kinds of characters: codes (octal) 4Ø - 176 comprise 95 graphic characters (including the space), each of which (in a typical output device) causes printing followed by advance of the printing position to the right. Codes (octal) Ø - 37 and 177 comprise 33 control characters, which cause various actions. These codes are now discussed in further detail (all character codes in octal).

The graphics are as follows:

1. The space (4Ø), which is unique among graphics in that its printed form is entirely blank;

2. The digits, 0 through 9 (6Ø - 71);

3. The upper case letters, A through Z (1Ø1 - 132);

4. The lower case letters, a through z (141 - 172); and

5. Special symbols and punctuation marks, as follows:

Section 3: MESSAGE STANDARDS

Exclamation point	!	(41)
Double quote	"	(42)
Number	#	(43)
Dollar	$	(44)
Percent	%	(45)
Ampersand	&	(46)
Apostrophe or close single quote	'	(47)
Parentheses	( )	(50 and 51)
Asterisk	*	(52)
Plus	+	(53)
Comma	,	(54)
Minus	-	(55)
Period	.	(56)
Slant	/	(57)
Colon	:	(72)
Semicolon	;	(73)
Less than	<	(74)
Equal	=	(75)
Greater than	>	(76)
Interrogation mark	?	(77)
Commercial at	@	(1ØØ)
Brackets	[ ]	(133 and 135)
Reverse slant	\	(134)
Caret	^	(136)
Underline	_	(137)
Open single quote	`	(140)
Braces	{ }	(173 and 175)
Vertical	¦	(174)
Overline	~	(176)

Section 3: MESSAGE STANDARDS

3.3 Control Characters

The control characters are each identified by a two or three letter mnemonic. They may be conveniently divided into six categories: format effectors (FE), device controls (DC), information separators (IS), deletion sentinels (DS), code extenders (CE), and communication controls (CC). These are now discussed in turn.

1. The purpose of format effectors is to control the layout of graphic characters on the two dimensional surface of a page, altering the natural progression to the right. It must, of course, be realized that not all printing devices can properly respond to all these controls.

Format Effectors

Backspace	(BS, 1Ø)	moves the printing position one to the left.
Horizontal tabulation	(HT, 11)	advances the printing position to the right to the next of a predetermined set of positions.
Line feed	(LF, 12)	moves the printing position one downward.
Vertical tabulation	(VT, 13)	moves the printing position downward to the next of a predetermined set of positions.
Form feed	(FF, 14)	moves the printing position to the top left of the (next) page.
Carriage return	(CR, 15)	moves the printing position to the extreme left of the current line.

Section 3: MESSAGE STANDARDS

2. Device controls are intended to operate ancillary devices associated with an output medium. Bell (BEL, Ø7) is intended to activate some form of attention-arousing device. The remaining device controls 1 through 4 (DC1, DC2, DC3, DC4, 21-24) have no standardized purpose.

3. Information separators have the purpose of dividing a message into logical subunits. Thus, they have meaning similar to that often given to such graphics as space, comma, semicolon, parentheses, brackets, and braces or to such indicators as end-of-record and end-of-file marks on magnetic tape. In their intended hierarchical order, from most to least inclusive, they are file separator (FS, 34), group separator (GS, 35), record separator (RS, 36), and unit separator (US, 37).

4. Deletion sentinels are a more or less miscellaneous category of controls having in one way or another to do with omissions that have been, or should be, made from a message.

Deletion Sentinels

Null	(NUL, ØØ)	is the "all zeros" character. It is what would naturally be expected to be "received" when nothing is being sent or to be found in a previously unused storage medium. For some media (e.g., paper tape), it is indistinguishable from the gaps between characters. It is intended to have no meaning, to be ignored, unless some exceptional condition can be inferred from its occurrence.

The same view is taken of:

Section 3: MESSAGE STANDARDS

Delete	(DEL, 177)	the "all ones" character. In some nonerasable media (e.g., paper tape), an erroneous character can therefore effectively be converted to a null by making it all ones.
Cancel	(CAN, 3Ø)	is supposed to imply that some portion of the message is to be discarded, presumably because of error.
End of medium	(EM, 31)	implies that the end of the meaningful information in a medium has been reached and subsequent characters should be discarded; thus, it is similar in intent to an end-of-tape mark.
Substitute	(SUB, 32)	is intended to be substituted for a character known to be in error; for example, an intermediate node in a transmission might use this character to replace characters received with bad parity.

5. Code extenders provide for adding more characters (except communication controls) to the code. Shift out (SO, 16) indicates that all following characters (except communication controls) are to be interpreted according to some alternate scheme until shift in (SI, 17) is encountered. Escape (ESC, 33) combines with some number of immediately subsequent characters (except communication controls) and the combination is interpreted as a single new character.

6. The remaining characters, the communication controls, are not intended to appear as part of a message but are to be inserted and removed, if necessary, by the network nodes in order to control the transmission activity on a channel. Briefly, they are:

Section 3: MESSAGE STANDARDS

Communication Controls

Start of heading	(SOH, Ø1)	
Start of text	(STX, Ø2)	
End of text	(ETX, Ø3)	for logical blocking of the message.
End of transmission	(EOT, Ø4)	for ending a transmission.
Enquiry	(ENQ, Ø5)	
Acknowledge	(ACK, Ø6)	
Negative acknowledge	(NAK, 25)	for initial "handshaking" in setting up a transmission.
Synchronous idle	(SYN, 26)	for holding synchronization on a synchronous channel.
End transmission block	(ETB, 27)	for dividing a transmission into blocks.
Data link escape	(DLE, 2Ø)	to provide for more communication controls.

The meanings of the control characters, of course, may be altered to fit the special needs of the various Octopus facilities, just as the printed form of the graphics must adapt to the capabilities of each printing device.

Section 4: THE USER

An employee desiring to become an authorized user of the Octopus network must make a specific request to Computation Department. A user is known to Octopus by a six digit user number, which is the same as his employee number; this number is used whenever the user needs to identify himself to the network. Each worker (and the Elephant) computer which a user is authorized to use keeps certain information about the user, which is now briefly summarized.

On most computers, the user has a *bank account*. As the user consumes resources of the computer (e.g., processor or I/O time, memory or storage space), the bank account is decremented according to a rate schedule. At specified intervals (usually once or twice a day) the bank account is replenished. The user is not permitted to use the computer whenever his account is exhausted.

Each computer keeps track of the files in secondary storage to which a user has access. In most cases, these files include *private* files accessible only to one user, *shared* or pooled files accessible to a group of users, and *public* files accessible to all users.

Each user is identified with a maximum security level at which he is permitted to operate. These levels are identified by a single letter mnemonic and include (lowest to highest) unclassified (U), protected as restricted data (P), administrative (A), confidential (C), and secret (S). Whenever a user interacts with a computer, he selects an *operating level* which may not exceed his maximum. Each file and I/O device is identified with a level, and a user's access to them is restricted according to his operating level; details are given in LTSS Chapter 4 and in connection with discussions of particular I/O facilities.

Section 4: THE USER

Each user has a *combination*, a sequence of six letters known only to him and a few security personnel. He is forbidden to write this combination down or to reveal it to others. If a computer requires the user to verify who he claims to be (when he supplies his user number), he must supply his combination; such verification is always required when operating at a classified level (confidential or higher).

REFERENCES

1. P. DuBois, J.G. Fletcher and M. Richards, *Livermore Time-Sharing System: Chapter 4, Files,* Lawrence Radiation Laboratory, Livermore, California, Computer Information Center, (October, 1970).

2. J.G. Fletcher and J.D. Lawrence, *Livermore Time-Sharing System: Chapter 3, Teletype,* Lawrence Radiation Laboratory, Livermore, California, Computer Information Center, (May 5, 1969).

INDEX

Bank account, 31

Character set:
  ASCII, 23
  escape convention, 25
  six-bit ASCII subset, 23
Combination, 32
Concentrators, 6
Control characters:
  code extenders, 29
    Escape, 29
    Shift In, 29
    Shift Out, 29
  communication controls, 29
    Acknowledge, 30
    Data Link Escape, 30
    End Of Text, 30
    End Of Transmission, 30
    End Of Transmission Block, 30
    Enquiry, 30
    Negative Acknowledge, 30
    Start Of Heading, 30
    Start Of Text, 30
    Synchronous Idle, 30
  deletion sentinels, 28
    Cancel, 29
    Delete, 29
    End Of Medium, 29
    Null, 28
    Substitute, 29
  description of, 27
  device controls, 28
    Bell, 28
  format effectors, 27
    Backspace, 27

Control characters:
  format effectors *(continued)*
    Carriage Return, 27
    Form Feed, 27
    Horizontal Tabulation, 27
    Line Feed, 27
    Vertical Tabulation, 27
  information separators, 28
    File Separator, 28
    Group Separator, 28
    Record Separator, 28
    Unit Separator, 28

Data acquisition subnetwork:
  definition of, 14
  terminal computers, 14

Elephant computer, 6
Employee number, 31

Files:
  pooled, 31
  private, 31
  shared, 31
  public, 31
File transport network:
  channels, 9
  information burst, 10
  initializing a transport, 9
transport success, 9

Headings:
  Binary bit, 21
  Bye bit, 21
  Chain bit, 23

Headings: *(continued)*
- Computer field, 19
- computer number, 21
- Control field, 23
- Device field, 19
- device number, 21
- Error bit, 21
- ID bit, 21
- messages, 19
- overview of, 19
- sink, 19
- Sink field, 19
- source, 19
- Source field, 19
- text, 19
- Type field, 23

Input-Output devices, 4

Memory:
- address, 4
- definition of, 3
- word, 3
- word size, 3

Message network, 10

Network interfaces, 5

Octopus network, 1, 7

Operating level, 31

Processor:
- definition of, 3
- instruction counter, 3
- instruction word, 3
- program counter, 3
- registers, 3

Remote I/O subnetwork, 14

Secondary storage:
- definition of, 4
- disks, 4
- drums, 4

Security levels:
- administrative, 31
- confidential, 31
- protect as restricted data, 31
- secret, 31
- unclassified, 31

Teletypewriter subnetwork:
- description of, 10
- teletypewriter communication, 12
- tentacles, 12

Terminal computer, 6

Transmission protocol:
- definition of, 14
- error status bits, 16
- Receive Permitted bit, 17
- Receive Requested bit, 16
- Send Permitted bit, 16
- Send Requested bit, 16
- status bits, 16

User number, 31

Worker computer:
- definition of, 5
- designator, 5
- time shared, 5

3/71

The LRL secured operating system is available to users. The system permits users the option of assigning levels of protection to files being manipulated within the OCTOPUS system. The secured system will in no way hinder the user who wishes to operate at "Protect-As-Restricted Data" (PARD).

1. COMBINATIONS Combinations will be made available to users requesting same. A request form may be obtained at the print out room located in the central computer facility Bldg. 113. The combination must be used for levels of protection greater than PARD. Users that do not wish to use a combination for PARD, but require it for classified runs should check the appropriate box on the "Computer Combination Request Form".

2. LOGGING ON

   ID(GLMRS)$\wedge$USER#$\wedge$SUFFIX$\wedge$ACC#$\wedge$PL$\wedge$COMBO

   PL (Protection Level): = U (unclassified)
   = P (PARD)
   = A (ADP or Official Use Only)
   = S (Secret)
   = K (System)
   = Blank (same as P)

   COMBO: A 6 alphabetic character string. Hold down "Red Button" (upper left keyboard) while typing your combination. This action will result in % being echoed. (In the future the "Red Button" will not be used). For G machine, hold down 'Control Key' while typing your combination.

3. CARD FILE ID

Col 1	Col 4	11	30	35	40	51-72
ID	User#	Filename	Mode	PL	Size	Ident.

   Mode = M, S, A, or BLANK

   PL = U, P, A, S, K or BLANK (BLANK is same as P)

4. FILE HANDLING WITHIN A GIVEN MACHINE

There will be three categories in the file structure:

1. Private
2. Need-to-know pools
3. Public Shared

If a user wishes to access a file, the system first searches his private index. If that fails, the next search is the Pool index, if he has one, and lastly the Public index.

Write Access

A user will be allowed write access only to any files in his private index, or the Public index provided the file has a read/write mode access. The user will be denied write access to all files in the Pool index. (Only one combo per pool can have write access to that pool.)

Read Access

A user will be denied read access to all files (regardless of categories 1, 2, and 3 above) unless these files have a PL equal to or below the operating PL of the user.

Open Access

A user will be allowed to open any file in his private index. A user will be denied access to any file of a PL higher than his operating PL in the Pool or Public indices.

Create

A user will be allowed to create a file in his private index only. The PL of the file will be equal to his operating PL unless he states otherwise. A user may create files of any PL if so stated.

Destroy

A user will be allowed to destroy a file in his private index only. This file may be of any PL. Files with a PL higher than P will be written over with disk pattern.

Give

A user may be allowed to give a file of a PL higher than (PARD) only to his Pool index (provided no name duplication exists) and to a special system user number (999999). The latter is for off line processing only.

NETWORK FILE HANDLING

File Transport - Host

From PDP-6

If Create:

Host will create file PL = source file PL given in message.

If Not Create (i. e., subsequent bursts):

Check that host file PL is ≥ source file PL given in message. (Source file PL is level of information on PDP-6 side.)

To PDP-6

Check that Host file PL being accessed is ≥ destination file PL given in message.

Return Messages to PDP-6 From Host:

Contain PL of host file involved (even error message rejecting PL).

Note: PDP-6 will send two PL's codes: A user operating PL and a file PL.

5. INPUT OUTPUT HANDLING

USER 1 Output

1. High speed output files will be appropriately labeled at the top and bottom of each page and on the identification pages if the PL of the output is A, S, or K.

2. No output that has a PL greater than (P. A. R. D.) will be permitted to go to an online printer.

3. A log is maintained of files (higher than PARD) going to HSP tape and Dd80 plot tape and made available to the operator when these tapes are removed from the machine.

TAPES

When a user requests a tape to be mounted, the system informs the operator of the operating PL of the user.

6. ELEPHANT FILES

Security requirements in regard to Elephant files are of the same general nature as those relating to other files. Details are in LTSS Manual Chapter 4, Section 4.1.11.

7. NEW LIBRARY ROUTINES (LRLLIB 7600-6600)

Call SETLVL (L U, LVL) after "open call"
Call GETLVL (LU, LVL) before "call assign"

LU = Logical Unit
LVL = PL : U = 1, P = 2, A = 3, S = 5, K = 7

Note: Use decimal number equivalent for LVL

TMDS (PARD or Unclassified Operation)

1. Each console will have a label attached (DATA DISPLAYED HEREON PROTECT AS RESTRICTED DATA.) This label is a reminder.

2. To insure that the console you have reserved is the same console that you are looking at, TMDS will RESERVE, ERASE, and display your user number and a message describing protection level. This message is a reminder.

   Therefore, if you do not see your user number, you know that the console number was in error and that you should not send data.

   The classification message is a reminder of the program's protection Level.

3. TMDS ignores all other data in the Packet/Block which contains the reserve command.

   This insures that the user number will appear.

4. TMDS will ERASE and RELEASE whenever a release command is sent.

5. TMDS will ERASE and RELEASE whenever 15 minutes have elapsed since the last datum sent.

   This provides that the console will become available should the user's abandon the console or should the user's worker computer go down.

6. TMDS has a manual erase button which can be used to erase all consoles.

   This button will be used in the event of extended PDP-6 failure.

   Users who need to erase data but are unable to via software (due to some system failure) should alert the PDP-6 personnel, or the Head Operator.

TMDS (PL's Higher than PARD)

The following modifications (additions) are proposed to protect classified material.

1. The first frame of information will contain the user number, the classification message and random sequence of characters.

   Before secret restricted data is displayed TMDS will verify that the user is indeed watching the console that he reserved; the user must type in this random sequence of characters within one minute or the console will be released.

2. A routine will be provided to permit the user to easily display the label "Secret Restricted Data" at the bottom of the screen. It is the user's responsibility to display this label, or to display a placard with this label.

3. Only those consoles whose locations are cleared by Security to receive secret data will be permitted by TMDS to display secret restricted data.

4. TMDS will automatically ERASE and RELEASE a console whenever 5 minutes have elapsed since the last datum sent.

5. The user may not link consoles (display the information simultaneously on more than one console.)

Please refer to LTSS Chapters 9 and 10 for details concerning System Calls. Further questions should be directed to S. O. S., X 8582, Bldg. 113.

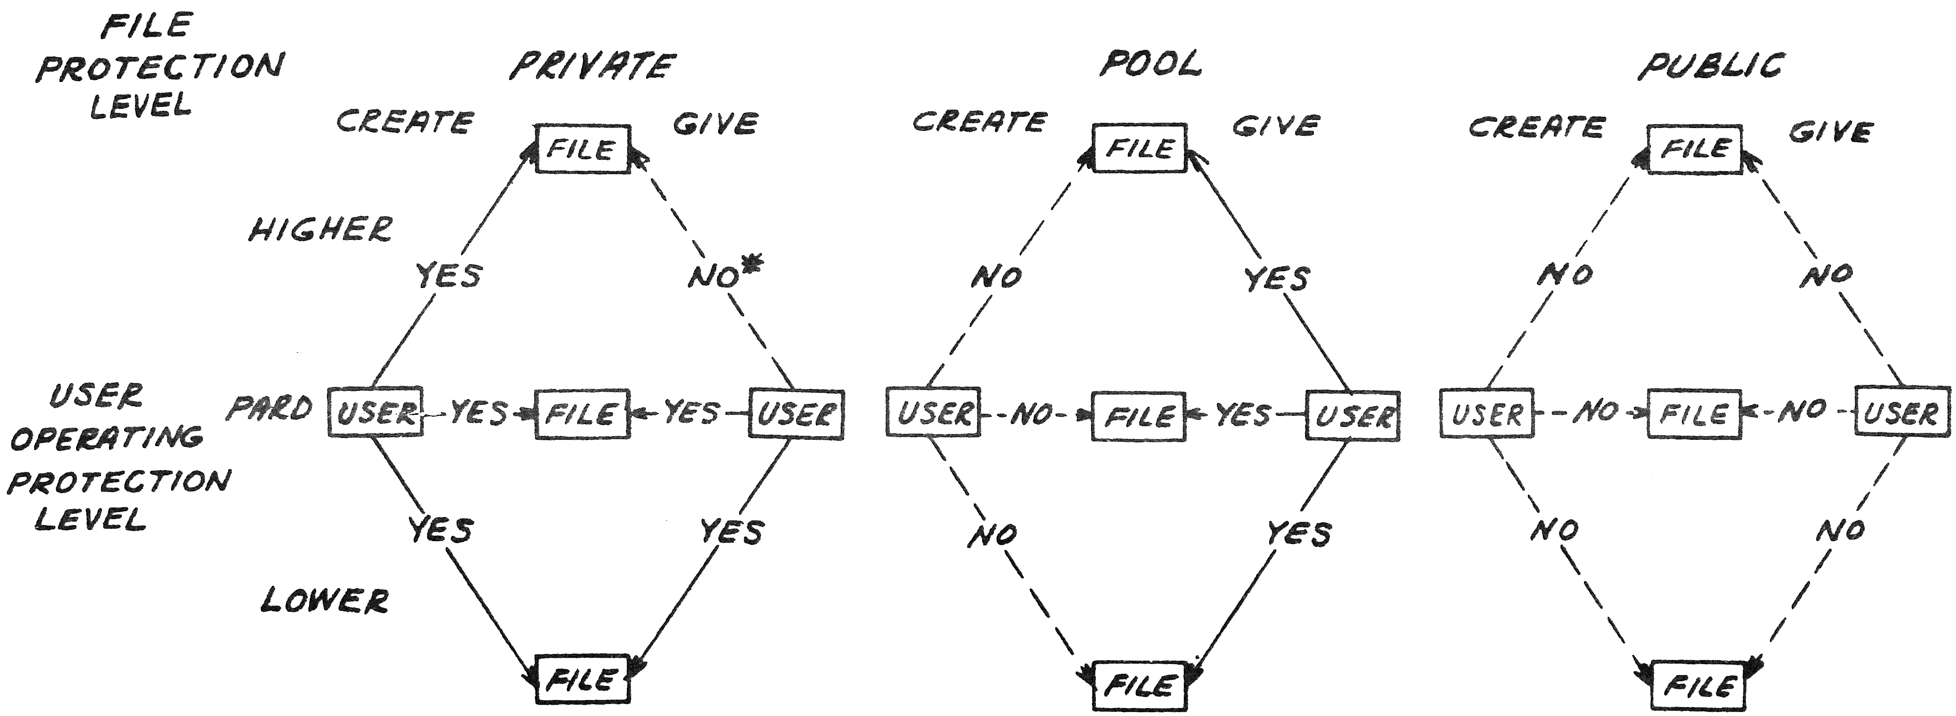

* USER MAY GIVE FILE ABOVE PARD

TO A SPECIAL PRIVATE

USER FOR OUTPUT PROCESSING

CREATE, GIVE RIGHTS/ PROTECTION LEVEL

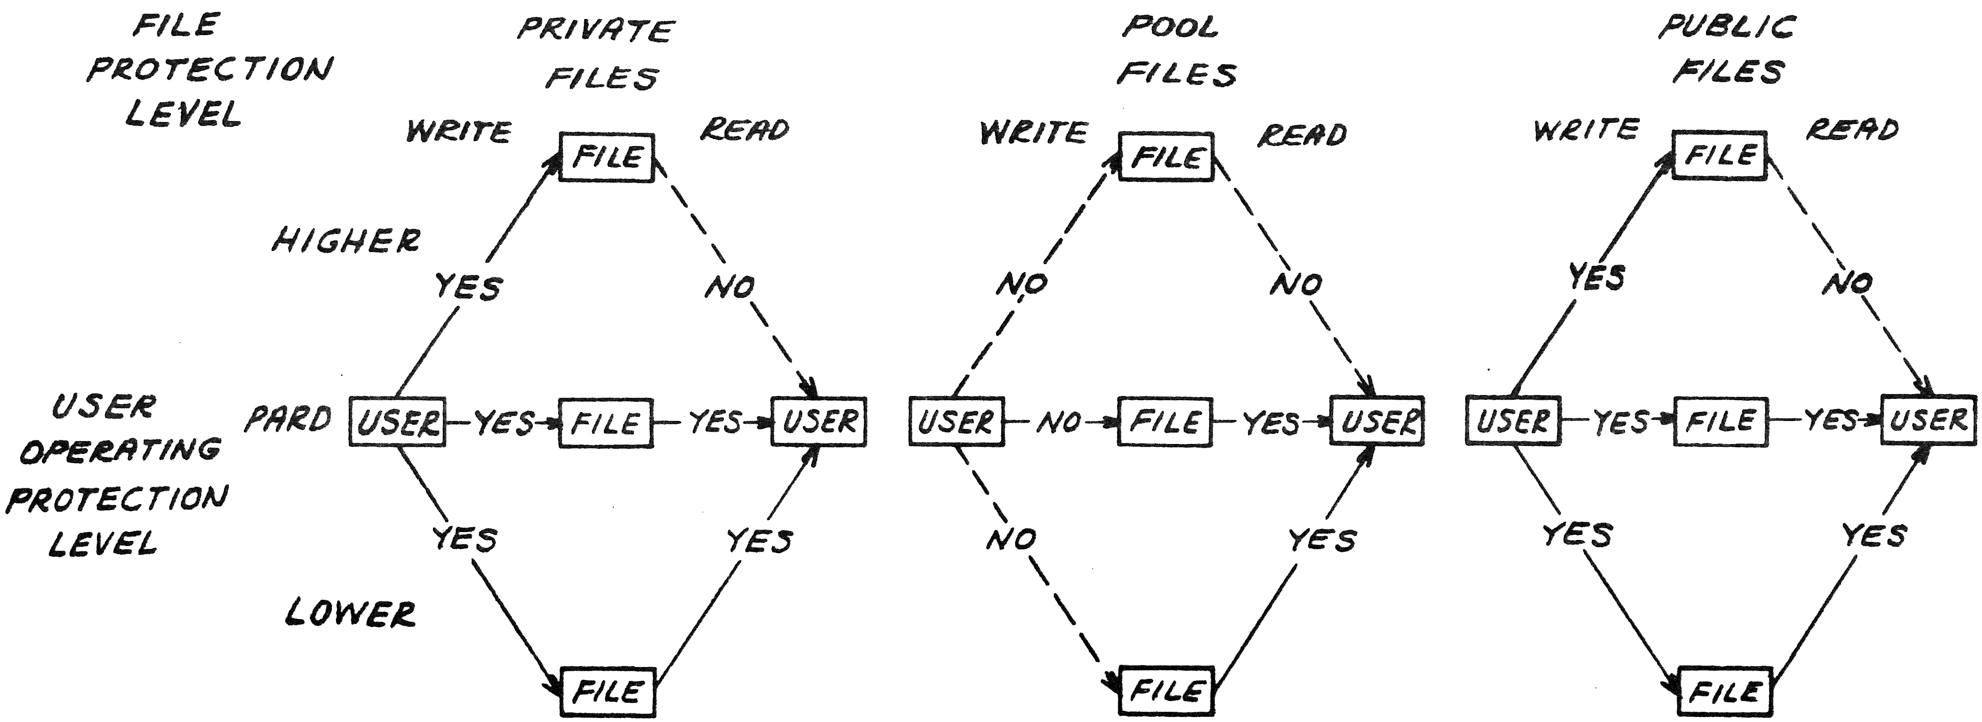

READ, WRITE ACCESS / PROTECTION LEVEL

CMR 3-71