AD-773 831

AN INTERFACE TO THE ARPA NETWORK FOR THE CP/CMS TIME-SHARING SYSTEM. VOLUME 1

Joel M. Winett, et al

Massachusetts Institute of Technology

Prepared for:

Electronic Systems Division Advanced Research Projects Agency

28 November 1973

**DISTRIBUTED BY:** 

National Technical Information Service U. S. DEPARTMENT OF COMMERCE 5285 Port Royal Road, Springfield Va. 22151

# Best Available Copy

UNCLASSIFIED Security Classification

Security Classification	AD77	3 83/
DOCUMENT CONTR	ROL DATA - R&D	
(Security classification of Atle, body of abatract and indexing at 1. CRISING ACTIVITY (Comparing mithon)	notation must be entered when the overall report in ca	NUMBER
	Unclassified	
Lecole Laboratory, M.I.T.	25. GROUP	
J ARPORT TITLE		
An Interface to the ARPA Network for the CP/CMS Tim	e-Sharing System – Volume 1	
4. DESCRIPTIVE NOTES (Type of report and inclusive dates) Fechnical Note		
5. AUTHOR(5) (Lest name, list name, initial)		
Winett, Joel M. and Sammes, Anthony J.		
5. REPORT DATE	74. TOTAL NO. OF PAGES 76. NO. OF	PEPS
28 November 1973	116 20	
Se. CONTRACT OF GRANT NO. F19628-73-C-0002	94. ORIGINATOR'S REPORT NUMBERIS	
APPA Order 600	Technical Note 1973-50 - Vol. I	
5. PROJECT NO. ANTA OLGET GOU	95. OTHER REPORT NO(5) (Any other number essigned this report)	re that may be
с. d.	ESD-TR-73-336	
10. AVAILABILITY/LIMITATION NOTICES		
Approved for public release; distribution unlimited.		
II. SUPPLEMENTARY NOTES	12. SPONSORING MILITARY ACTIVITY	
Supplement to ESD-TR-73-337	Advanced Research Projects Ager Department of Defense	icy,
13. ABSTRACT		
The interface to the ARPA network consists of Host to Host communications, a LOGGER/SERVER the network, and a user subroutine package for com a logged on CP virtual machine operating in the CM run in separate virtual machines; the NCP handling handling pseudo I/O operations with CP through a so virtual machines communicate with the NCP virtual virtual machine communications facility.	a Network Control Program (NCP) which ha for providing access to a CP virtual machine imunicating with other Hosts on the network IS environment. The NCP and the LOGGER the I/O operations with the IMP and the LOG oftware supported virtual terminal device. Of I machine through a special virtual machine	ndles from from each GER CMS to
This report describes the routines which make flow charts on the logic for each of these routines.	up the NCP and the LOGGER. Volume 11 inc	ludes
Reproduced by NATIONAL INFORMAT U & Departm Springfiel	TECHNICAL ION SERVIUE and of Comme ca id va 22151	
14. KEY WORDS		
ARPA network time-sharing system	IBM 360/67 victual machines	
	1.	

UNCLASSIFIED Security Classification

#### MASSACHUSETTS INSTITUTE OF TECHNOLOGY

LINCOLN LABORATORY

## AN INTERFACE TO THE ARPA NETWORK FOR THE CP/CMS TIME-SHARING SYSTEM VOLUME I

J. M. WINETT A. J. SAMMES Group 23

TECHNICAL NOTE 1973-50

28 NOVEMBER 1973

FFR

Approved for public release, distribution unlimited.

LEXINGTON

ii.

MASSACHUSETTS

The work reported in this document was performed at Lincoln Laboratory, a center for research operated by Massachusetts Institute of Technology. This work was sponsored by the Advanced Research Projects Agency of the Department of Defense under Air Force Contract F19628-7(4-C-0230 (ARPA Order 600),

This report may be reproduced to satisfy needs of U.S. Government agencies.

### Table of Contents

#### Volume I

1. Introduction	1
2. The ARPA Network	1
2.1 General Description	1
2.2 Network Concepts	5
2.3 Requirements of the HOST Software	8
2.4 POST-to-INP Protocol	ų
2.5 HOST-to-HCST Protocol and the NCP	13
2.6 Process-to-Process Protoccls	16
3. The CP-/CMS Time-Sharing System	17
4. The Lincoln Laboratory IBM 360/67 Network Software	19
4.1 System Lesign	21
4.2 The NET Package	23
4.3 The logger-Server	33
4.4 The Network Control Program	37
4.4.1 NCPMAIN	39
4.4.2 NCPXFEP	41
4.4.3 NCPIMPI	44
4.4.4 NCEIMPO	46
4.4.5 NCPMONIT	47
4.4.6 NCPINIT and Error Recovery	48
4.5 The Treatment of Errors	49
4.6 Concluding Remarks	50

v

# Preceding page blank

Appendices

A.	IMF Interface Specifications	54
в.	HCST-to-IMP Messages	56
•••	IME-to-HOST Messages	58
D.	The HOST-to-HOST Protocol	61
Ē.	CF Support for a Virtual Terminal Device	66
F.	CF Virtual Machine Communications Facility	69
G.	The NET Boutine	71
H.	NET User Macros	74
I.	The MLCIST Program	76
Ј.	The IRYNEI Program	77
К.	The TELNET LOGGER/SERVER	79
L.	The User TELNET	86

Volume II: Flowcharts

A. NCPINIT Flowchart
E. NCPMAIN Flowchart
C. NCPXFEF Flowchart
D. NCFIMPI Flowchart
F. NCPIMPO Flowchart
F. NCFMONIT Flowchart
G. LOGMAIN Flowchart
H. LOGDFV Flowchart

I. LOGST Flowchart

# List of Figures

Figure 1	Network Topology - Fall 1971	2
Figure 2	ARPA NET Logical Map - Fall 1971	3
Figure 3	Hosts on the ARPA Network	4
Figure 4	HCST-to-IMP Message format	10
Figure 5	HCSI-to-HOST Message Format	12
Figure 6	HOST-to-HOST Control Commands	14
Figure 7	System Lesign	22
Figure 8	NET Fackage	24
Figure 9	NET to NCP Control Block Format	26
Figure 10	NET Command Format	26
Figure 11	Interrupt Stack Format	26
Figure 12	NET Commands	28
Figure 13	NCP to NET Interrupt Codes	28
Figure 14	Pormat of Status Bits	30
Figure 15	NET Fleg Bits	32
Figure 16	LOGGEF Foutine Flags	34
Figure 17	The Network Control Program	38
Figure 18	Port Table Format	40
Figure 19	Polling and Device Flags	40
Figure 20	Input and Output Euffer Formats	40
Figure 21	Accounting Fecord Format	42
Figure 22	RV Table Format	42
Figur∈ 23	NCP Errcr Codes	51
Eigure 24	Major Interfaces in the System	52

10.00

ainete.

**16 18** 

# List of Figures

(Continued)

Figure B.1	IMP-to-HOST Message Pormat	56
Figure C.1	HOST-to-HOST Message Format	58
Figure G.1	Interrupt Stack Control Block	73
Figure G.2	Interrupt Codes	73
Figure K.1	LOGGER ASCII/EBCDIC Code Mapping	82
Figure L.1	Serving Hosts on the ARPA Network	94
Figure L.2	The EBCLIC Code	95
Figure 1.3	The USASCIT Code	96
Figure L.4	TELNET ASCII/EBCDIC Code Mapping	97
Figure L.5	2741 Keyboard Control Character Mappings	10 1
Figure L.6	2741 Character and Control Codes	10 3
Figure L.7	Code for Characters on a PN Print Train	104
Figure L.8	Code for Characters on a TN Print Chain	105

.

#### An Interface to the ARPA Network

for the CP/CMS Time-Sharing System

#### 1. Introduction

The interface to the ARPA network consists of a Network Control Program (NCP) which handles Host to Host communications, a LOGGUR/SERVER for providing access to a CP virtual machine from the network, and a user subroutine package for communicating with other Posts on the network from a logged on CP virtual machine operating in the CMS environment. The NCP and the LOGGER each run in separate virtual machines; the NCP handling the I/C operations with the IMP and the LOGGER handling pseudo I/O operations with CP through a software supported virtual terminal device. CMS virtual machine through a special virtual machine to virtual machine communications facility.

All routines were written in assembly language for the IBM 360/67 system and operate with CMS in a CP virtual machine environment.

A brief description of the ARPA network is given, followed by an outline of the CP/CMS time-sharing system. This provides the fram work for a more detailed consideration of the interface software and the automatic initialization and error recovery features that are incorporated in the system design.

#### 2. The ARPA Network

2.1 General Description

The ARPA Network is a store and forward, message processing network which interconnects a number of large, multi-processing computer systems throughout the United States. The computer systems, known as HOSTs, are linked together for the purpose of sharing resources among member sites. (See ROBERTS & WESSLER, Ref. 17). At each node of the network is an Interface Message Processor or IMF which controls all traffic through the node. (See HEART, KAHN, ORNSTEIN, CROWTHER & WALDEN, Ref. 9.) IMPs are interconnected by wide-band (50 kilobits per second) leased communication lines and provide for the connection to the network of up to four HOSTS at each site. The IMP is a slightly modified H =ywell DDP-516 computer and has been developed by BOLT, BERANEK & NEWMAN (Ref. 1).

The topology of the network, at the time the system was implemented, was as shown in Fig. 1 and Fig. 2. The computer systems connected were as listed in the table of Fig. 3.



Fig 1: Network Topology - Fall 1971



Fig. 2: ARPA Net Logical Map - Fall 1971

TN73-50-I(3)

100 -in orbeits

SITE	I HOST CCMPUTERS	NETWORK IDENTIFIER
UCLA	Sigma 7   IBM 360/91	01
SRI	PDP 10	02
UCSB	IBM 360/75	03
UTAH	PDP 10	04
BBN I	DDP 516 PDP 10 PDP 10	05   45   85
MIT	GE 645 PDP 10	06 46
RAND	IBM 1800	07
SDC	IEM 360/67	08
HAFVARD	PDP 10	09
I LINCOLN	IBM 360/67 TX-2 TSP	0 A 4 A 8 A
STANFORD	PDP 10	08
ILLINOIS	PDP 11	0c
CASE	PDP 10	0 D
CARNEGIE	PDP 10	0 E
PAOLI	Burroughs 6500	OF

Fig. 3: Hosts on the ARPA Network - July 1971

4

Since that time, the network has been expanded, and a more recent status report is given in KIRSTEIN (Ref. 10).

2.2 Network Concepts

The HOST computers typically differ from one another in type, speed, word length and operating system, though the notion of multiprogramming dominates most of the systems. Such HOSTs could concurrently support several users, with each user running one or more processes. A 'process', in this context, is taken to mean any software entity in a HOST with which an instruction counter is associated, and includes each concurrent activation of a re-entrant program. It corresponds closely to the term 'activity' used by SPOONER (Ref. 18).

A fundamental requirement of the design is to provide for process-to-process communication over the network. It was decided that such communication should be based, not on a solitary message, but rather upon a sequence of messages and that this sequence should conform to some standard formatting procedure or 'protocol'. A number of such process-to-process protocols have been defined and included:

- a. An Initial Connection Protocol (ICP) which enables any process to gain access to specific processes (such as the system logger) at another HOST.
- b. A Telecommunication Network (TELNET) protocol which enables any terminal serving process to be connected to any remote keyboard-printer terminal by mapping through a Network Virtual Terminal (NVT).
- c. A Data Transfer protocol which specifies standard methods of formatting data for transmission through the network.
- d. A File "mansfer protoce" which specifies methods for reading, writing and updating files sorted at a remote BOST.
- e. A Graphics protocol which facilitates the transfer of graphics display information.

Processes which desire to communicate with one another have to be coupled by one or more 'connections'. A connection is defined to be simplex (uni-directional) and connects an output port of one process to an input port of another. Once such a connection has been established, messages passed to the output port of one process are automatically received at the input port of the other. Difficulties exist in uniquely identifying ports (network-wide) in that each HOST, typically, has a aifferent internal scheme for naming its processes. For this reason, a standard, intermediate name space is used, with a separate portion of the name space allocated to each HOST. The elements of the name space are called 'sockets' and each HOST is responsible for mapping sockets to ports for its own internal processes. A socket is specified by a 40-bit number, the first 8 bits of which identified the HOST. In the Lincoln Laboratory implementation, a further 24 bits are used to identify the internal process and the reamining 8 bits, the port within the process. The 'gender' of the socket (that is, whether it is send or receive) is marked by the last bit, thus giving each process a maximum of 128 send and 128 receive sockets.

Although it would be possible for processes to establish connections on their own behalf, it is considered desirable that control be co-ordinated in a standard manner, throughout the network, by a special purpose Network Control Program (NCP) resident in each HOST. Processes within a HOST then communicate with the rest of the network through their local NCP, which establishes, breaks, and controls data flow over connections as required.

In order to accomplish its tasks, the NCP of one HOST is required to communicate, in a well-defined manner, with the NCPs of other HOSTs. To this end a HOST-to-HOST protocol was defined by the Network Working Group (NWG), which specifies methods of establishing and breaking communications raths, the managing of buffer space, and a method of providing interrupt facilities. This protocol is known as the 'second-level protocol' and is transparent to the higher level process-to-process protocols outlined previously.

The medium for communication between HOSTS is the 'link'. Every HOST is considered to have 256 receive links from every other HOST, which the NCP could assign to connections as required. Links are permanently available to the NCT and are specified by the 8-bit HOST identifiers concatenated with an 8-bit link number. Link number 0 is used throughout the network as the 'control link', over which NCPs can issue 'control messages' to one another in accordance with the standards specified by the second level protocol. Control messages are always interpreted by an NCP as a sequence of one of more 'control commands' to effect such activities as the initiation of a connections, the assignment of buffer space, or the notification of an interrupt.

Each NCP maintains a 'rendezvous' (RV) table which records the current assignment of sockets to internal pr cesses and links to connections. Incowing messages on a link, other than link 0, are passed by the NCP to that port of that internal process specified by an appropriate entry in the RV table. If no such entry exists, the message is rejected. Outgoing messages, from in internal process, are associated with the appropriate link number, obtained from the RV table, before being transmitted to the appropriate HOST.

In order to effect any transmission of messages, via the network, an NCP has to communicate with its local IMP. A well defined procedure, known as the HOST-to-IMP or 'first level' protocol, has been specified for this purpose. Two types of messages are permitted by this protocol:

- a. 'Regular' messages which are always associated with a link and which are passed by the IMPs, through the network, to a destination HOST. Such messages provide the mechanism for all HOST to HOST communication.
- b. 'Irregular' messages which are communications between a HOST and its local IMP. Such messages are not, in general, passed through the network to a distant HOST, but are used for local control purposes.

The protocol is implemented through the medium of a 32-bit 'leader', appended to all messages. Fields in the leader specify the message type, the destination or source HOST and the link number. Regular messages are of message type 0, and irregular messages are of type non-zero, the particular value of the type defining the function of the irregular message. The protocol is transparent to the higher level HOST-to-HOST and process-to-process protocols referred to above.

Any regular message, passed by a HOST to its local IMP, is split, by that IMP, into one or more 'packets' for transmission through the network. Each packet is preceded by a 64-bit routing 'header' which contains such information as the destination and source HOSTs, the link number, the message and packet numbers and a set of flags.

A packet is passed, from IMP to IMP, through the network until it reaches its destination, as specified in its header. Each intermediate IMP selects the next leg of the transmission path, for the packet, on the basis of network topology, local tr fic density and current line qualities. The destination IMP reconstructs the message from the individual packets, re-ordering them as necessary in accordance with the packet numbers. Such re-ordering is required because of the potentially different routes, and hence different transmission times, taken by each packet. When the messages have been reconstructed, the destination IMP passes it to the destination HOST and transmits an acknowledgement message to the source IMP. This acknowledgement message is known as a Ready For Next Message (PFNM). On

WILL WARRAND STATE

receipt of the RFNM, the source IMP advises the source HOST that the message had been passed and that the link associated with it is now free for further transmissions. The packet mechanism and the IMP to IMP communication procedures are transparent to all the higher level protocols.

At the hardware level, a HOST is connected to its IMP by a specially designed interface, across which information can be transmitted, bit serially, in either direction. It is this interface which resolved the differences in word lengths and speeds between the HOSTs and IMPs. Transmissions consist of a bit train, containing no special indication of word or byte boundaries, but interrupted as required while the sender fetches a new word from memory, or the receiver writes an assembled cre into storage.

The Lincoln Laboratory 360/67 is connected to the IMP through a specially built interface box. This tox is attached as a control unit on the multiplexer channel and provides a duplex connection to the IMP. Thus the IMP appears as two devices, one for transmission of data and one for the reception of data, both of which can operate simultaneously. The specifications of data transfer to and from the IMP through this interface box is described in Appendix A.

#### 2.3 Fequirements of the HOST Software

The network control software, resident in each HOST, is required to observe five major interfaces and effect the protocols associated with them. These are:

- a. The hardware interface between a HOST and its IMP, and the protocol which specifies the control of that interface.
- b. The software interface between a Network Control Program and its local IMP, and the NOST-to-IMP protocol.
- c. The software interface between Network Control Programs, and the HCST-to-HOST protocol.
- d. The software interface between and internal process and its local Network Control Program, and the protocol which describes the form of messages between them.
- e. The software interface between processes, and the various process-to-process protocols.

Of these protocols, the HOST-to-IMP, HOST-to-HOST and process-to-process are network defined; the remainder are implementation dependent and vary in detail from sive to site. An outline description of the network defined protocols follows; the implementation dependent ones are described in Section 4.

#### 2.4 HOST-to-IMP Protocol

The HOST-to-IMP protocol was specified by BOLT, BERANEK & NEWMAN (Ref. 2) who designed the IMP software and modified the Honeywell DDP-516 computer. This protocol is based on the concepts of r 'message' and a 'link'.

A message consists of not more than 8096 bits of data which includes 32 bits for a leader and at least one bit fcr "padding".

Padding is added by the hardware interfaces to resolve the word length mis-match problem of dissimilar HOSTs and IMPs. The source HOST to IMP interface appends a single one bit to the end of a message, followed by sufficient trailing zero bits to complete an integral number of IMP words. The destination IMP to HOST interface appends additional trailing zero bits to complete an integral number of destination HOST words. The end of the message can thus be located by a destination. HOST by searching its input buffer, back from the last word transferred, until it finds a non-zero bit. The bit pricr to this will be the last bit of the message sent by the source HCST. No network significance is to be placed on HOST word or byte boundaries.

Two classes of message type are specified, as stated previously; regular messages which have a zero message type and which provided the medium for all HOST to HOST communications; and irregular messages which have a non-zero message type and which are used for control purposes between a HOST and its IMF. (It should be noted that messages of non-zero message type are referred to in BOLT, BERANEK & NEWMAN (Ref. 2) as 'control messages'. This term is also used in the HOST-to-HOST protocol for regular messages on link zero. To avoid ambiguity, the term 'irregular message' is used throughout this report to refer to message' is reserved for HOST-to-HOST protocol use.)

The 32-bit leader determines the type of the message and the link, if any, with which it is to be associated (Fig. 4).

Regular messages are always to be associated with a link and this is fully specified by a destination HOST identifier, a source HOST identifier, and a link number. Of these, the originating HOST specifies, explicitly, only the destination HCST identifier in the HOST identifier field, and the link number. At its destination, the destination IMP regenerates the leader from the packet header, replacing the value originally in the HOST identifier field by the source HOST

	TN73-50-I(4)					• • • <u></u> _	
!	4	ų	8	8	8	variable	1
	FLAGS	MESSAGE TYPE	HOST IDENTIPIER	LINK NUMBER	NOT ASSIGNED	I DATA	
	k <		-32-Bit Leader		>	L • • •     	. 1

)   	HOST TO IMP MESSAGES	I IMP TO HOST MESSAGES		
MESSAGE MEANING		MESSAGE TYPE	MEANING	
1 0	Regular message HOST to HOST.	0	Regular message HOST to HOST.	
	Error in leader.	ï	Error in leader.	
2	HOST going down.	2	IMP going down.	
3	-	3	Elocked link.	
4	No operation.	ц	No operation,	
5	Regular message for discard.	5	RFNM.	
E	-	6	Link table full.	
7	-	7	Destination IMP or HOST dead.	
9	Erro, not in leader.	8	Error not in leader.	
<u>ġ</u>	-	9	Incomplete transmission.	
10-15	-	10-15	-	

Fig. 4: HOST-to-IMP Message Pormat

A STATE OF A STATE OF A STATE

identifier. Once a message had been assigned to a link, the link is said to be 'blocked' and the IMP will refuse further transmissions over it until the first message has either been delivered or rejected. Notification of these events is by irregular IMP to HOST messages; a type five indicating successful delivery (RFNM); a type seven, that the destination HOST did not receive the complete transmission. Each of these messages unblocks the link. Other irregular messages in the protocol enable an IMP to indicate to its BOSI that a link is blocked; that errors have been noted in previous messages; or that its link table is full. The HOST can indicate to its IMP that it is closing down; that errors have been noted in previous messages; or that a particular message is for test surposes only. The messages are summarized in Fig. 4 and described in more detail in Appendix B and C.

With reference to the leader, the PLAGS field, of four bits width, is set to zero by source HOSTs and ignored by destination HOSTs. The field is to be used only to communicate with IMP background programs for diagnostic, maintenance or measurement purposes. The fields, MESSAGE TYPE (four bits), HOST IDENTIFIER (eight bi\*s) and LINK NUMBER (eight bits) are to be used as described above. The remaining eight bits of the leader are unassigned. The DATA field, of variable width, is unused in irregular messages and is formatted according to the HOST-to-HOST protocol conventions is regular messages.

The IMPs exercise traffic control and restrict the total amount of data in transmission through the network by:

- a. Refusing additional messages over a 'blocked' link.
- b. Limiting the maximum size of a message to 8096 bits.
- c. Limiting to 63 the number of transmit and receive links a HCST can concurrently have active.
- d. Rejecting any message which required longer than 15 seconds for re-assembly from its constituent packets.
- e. Rejecting any message which has not been collected by a destination HOST within 40 seconds of re-assembly.

The constraint of a regular message to not more tha. 8096 bits in length is imposed by the IMP sub-network. The second-level or HOST-to-HOST protocol passes this constraint to the user process.

TN73-50-I(5)				
	Regular Messa 32 Bi	ige LEADER		
FIILER (M1) E bits	CONNECTION BYTE SIZE (S) 8 bits	NUMBER OF BYTES (C) 15 bits		
FILLER (M2) 8 bits				
TEXT				
	2 Bytes Each of S	5 Bits per Lyte		
1		FILLER (M3)		
   		variable		

Fig. 5: HOST-to-dOST Message Format

and the state

a vital at the second of the

in the

14.3

#### 2.5 HOST-to-HOST Protocol and the NCP

An early design for the HOST-to-HOST protocol was reported on in CARR, CROCKER & CERP (Ref. 4), but the version implemented was formulated by the Network Working Group and is specified in CROCKER (Ref. 7).

This protocol is based on the concept of a HOST-resident Network Control Program (NCF), which can communicate with other HOST-resident Netsork Control Programs, by means of 'control messages', issued over 'control links', for the purpose of effecting 'connections' between 'sockets' assigned to HCST-resident processes.

All HOST-to-HOST communications are regular messages, control messages being sent over link zero and interprocess transmissions over non-zero links. The transmitting NCP is responsible for segmenting interprocess transmissions into regular messages, not more than 8096 bits in length. The receiving NCP is therefore responsible for concatenating successive regular messages into a single, or differently divided, transmission for delivery to the receiving process. Because of the differing word lengths of the communitating HOSTs, the DATA field of regular messages (Fig. 4) will not, necessarily, start or end on word bo idaries at the receiving HOST. In order to concatenate successive regular messages, the receiving NCP has to be prepared to employ extensive bit shifting; an expensive operation in terms of processor time. To alleviate this situation, the notions of a 'message header! and a 'connection byte size! have been introduced.

The message header embraces the regular message leader and extends the required information at the beginning of each message to a total of 72 bits. The figure of 72 was chosen as the lowest common multiple of most HOST byte and word lengths; these being 8, 18, 24 and 36 bits. This header reduces the bit shifting problems, at the beginning of a message, for the majority of HOSTs. The format of the header is given in Fig. 5.

It consists of the 32-bit leader, specified ty the HOST-to-IMP protocol; an 9-bit filler field (M1); an 8-bit connection byte size field (S); a 16-bit number of bytes field (C) and a further 8-bit filler field (M2). The text field of a message consists of C bytes each of S bits per byte. It was decided that, for each connection, a connection byte size should be agreed upon, by the communicating procedures, and that this would typically be the lowest common multiple of the respective RCST word lengths. The text of messages would then be an integral number of these bytes. This procedure reduces the bit shifting problems at the end of a message. In addition, a variable length filler field (M3), is specified at the end of a message which can be used by the transmitting HOST to fill the message

A P. Mide Bulling on White States States and the second

The second s

1 8 1 NOP (0)	   			T	N73-50-1(	(6)
8   RIS (1)	     recei	32 ive socket	   send :	32 socket	   8   1ir	B n <b>k</b>
8 SIR (2)	l I send	32 socket	32 receive socket		8   Size	
8   CLS (3)	l I my so	32 ocket	32 your socket		     	
8   ALL (4) 	8   link	nessage :	6 space	32     bit space		
8   GVB (5)	8   link 	8 message 1	fraction	8 bit fraction		
8   RET (6)	8   link 	nessage s	16 3 message space bit sp		2 ace	
8   INR (7) 	   8   link 					
6   INS (8) 	8   1ink 	   				
8 ECO (9)	8 data					
8   ERP (10)	g data					_
6   ERR (11)	8 code		80 data	3		
8   RST (12)	l 1 1					
8   RRP (13)	t ( (	Fig. 6: HO	)ST-to-HO	ST Control	L Comman	nds

out to a word boundary. All the filler fields have a zero value.

A connection byte size is specified when the connection is first established, on behalf of processes, by the co-operating NCFs. Connections between NCPs, on links number zero, are not, however, explicitly established but are considered to exist when a HOST becomes active. The connection byte size for these connections is implicitly defined to be 8 bits, and the text field always consists of an integral number of control commands.

The format for each control comwand is given in Fig. 6; the first field of 8 bits, known as the 'op-code', defines the function of the command and successive fields specify parameters to it. The figure shows the three character mnemonic name of each op-code; its numeric value in parentheses; the size of the parameter fields and the type of parameters required. Op-codes are defined for such functions as the establishment of a connection, RTS (receiver to sender) and STR (sender to receiver); the termination of a connection, CLS (close); the control of buffer space, ALL (allocate), GVE (give back) and RET (return); the notification of interrupts, INR (interrupt receive link) and INS (interrupt send link); the reporting of errors, ERP (error) and the determination of status, ECO (echo), EFP (echo reply), FST (reset) and RRP (reset reply).

When two processes desire to establish a connection, the prospective sender requests its NCP to issue an STR command and the prospective receiver requests its NCP to issue an RTS command. The STR command specifies the lccal send socket, the foreign receive socket and a connection byte size. The NTS command specifies the the foreign send socket and a local receive socket, previously unassigned link number for use by the connection. Once these commands, known as 'requests for connections' (RFCs), have been exchanged, and provided that the specified sockets match and were not already in use, the NCP of the prospective receiver issues an ALL command specifying the amount of buffer space it is prepared to assign to the connection. On receipt of this command, the NCF of the prospective sender is permitted to transmit data, on behalf of its process, up to the amcunt specified in the command. It is expected that as the NCP of the prospective receiver passes the data to its process and releases its internal buffer space, it will issue further ALL commands to the sender, to enable transmissions to continue. The NCP of the prospective receiver can, at any time, request a return of all or part of the buffer allocation by means of a GVB command and this is to be acknowledged by a RET command. All NCPs are to maintain message and bit counters for each of their send connections which are to reflect the current buffer allocation at the receiving site. They are specifically prohibited from issuing any message which would result in either counter being decremented below zero. This flow control mechanism does not apply, however, to control messages issued on links number zero.

Connections can terminate by either NCP issuing a CLS command which the other NCP is to echo. Care is required to ensure that a CLS command, issued by a sending NCP on link zero, does not arrive at the receiving NCP before the last data message on the connection to which it refers.

An NCP is permitted to issue an ERR command whenever it detects a HOST-to-HOST protocol error by another NCF. The code field of this command specifies the type of the error and the data field gives additional information about it. A number of error types are network defined but it is expected that implementers of NCPs would publish details of their own additional error types. (See Section 4).

Further details of the HCST-to-HOST protocol are given in Appendix C.

#### 2.6 Frocess-to-Process Protocols

Two process-to-process protocols were defined at the time the system to be described was implemented; the Initial Connection Protocol (ICP) and the Telecommunication Network protocol (TELNET). These protocols are used both by the LOGGER SERVER and the user TEINET processes discussed in Section 4.

A family of ICPs is formally specified, on behalf of the Network Working Group, by DOSTEL (Ref. 15). The protocols require a 'listen' function to be defined which permits a process to specify to its local NCP the identity of a socket for which it is prepared to accept a connection from any other process. On receipt of an appropriate request for connection to such a socket, the local NCP is to issue the matching RFC, thus completing the connection, and to advise its process accordingly. This function is effected in the Lincoln Laboratory system by the process-to-NCP command, enable connect (ENC) and enable listen (ENL) described in Section 4.

Using this protocol, a server process listens on one of its send sockets, L, which is well known throughout the network. A user process, wishing to communicate with the server process, requests its NCP to establish a connection to L from one of its receive sockets, U. On establishment of the connection, which is to have a byte size of 32 hits, the server process has two options. It could either request its NCP to close the connection, indicating that it is not prepared to communicate further with the user process, or, it could send a single 32-bit even number S, identifying a pair of sockets, S and S+1, which are to be used for all subsequent transmissions. If the server process adopts the latter option, both processes are to request their respective NCPs to close the connection between L and U and to establish new connections between S and U+3 and S+1 and U+2. The protocol thus provides a means of establishing duplex communications between process. The well known socket, L, is called the Initial Connection Socket and is defined to be X'00000001' for the LOGGER SERVER, where X signifies a hexadecimal number. U+2 and U+3 are specified for use, rather than U and U+1, in order to avoid a race condition.

The TELNET protocol, described by O'SULLIVAN (Ref. 14) on behalf of the TEINET committee, was designed to make a process at a using site appear to a process at a serving site as logically equivalent to a terminal directly connected to the serving site. The protocol is effected through the medium of a Network Virtual Terminal (NVT) which normally uses a standard, intermediate terminal code. Both the serving and using processes map their local codes to this intermediate code and must conform to the conventions of the NVT. This approach eliminates the need for a HOST to keep information about the characteristics of terminals and terminal handling conventions at other host sites.

The intermediate data code of the NVT is defined to be the 7-bit ASCII code transmitted in 8-bit bytes with the high order hit set to zero. Connections between processes are established using the Initial Connection Pretocol and have a connection byte size of 9 bits. A number of 8-bit TEINET control codes are also defined, each with the high order tit set to one. These are used for such purposes as the notification of a break or reverse break signal; a request to echo or cease echoing data characters, or a warning to suppress local printing of input characters. In addition, a data type control code can be sent as the first byte of data over a connection to signify whether the standard NVT code is to be used or some defined alternative such as EECDIC. A special synchronising code is also defined, which can be inserted in the data stream and which is always associated with a network interrupt (INS) command. This enables stacked input to a serving process, up to the synchronising code, to be discarded or temporarily ignored by those systems which recognised an 'attention' key as having a specified interrupt function.

#### 3. The CP/CMS Time-Sharing Systam

The Lincoln Latoratory CP/CMS time-sharing system is a modified version of IBM's CP-67/CMS time-sharing system (Ref. 6) and consists of two independent components: the Control Program (CP) and the Conversational Monitor System (CMS). The Control Program creates a set of 'virtual machines' which are time shared with one another, and the Conversational Monitor System provides a conversational operating system which can be used to control a virtual machine.

A virtual machine is a functional simulation of the real IBM 360/67 and its associated input-output devices. The virtual memory and virtual devices of a virtual machine are specified by a predefined configuration which can differ for each virtual machine. The Control Program allocates the CFU resources of the real Lachine, for a short period of time, to virtual machines in turn, in accordance with a scheduling algorithm.

All input/output commands from a virtual machine are trapped by CP and translated to real I/O operations. Unit record input-output is normally spooled onto disk. Virtual memory is divided into 4096-byte blocks, held on paging drums and paged in and out of real memory as required. Special hardware is provided on the IBM 360/67 for paging purposes which performs 'dynamic address translation'.

CP maintains a record of all permitted virtual machines and identifies each one by means of a USERID. Associated with this USERID is a password, accounting information and a table describing the configuration of the related virtual machine. Provided that the number of active virtual machines is not the maximum that the system can support at any one time, a user could request CP to establish his virtual machine by logging in on an unassigned terminal. "Logging in" involves the typing of a USERID and password which are used to validate an authorized user. The terminal enables a user to control his virtual machine, by means of CP console functions, in a manner similar to that of an operator at the console of the real machine. The virtual machine appears to a user and his programs to he indistinguishable from the real machine.

Having established the virtual machine, the user can employ CP console functions for such purposes as the Initial Program Loading (IPI) of an operating system; the display or modification of any portion of his virtual memory; The stopping or restarting of his virtual machine; the sending of messages to other terminal users, the setting of special virtual machine facilities; and the querying of CP for systems information.

CMS is one of the operating systems which can be loaded into a virtual machine by means of the TPL function. It is a single user conversational system which provides a comprehensive command language and programming facility. Commands exist for file creation and manipulation, program compilation, execution control, debugging and various utility operations.

The file handling commands enable the user to create, copy, move, combine, edit, print and erase disk files

A CARLES AND A CAR

and other compands provided access to tape units, line printers, card readers and card punches. Compilation commands are available for Assembler, Fortran IV and PL/1, and the execution control commands allow the user to load relocatable object programs and execute them under terminal control. Executing programs can be halted at pre-determined points, examined, modified and restarted using the debug commands. Utility functions provide for tape and disk copying, sorting, dumping and printing.

An EXEC facility which includes substitutable parameters and conditional statements allows a set of commands, stored in a file, to be obeyed by the command interpreter. Control can be transferred from the CMS command interpreter to the CP console function environment by means of ar "attention" key, thus allowing locping programs to be aborted.

Both CP and CMS commands can be issued from the terminal or called from user programs. All the network software was designed to operate in a virtual machine environment, under control of the CMS operating system,

#### 4. The Lincoln Laboratory IBM 360/67 Network Scftware

Support for the ARPA network on the IBM 360/67 computer at Lincoln Laboratory is provided through a Network Centrol Program (NCP) which runs in a virtual machine of the CP-67 operating system. The NCP virtual machine has direct access to the Interface Message Processor (IMP), and hence to the network, and communicates with other virtual machines over a pair of virtual levices, one for sending and one for receiving. The virtual device used for sending can be directed to a particular virtual machine in the same way that user virtual machines direct transmission to the NCP virtual machine. Transmissions to the NCP are identified is to the sending user and thus are appropriately processed by the NCP.

The implementation of the NCP in a virtual machine rather than as part of the core resident system facilitated its development since it could operate independently of other virtual machines and, hence, not impact the reliability of the operating system. The NCP is written as a user program which calls on CMS system facilities for program initialization, including program re-initialization in the event of a catastrophic error, for disk file services, and for virtual machine to virtual machine communications.

By operating the NCP in a virtual machine, core memory is not dedicated to support the network operation; system resources are only required when network activity cccurs. In addition, operating the NCP in a virtual machine is particularly convenient since the protocol by

7 No. 1 House & Rowsell & Statement of State

which NCP's communicate is being changed and will continue to evolve. The added overhead to the operation of the NCP caused by these techniques does not preclude useful network operation and experimentation. A number of improvements can be implemented to reduce this overhead by giving the NCP special priorities if the use of the network is limited by the performance of the NCP.

The NCP keeps a table (the rendezvous table - RV table) whose entries reflect the state of connections between the local NCP and a foreign NCP. Each entry contains the identification of a foreign socket and a local socket, and the userid of the local virtual machine which is to be associated with the local socket. Entries are made in the RV table either on the receipt of a Reguest For Connection (RFC) control message from a foreign NCP or through a reguest from a local virtual machine.

The communication protocol between a user virtual machine and the NCP resembles the NCP to NCP protocol. Data carnot be transmitted until a connection is initiated by both the sender and the receiver and the NCF to NCP requirements for flow control are met. The specifies the number of bits it is prepared to NCF buffer for a receive connection and stores messages received from a foreign NCP until a local user makes a request for a message. If no message has been received when a user makes a receive request, this fact is returned to the user. The NCP will not accept a message from a user for tranmission if the receiving NCP has not adequate buffering. In this way, provided the requirements for flow control between NCP and NCP are passed to the user. The NCP also notifies the user of any asynchronous change in the state of a connection so that he can take appropriate action. For example, the user is notified a) when a connection he has initiated has been completed by the foreign host, b) when a message is received and there had not been any previous messages buffered, and c) when a connection is closed by the foreign host. In addition, the user is notified when a network interrupt has been received for the connection.

A set of low-level user routines (NET) have been provided to communicate with the NCF. These routines follow the protocol defined for network communications through the NCF. For each user request to the NCP, a reply from the NCP is sent to the user virtual machine indicating either successful or unsuccessful completion of the request and the current status of the connection. These roltines handle the stacking of asynchronous interrupts which are unstacked through a call to the NET routines.

One user virtual machine which has been developed is the LOGGER virtual machine. The LOGGER controls a number of virtual terminals which provide access to the CP system

the second second day and second s

in a manner similiar to physical terminals, but instead of a hardware terminal, CP is interfaced with a specially designed virtual terminal device which is driven by software. Thus, information normally typed on a terminal by CP and information normally keyed into a terminal is transferred between the LOGGER and CP through one of these virtual terminal devices. The protocol for controlling a software terminal is described in Appendix E.

When the LOGGER is initialized it is prepared to communicate with the NCP over the socket defined for LOGGER operation under the Initial Connection Protocol (ICP). Communication under the ICP will result in a new pair of sockets being defined through which the LOGGER will provide CP services to the network. Once this pair of sockets is defined, the LOGGER acts as a SERVEP interfacing a virtual terminal with a remote process through the NCP.

During the initial login sequence, the LOGGER analyzes received network messages for 'userid' and 'password' and compares these with entries in a special file. If the userid does not appear in the file, an invalid userid is passed to CP. If the password matches the network password corresponding to the valid network userid specified in the file, the system password is sent to CP. Since network passwords and valid system passwords do not have to be identical, access to the CP system from the network is specially controlled and may be restricted to a subset of authorized system users. Though the operation of the LCGGEE/SERVER as a separate virtual machine introduces another level of dispatching overhead, this does not appear to be significant and does not limit its use.

4.1 System Lesign

Figure 7 outlines the major elements of the network software and the interfaces between them.

The Network Control Program (NCP) resides in a virtual machine (VM1 in the figure) to which is attached the IMP hardware interface and a terminal. The terminal is used for establishing, and monitoring the activity of the NCF. Processes in other virtual machines communicate with the NCF through the CF virtual machine to virtual machine communications (VMCOM) interface. This facility was specially designed for efficient virtual machine to virtual machine communications.

In an earlier version of the network software, the CP facility for transferring card image from the virtual card punch of one virtual machine to the virtual card reader of another virtual machine was used. This process used the CP spocling facility which limited the rate of virtual machine inter-communications to the speed of writing a spooled file to the disk and then reading the



Fig. 7: System Design

and the second second states of the second second

spooled file. With the VMCOM facility, data is moved from the virtual memory of one virtual machine to a CP free storage area, and then from the CP free storage area to the virtual memory of another virtual machine. A description of the VMCOM facility is contained in Arpendix F.

In Figure 7, three virtual machines are shown connected to the NCP through the VMCOM interface. VM3 represents a virtual machine, controlled by a user at his terminal, executing a process which is in communication with the network. All user processes communicate with the NCP by means of a subroutine and macro package, NET, which effects a locally defined process-to-NCF protocol.

VM2 represents a virtual machine containing the logger-server process. This process uses the NET package to communicate with the NCP and observes the TELNET and ICP process-to-process protocols for all network transmissions. A modification was made to CP, by the system maintenance group, to provide a CP virtual terminal interface for use by the logger-server. This interface enables the logger-server to establish and service virtual machines as though it, itself, was a set of local terminals each controlled by a separate user.

The figure shows two virtual machines, VM4 and VM5, established from the network, by the logger-server, through the virtual terminal interface. Both virtual machines contain user processes which are executing under the remote control of network users or their local processes. In addition, VM4 represents the case where the executing user process is itself accessing the network by means of the NET package.

This design allows for more complex configurations to be established than that shown in the figure. Any number of active virtual machines could use the NET/VMCOM interface and each user process is permitted a miximum of 256 connections to the network. The logger-server currently supports up to three virtual machines, at any one time, as this is considered to represent the maximum, in terms of local resources, that should be allocated to the network.

Brief descriptions of each of the major elements follow, together with some implementation details.

4.2 The NET Package

Figure 8 shows in block schematic form the way in which a user process communicates with the Network Control Program.

A set of high-level NET commands are defined which can be employed by a process to call subroutines which used the NET package. Each high-level command is in the form of a macro, which expands into one or more primitive NET



Fig. 8: NET Package

and the second secon

commands. These commands, which can be used directly by the user process, are interpreted by the NET package and issued as requests to the NCP through the medium of the VMCOM facility. Communications between NET and the NCP consist of an 80-byte control block together with a data block of from 0 to 1024 bytes.

The format of the control block is shown in Figs. 9, 10, and 11 as a set of field names, each followed by the field byte size in parentheses.

The CODE field, one byte. is used to specify the type of control block and the remaining fields are used for parameters set by either NET or the NCP. The CODE is given by a numeric value, in the range 0 to 12, and the significance of each is indicated in Fig. 12.

CODEs are defined which allowed NET to request the NCP to:

- a. Issue, on its behalf, certain network commands such as STR (0), RTS (1), CLS (4), INS and INR (6).
- b. Transmit, on its behalf, interprocess messages over established connections (2).
- c. Deliver to it, interprocess messages from the network which the NCP had buffered until required (3).
- d. Advise on the status of particular sockets (5).
- e. Enable specified sockets to accept foreign STRs (9) or RTSs (10) without having to know in advance the foreign socket numbers. (The 'listen' function referred to in Section 2.6).
- f. Disable sockets which previously had been enabled (11).
- g. Purge all local references to connections not correctly closed by foreign HOSIs (12).

The NCP replies to each NET request with a reply control block (7) which contains appropriate status information in certain of its parameter fields. When the NET request is for the delivery of an interprocess massage (3), the reply control block is followed by a data block, which represented the message. The NCP can also pass asynchronous interrupts to NET, at any time, by means of an interrupt control block (8). A number of interrupt codes are defined to notify the process of such events as:

a. The receipt of a network interrupt (INS or INR) on che of its current connections.

the star of the second second second

TN73-50-I(9)

NET TRANSMISSION CODE (1) USER ID (8) LOCAL ID (3) ) LOCAL TAG (1) ) FOREIGN HOST (1) ) FOREIGN TAG (1) ) FOREIGN TAG (1) ) EIT COUNT (2) INTERRUPT CODE (1) STATUS (4) EIT ALLOCATION LEFT (4) MESSAGE ALLOCATION LEFT (2) CONNECTION BYTE SIZE (1)

LOCAL SOCKET (4)

FOREIGN SOCKET (5)

Fig. 9: NET to NCP Control Message Format

TN73-50-I(10) CHARACTER COLE (4) LCCAL ID (3) LOCAL SOCKET (4) LOCAL TAG (1) FOREIGN HOST (1) FOREIGN ID (3) FOREIGN SOCKET (5) ) FOREIGN TAG (1) \_\_\_\_\_) FLAG (1) EJI COUNT SENT CR RECEIVED (2) EUFFER ADDRESS (4) STATUS (4) BITS LEFT (4) MESSAGES LEFT (2) CONNECTION BYTE SIZE (1) UNUSED BYTE (1) ADDRESS OF INTERRUPT STACK (4)

Fig. 10: NET Command Format

TN73-50-I(11)

INTERRUPT CGEE (1) NEXT ENTRY IN INTERRUPT STACK (3) LOCAL SOCKET (4) FOREIGN SOCKET (5) SPARE (3) STATUS (4)

Fig. 11: Interrupt Stack Format

- b. The acceptance, by a foreign HOST, of a request to establish a send connection to it, (i.e., the receipt of a foreign RTS which matched a previously issued STR).
- c. The loading of NCP buffers for one of its receive connections.
- d. The closing of one of its connections by a foreign HOST.
- e. The matching of 'enabled' sockets with appropriate foreign STRs or RTSs.
- f. The shutdown of the local NCF.
- g. The resetting of any foreign NCP associated with one of its connections.

The interrupt codes are listed in Fig. 13 and any interrupt control block sent from the NCP to NET would contain one of these codes in the second byte of the BIT CCUNT field, which doubles as the INTERRUPT CODE field (Fig. 9).

The USERID field it employed by both the NCP and NET to identify the virtual machine associated with the NET request. NET accesses the CP USERID, on behalf of its process, and completes this field accordingly. Sockets are defined (Section 2.2) to consist of an 8-bit HOST identifier concatenated with a 32-bit socket number. Local sockets are specified by the 32-bit socket number alone, since the local NCP can add its own HOST identifier as necessary. The LOCAL SOCKET field is ID and LOCAL TAG divided into LOCAL fields for convenience of mapping to process identifiers and ports, and the FOREIGN SOCKET field is similarly divided, but includes a FOREIGN HOST field. Both the LOCAL SOCKET and FOFEIGN SOCKET fields are normally completed by NET, except in the case of enable requests where NET completes only the IOCAL SOCKET field and the NCP issues an interrupt control message, when the enabled socket is matched, with the FOREIGN SOCKET field completed. The BIT COUNT field is used by NET when sending an interprocess message to the NCP and by the NCP when delivering an interprocess message to NET.

The STATUS field is always completed by the NCP whenever it issues a reply or interrupt control message. The field is divided into 4 groups, each of 8 bits; the first two groups describe the reply status of the previous NET request, and the second two groups refer to the connection status of the associated entry in the RV table. Each bit, when set, indicates that a particular event has occurred or a particular state exists and the significance of each is given in Fig. 14.

The Internet of the State of Concerns
-	TN7	3-50-1(12)		
	NET CHA C	COMMAND N RACTER CDE	CF CONT MESSAGI CODE	TPOJ. E CESCRIPTION
1	CON	(Connect)	0	Pequest NCF to issue an STR.
1	LIS	(Listen)	1	Request NCP to issue an RTS.
1	SND	(Send)	2	Request NCP to transmit a buffer.
1	RCV	(Receive)	3	Accept a buffer from the NCP.
1	CLS	(Close)	4	Pequest NCP to close a connection,
ł	STA	(Status)	5	Request status from the NCP.
1	INT	(Interrupt)	6	Pequest NCP to issue an INS or INR.
1	NOP	(No-operation	) 7	Request status from NET. Reply code from NCP.
1	-		8	Interrupt code from NCP.
1	ENC	(Enable conne	ct) 9	Request NCP to accept a foreign STR.
1	ENL	(Enable liste	n) 10	Request NCP to accept a foreign RTS.
1	DIS	(Disable)	11	Request NCP to disable an enabled socket.
1	PUR	(Purge)	12	Request NCP to purge an entry.

Fig. 12: NET Commands



Fig. 12: NCF to NET Interrupt Codes

28

Those bits marked with an asterisk (\*) in the RV status are reset by the NCP after it has replied to an NCP status request.

The BITS LEFT and MESSAGES LEFT fields are completed by the NCP, whenever it replies to a NET request associated with a send connection, to notify NET of the current foreign allocation for that connection. The CONNECTION BYTE SIZE field is completed b, NET for all connection requests, though only issued to the network, by the NCP, for send connections. A zero value on receive connection requests implies acceptance by NET of any connection byte size.

One advantage of this NCP mechanism is that a virtual machine, containing a process in communication with the network, can be 'logged off' without breaking network connections. All incoming messages for the process are buffered by the NCP and all information relating to interrupt control messages are recorded in the status flags. On resuming activity, the process can determine its current connections by requesting status from the NCF and can continue operations by taking the appropriate action on its connections.

Figure 10 shows the format of the parameter list to be used by a process when issuing a primitive NET command to the NET package. WET is accessed by a subroutine call which is associated with a pointer to the parameter list. The CHARAC TR CODE field of the parameter list contains the mnemonic names of the command required and this is interpreted by NET which issues the appropriate control message to the NCP. The mnemonic names of all permitted commands are shown in Fig. 12 associated with the control message codes which are generated by NET. All commands cause the generation of a control message, with the single exception of NOP which is used to obtain the NET interrupt status. The LOCAL SOCKET, FOREIGN SOCKET, BIT COUNT, STATUS, BITS LEFT, MESSAGES LEFT and CONNECTION EYTE SIZE fields are all used for the same purposes as, and mapped by the NET package to, the fields of the same name in the control message. The PLAG field is reserved for use by the high-level NET commands and the BUFFER ADDRESS field contains a pointer either to a buffer holding the message for transmission in a SND command or to a huffer prepared for reception of a message in a RCV command.

NET maintains in free store an interrupt stack on behalf of its process. The ADDRESS OF INTERRUPT STACK field holds a pointer to the head of this stack, each entry of which is to be in the format shown in Fig. 11. If there are no entries, the pointer has a value of zero. NET will read interrupt control messages, from the NCP, either on receipt of a NOF command from the user process or when awaiting a reply control message from the NCP. Information contained in each interrupt control message is transferred to the appropriate fields of a separate

A THOM DATA ALL LAND OF THE AND

TN73-50-I(14) GROUP 1 - NET STATUS OURIMP EQU X 80 1 OUR IMF DEAD ALLSTR EQU X 140 1 ALL OUR FREE STORE IS IN USE LNKTBL EOU X 1 20 1 OUR IMFS LINK TABLE IS FULL FORIME EQU X ' 10 ' FOREIGN IMP/HOST IS DEAD FORLINK EQU X \* 08 \* ALL LINKS TO HIM ARE IN USE X 04 1 TAGMIS EQU TAG IS OF WRONG POLARITY FOR RFC RFCENB X '02' ATTEMPT TO ISSUE AN RPC ON ENABLED SOCKET EOU ATTEMPT TO REPEAT AN FFC RFCRFC EQU X'01' GROUP 2 - NET STATUS ¥ SOCKMIS EQU X \* 80 \* SOCKET MISMATCH LINKMIS X 40 1 EQU IINK IS OF WRONG POLARITY X 20 1 TXLEJ EQU TRANSMISSION REJECTED - SFU RV STATUS CMDREJ EOU X 10 1 COMMAND REJECTED - SEE RV STATUS X \* C8 \* DRAINON EQU DRAIN SET ON X 04 1 TTEMPT TO ISSUE CLOSE ON CLOSED SOCKET CLSCLS EQU BADFAF EQU X 02 BAD PARAMETERS X'01' NOTPER EQU NOT PERMITTED - HOST NOT KNOWN ETC \* GROUP 3 - NET AND BV STATUS \* OURCON EQU X • 80 • OUR CONNECT/LISTEN ISSUED X 40 1 HISCON EQU HIS CONNECT/LISTEN ISSUED BUFFLD EQU X \* 20 \* **BUFFERS ARE LOADED** X 10 1 INTERRUPT RECEIVED INTERPT EQU OUR CLCSE ISSUED OURCLS EQU X 108 1 HISCLS EOU X\*04\* HIS CLOSE ISSUED CEASE X • 0 2 • CEASE ON LINK ISSUED EQU RFNM X 01 1 EQU RFNM AWAITED GPOUP 4 - NET AND RV STATUS X\*80 \* ALLOCATE SENT OR RECEIVED ALLOCOK EQU X 40 \* EQU ENABLED FOR CONNECT OR LISTEN ONLY ENABLED ALLCONOK EQU X 20 1 ALLOCATE TOO SMALL FOR CURRENT ORDER \*\*\* X 10 1 MSGFAIL EQU LAST MESSAGE FAILED IN TRANSMISSION \*\*\*\* SNDNOK EOU X\*08 \* SEND CEDER HAS BIT CCUNT >8000 \*\*\*\* BYMIS X • 94 • BYTE SIZE MISMATCH EQU X 02 1 UNUSED EQU INTEPLST EQU X . 01. NETIN INTERRUPT LOST (\*) - THE INTERPT BIT IS CLEARED BY STATUS (\*\*) - THE ENABLED BIT IS CLEARED ON OUR LIS OR CON (\*\*\*) - THE ALLOCNOK BIT IS CLEARED WHEN MORE ALLOCATION IS RECEIVED (\*\*\*\*) - THE MSGFAIL AND SNDNOK BITS ARE CLEARED ON THE NEXT SND

Fig. 14: Format of Status Bits

30

and the second second second second second second

entry in the interrupt stack, the entries being chained by means of the NEXT ENTRY IN INTERRUPT STACK field. The last entry contains a value of zero in this field. It is the responsibility of the user process (or the high-level NET commands) to take action on the stack entries, to return to free storage the block associated with each one and to maintain the pointer in the ADDRESS OF INTERRUPT STACK field.

In addition to the full status reply in the STATUS field of the parameter list, a return code is placed in general purpose register 15 in order to conform with the standard CMS interface. This, together with the standard calling procedure, allows the NET package to be used by a process written in any of the languages supported by the CMS system. See Appendix G.

It is expected that most processes would access NET through the set of macros which effect the high-level NET commands (see Appendix H). For each connection required, a NETSOCK or NETSOCKV macro is to be specified which generates a buffer for the NET command parameter list (Fig. 10) and associates identifiers with certain cf its fields. Before opening a connection a NETID macro is to be issued, which initializes the PLAG field of the parameter list and obtains a value for the LOCAL ID field of the LOCAL SOCKET. The CP UTABLE address for the virtual machine is used for the ID field. A NETOPEN macrc can then be issued to establish the connection whose polarity is determined by the value in the LOCAL TAG field. This can be followed by NEIREAD or NETWRITE macrus, as appropriate to the polarity, in order to transfer interprocess messages. Flags (Fig. 15) are set in the PLAG field of the appropriate parameter list, resociated with each connection, by a routine which analyzes and collapses the interrupt stack each time a high-level command is issued. A NETNOP macro is also provided for this purpose. Network interrupts can be issued by means of the NETINT macro, status obtained by means of NETSTAT and the connection closed by means of NEICLOSE.

Most macros expand into code which calls subroutines of the same name as the macro. Each subroutine is associated with an appropriate parameter list and generates one or more calls to the NET package.

The arrival of interrupt control messages, from the NCP, is detected by an interrupt from the VMCOM device on the virtual machine. A process, therefore, which wishes to suspend itself awaiting some network activity can issue a CMS WAIT SVC on the VMCOM device until an interrupt control message has arrived. A NETWAIT macro is provided for this purpose. When the interrupt arrives and control is returned from NETWAIT, the process can then issue a NETNOP macro and act on its flags accordingly. An example of this procedure, together with the use of most of the macros, is given in the TELNET process.

Control Control of the Advantage of the

Contraction of the second second

*	EQUS	FCR FIAGS	IN	NETPLIST	TN73-50-I(15)	]
OPEN SEND NETIN LISTEN BUFFERS CLOSED ENABLE DEAC	EQU EQU EQU EQU EQU EQU EQU	X * 80 * X * 40 * X * 2C * X * 10 * X * 08 * X * 04 * X * 02 * X * 01 *		OPEN/NOTOPEN SEND/RECEIVE NETWORK INTERN LISTEN ISSUED BUFFERS LOADEN LINK CLOSED ENABLE MATCHEN NCP DIED OR RE	RUPT D D ESET ISSUED	

- OPEN Set when CPEN or OPENE macro successfully issued Cleared on successful close.
- SEND Set when CPEN or OPENE macro issued on send socket Set when enable on send socket issued Cleared on successful close.
- NETIN Set on receipt on network interrupt Cleared on a netstat before STA issued.
- LISTEN Set on receipt on a LISTEN interrupt Cleared on a netstat before STA issued.
- BUFFERS Set on receipt of a Buffers loaded/ready interrupt Cleared on a netread before RCV issued Cleared on a netwrite before SND issued Cleared on a netstat before STA issued.
- CLOSED Set on receipt of a socket closed interrupt Cleared on successful net close.
- ENABLED Set on receipt on an enabled completed interrupt Cleared on a netopen.
- DEAD Set on receipt of an NCP dead interrupt Not reset.

Fig. 15: NET Flag Bits

32

## 4.3 The Logger-Server

The NET system is designed to operate in the CMS environment which supports a single user process only. When more than one user process is required to be concurrently active in a virtual machine, as in the case of the logger-server, a polling or scheduling routine is needed and the high-level NET commands cannot be used. One reason for this is because the high-level commands deny access to the interrupt stack maintained by NET and for multi-process working it is necessary for the polling routine to associate NET interrupts with the separate processes. Rather than complicate the high-level commands for the logger-server system.

The logger-server is designed to be interrupt driven by a simple polling routine which activates a number of processes in turn is suspends itself whenever all outstanding interrupts have been serviced. Four devices can raise interrupts on the logger-server virtual machine and thus cause the polling sequence to be resumed. These are the VMCCM device, whose interrupt signifies the receipt of an interrupt control message from the NCP, and three CF virtual terminals whose interrupts indicated the presence of messages for their associated remote users. Associated with each process is a table which contains sufficient details on the state of the process to enable it to be restarted from its current suspension point. The table includes copies of the program status word (FSW), all general purpose registers, buffers, flags and pointers (see Fig. 16).

The poller accesses two flags in each table: the WAIT flag and the START flag. If the WAIT flag of a process is set, that process will not be re-activated during the current poll, but if it is clear, the poller will set it and 'est the state of the START flag. If the START flag is clear, the poller will restore all registers from the table and re-activate the process at the point defined by its PSW, but if it is set, the poller will clear it and activate the process at a start address specified by a start entry in the table.

Four processes are supported by the poller; an enabling process and three CP virtual terminal control processes. The control processes are all identical in function and are implemented by a single segment of re-entrant code. An interrupt handler is implemented which recognizes all interrupts on the virtual machine and associates each of the four devices with a particular process. The VMCOM device is associated with the enabling routine and each of the CP virtual terminals is associated with one of the control processes. On receipt of an interrupt, the interrupt handler examines the START flag of the associated process and if this is clear it clears the WAIT flag. If, however, the START flag is set, it ignores the interrupt. Since a device interrupt will

TN73-50-I(16)

P N.T.

11 m 4			
WTI	EQU	X • 80 •	WAIT FLAG FOR POLLER
DII	EQU	X * 40 *	DEVICE END
UEi	EQU	X * 20 *	UNIT EXCEPTION
UC1	EQU	X * 10 *	UNIT CHECK
AT 1	EQU	X*08*	ATTENTION
UN 1	EQU	X*04*	INEXPECTED INTERRUPTS
ST1	EQU	X *02 *	START FLAG - USED FOR FNARTING
INT	EQU	X .01.	INTERRIPT STACK FLAG
*			
HI2	EQU	X * 80 *	HIO OUTSTANDING
MO 2	EQU	X 40 4	MORE TO COME FROM NOT BUEREDS
AW2	EQU	X * 20 *	AWAITING DESPATCH TO CD
ADDCR	EQU	X 10 1	CASKLAGE RETURN TO BE ADDED
CR	EQU	x . 05 .	CARRIAGE REFURN DETECTED AS LAST IN DUD
BUFPURGE	EQU	X*04 *	BUFFER PURGE IN PROCEESS
HINT	EÇJ	X . 02 .	HIO INTERRIPT AWATTED
REPLCG	EQU	X .01.	REPEAT LIGTN FIAG
*			
INTS	EQU	X * 80 *	INTERRUPT ON SEND ITNK
LISS	EQU	X + 40 +	LISTEN ON SEND LINK
BUFS	EQU	X * 20 *	BUFFFRS LOADED ON SENE TIME
CLSS	EQU	X • 10 •	CLOSE CN SEND LINK
ENCS	EQU	X . 08 .	ENABLE CONNECT ON SEND LINK
ENLS	EQU	X . 04 .	ENABLE ITSTEN ON SENT TINK
NCPS	EQU	X 02	NCP DIED ON SEND ITNK
CHKPW	EQU	X*J1+	CHECK PASSWORD FLAG
*			
INTE	EQU	X 80 1	INTEFRUPT ON RCV LINK
LISE	EQU	X * 4G *	LISTEN ON BOV LINK
BUFR	EQU	X * 20 *	BUFFERS IGADED ON RCV ITNK
CLSE	EQU	X • 10 •	CLOSE ON RCV LINK
ENCE	EQU	X . 08 .	ENABLE CONNECT ON BOUTTNE
ENLR	EQU	X*C4 *	ENABLE LISTEN ON RCV LINK
NCPR	EQU	X . 02 .	NCP DIED ON RCV LINK
DOWN	EQU	X • 0 1 •	LOGGER GOING DOWN
*			

\*

Fig. 16: LOGGER Routine Flags

34

also re-start the poller, if it is suspended, the process which services the device will eventually be re-activated.

In the case of VMCOM interrupts, which signifies the arrival of an interrupt control message from the NCP, the enabling routine issues a NO2 command to NET followed by a subroutine call to an interrupt stack handler. This handler analyzes the NET interrupt stack and generates from it separate interrupt stacks for each of the processes. When a new entry for a particular process is received, the WAIT flag for that process is cleared and its INTERRUPT STACK flag set.

A process suspends itself, when awaiting network or device activity, by means of a special wait subroutine (WEACK) which saves all process registers and the current PSW in the appropriate table before returning control to the poller.

On first becoming established, the logger-server calls an initializing subroutine which attempts to close all connections between itself and the network, purge all references to any of its sockets from the NCP, and log out all CP virtual terminals. Following this, the process tables are initialized and the WAIT and START flags set for each process except the enabling routine, for which only the STAFT flag is set. Control is then passed to the poller.

The poller can activate only the enabling routine at this time, and this causes an enable connect (ENC) command to be issued to NET for the initial connection socket (ICS) of the logger-server. The routine then suspends itself by means of the WEACK subroutine.

When a network user completes a connection to the ICS of the logger-server, the NCP sends an appropriate interrupt control message. This causes the enabling routine to be re-activated by the mechanisms outlined above. On re-activation, this routine accepts the request for connection from the network user if there is at least one CP virtual terminal unassigned. It will issue the appropriate NET commands, in accordance with the Initial Connection Protocol (ICP) described in Section 2.6, to establish duplex communications between the sockets of the next unassigned CP virtual terminal process and the network user process. On the successful completion of this protocol, the enabling routine clears the WAIT bit associated with the newly assigned control process, issues an enable connect for the initial connection socket, and suspends itself, ready to repeat the cycle of events. If there is no control process unassigned, the enabling routine refuses the request for connection, and then issues the enable connect command fcllowed by the WBACK subroutine call.

A DATE OF THE OWNER

The newly assigned control process will, in its turn, be activated by the poller at its start address and will commence to transfer information between the network user process and the CP virtual terminal interface, in accordance with the TELNET protocol and the CP terminal conventions. Initially, the control process is concerned with the logging on sequence, the verification that the required virtual machine is marked for network use, and the mapping of a network password to the local one. Subsequent activity is concerned with the actual passage of messages between the network user process and its local virtual machine; the translation of characters, as required, between ASCII and EBCDIC; and the interpretation of special control characters such as 'break', 'reverse break' and 'synchronize'.

Whenever network or device responses are required, the process suspends itself, using the WBACK subroutine, until the response has been received. During this period, one of the other processes can be re-activated and thus the poller is able to support the enabling process and the three CF virtual terminal processes concurrently, without any apparent loss in response.

The system is implemented in three program segments; LOGMAIN, LOCDEV, and LOGST. LOGMAIN contains a small initializing routine, the poller, the enabling routine, the interrupt handler, the NET interrupt stack handler and various subroutines for stack analysis and character code conversion. LOGDEV is the re-entrant segment which effects the three CP virtual terminal control processes and LOGST is the main initializing routiine.

The logger-server can be closed down by issuing to its virtual machine an 'external' interrupt. This causes the WAIT flag to be cleared for all processes in use and the DOWN flag to be set. On re-activation, each CP virtual terminal control process will issue a 'LOGGER CLOSING DOWN' message to its network user and then close its network connections. When all processes have terminated, the logger-server halts.

Error recovery is directed at maintaining a service to the network as a whole, for as long as possible. Errors detected in CP virtual terminal control processes are either network command rejections by the NCP, through NET, or virtual terminal transmission rejections by CP. In either case, the recovery of the failing process is not attempted but the virtual machine associated with it is logged off, its network connections are closed, and its process table re-initialized. In this way, the integrity of the other processes are maintained, the network user is informed of the failure by the closing of his connections and the facilities of the failing process are recovered for re-use. Errors detected by the lenabling process are network command rejections only. Certain rejections are recoverable as in the case of an initial SND command failing because of an

insufficient allocation from the network user. In such cases the command is repeated, at a reasonable frequency, until it is either accepted or a repeat count has become exhausted. In the latter situation, the connection is then closed. Other rejections are not recoverable, as in the case of an enable connect command failing. Under these circumstances, the logger-server is shut down after a trace flag has been set, indicating the source of the error, and a core dump output to the line printer. This procedure is also followed if a 'program' interrupt condition, signifying a fault in the prcgram code, is signalled by CP.

4.4 The Network Control Program

As in the case of the logger-server, the network control program (NCP) is designed to be interrupt driven by a simple poller which supports a number of concurrent processes in the NCP virtual machine. An outline, in block schematic form, of the inter-relationships between these processes is shown in Fig. 17.

Each process is implemented by a separate segment of code and is identified on the diagram by its segment name.

Communication with user processes, in other virtual machines, is through the VMCOM interface which is described in Section 4.2. Communication with the IMP is through the IMF hardware interface, which was designed by BRYAN (Ref. 3) and is described in Appendix 12.

NCPINIT performs initialization and error recovery tasks and NCPMAIN contains the poller, the interrupt handler and various subroutines associated with accounting and recording activities. Processes are refactivated according to the state of flags in the PORT TABLE which are updated by other processes and the interrupt handler.

NCFXFER accepts control messages from the VMCOM interface and interprets them in accordance with the NET to NCP control message protocol (Fig. 9). Depending upon the type of request and the current state of the RV TAELE, NCFXFER adds items to the IMP'S QUEUE for transmission over the network, adjusts entries in the RV TAELE and issues a reply to the user process via the VMCOM interface.

NCFIMPO is concerned solely with the output of the I.'- S QUEUE to the network and the associated control of the IMP hardware interface. A facility is included whereby this output can be channeled directly into the input buffers of NCFIMPI, for test purposes or for local use when the IMP is not available.

NCPIMPI processes all input from the IMP and implements the HOST-to-IMP protocol. The HOST-to-HOST protocol is



Fig. 17: The Network Control Program

A STATE

implemented partly by NCPIMFI and partly by NCPXFER. Depending upon the type of message from the network and the current state of the RV TABLE, NCPIMPI loads tuffers for collection by local user processes, sends interrupt cortrol messages over the VMCOM interface, adjusts entries in the RV TABLE and add items to the IMP'S QUEUE in reply to network messaces.

NCPMONIT provides facilities whereby entries in the RV TABLE can be inspected, certain network commands generated, NCP status monitored and the system cleanly closed down.

4.4.1 NCFMAIN

The polling mechanism of NCPMAIN is similar to that of the logger-server. Interrupts received from external devices are flagged in a PORT TABLE which is used to determine the next process for re-activation. Each device is associated with a process through an entry in this table. Entries associate the input channel from the IMP with NCPIMFI; the output channel to the IMP with NCPIMPO; the VMCCM interface with NCPXFFR; and the console terminal with NCPMONIT. No entry is required for the output portion of the VMCOM interface.

Each entry in the PCPT TABLE has the format shown in Fig. 18 which was originally designed for duplex devices.

It was subsequently found more convenient to implement simplex entries but the carlier structure was retained. The format is described as a set of field names, each followed by its field byte size in parentheses. Two sets of flags are defined; POLLING PLAGS which are used by the poller to determine whether the process should be reactivated, and if so, at which address; and DEVICE FLAGS which indicate the current state of the associated device. The significance of both sets of flags is given in Fig. 1°.

Other fields contain the start address of the process, save areas for its PSW and registers, and a buffer address. These fields are duplicated, in each entry, to provide facilities for both input and output channels. The CMS DEVICE ADDRESS and DEVICE ID are also specified in the POPT TABLE entry.

The format of the buffers is given in Fig. 20, and provides fields for a variable length MAIN BUFFER, the current CHANNEL STATUS WORD (CSW) of the device, and workspace for the process.

On receipt of an interrupt from a device, the interrupt handler sets the associated DEVICE FLAGS, copies the CSW to the appropriate field in the buffer, and clears the WAITSET bit in the PCILING FLAGS. The MAIN EUFFER is used by the device for any input or output transfers.

States - Superior - States

TN73-50-I(18)

POLLING FLAGS (1) DEVICE INPUT ROUTINE ADDRESS (3) DEVICE INPUT ROUTINE PSW (8) DEVICE INPUT ROUTINE PEGISTERS (64) INPUT BUFFER ADDRESS (4) DEVICE FLAGS (1) DEVICE OUTPUT ROUTINE ADDRESS (3) DEVICE OUTPUT ROUTINE PSW (8) DEVICE OUTPUT ROUTINE REGISTERS (64) OUTPUT BUFFER ADDRESS (4) SFAFE (2) DEVICE ADDRESS (2) DEVICE ID (4)

Fig. 18: PORT TABLE FORMAT

FCLLING FLAGS

TN73-50-I(19)

BUFIN	EQU	X 80	INPUT BUFFER LOADED
BUFCUT	EQU	X • 40 •	OUTPUT BUFFER LOADED
STATIN	EQU	X'20'	INPUT STATE BUSY
STATOUT	EQU	X'1C'	OUTPUT STATE BUSY
WAITSET	EQU	X • 08 •	WAIT STATE SET
SNDSET	EQU	X '04'	SEND STATE _ NOT RECEIVE

DEVICE FLAGS

INPROG	EQU	X*80*	IN PROGRESS STATE
DEVENL	EQU	X * 40 *	DEVICE END
UEX	EQU	X ' 20 '	UNIT EXCEPTION
UCHK	EQU	X ' 10'	UNIT CHECK
ATTN	EÇU	X'08'	ATTENTION
UNEXBIT	EQU	X • 0 4 •	UNEXPECTED BITS

Fig. 19: Folling and Device Flags

POINTER TO MAIN BUFFEF (4) SIZE OF MAIN EUFFER (4) CHANNEL STATUS WORD (4) WORKSPACE (variable) MAIN EUFFER (variable)

TN73-50-I(20)

the state of the s

Fig. 20: Input and Output Buffer Formats

The poller, if suspended on a CMS WAIT, is restarted on the receipt of an interrupt and re-activates the first process whose WAITSET bit is clear. The re-activation is either at the start address of the process or at its current PSW depending upon the state of its STATIN or STATOUT flags. Additionally, an output process requires its BUFOUT flag to be set before the poller re-activates it.

After servicing the data buffer and any current interrupts associated with a device, the process saves its registers and PSW in the appropriate fields and returns control to the poller, awaiting further interrupts. Process activity typically includes the initiation of further transfers over the device channel and the setting of certain POLLING FIAGS in the PORT TABLE entries of other processes.

NCFMAIN contains a subroutine for writing accounting records to disk, in the format shown in Fig. 21.

The subroutine is called by NCPXFER and NCPIMPI to write a record each time any connection becomes fully established or fully closed. The records lare completed from the appropriate entry in the FV TABLE and include the time at which the first Request for Connection (RFC) is made; the time at which the connection became fully established; the time at which the first close is received; and the time at which the connection became fully closed. Other fields in the record contain the total number of messages and bits issued over the connection; the local USEFID; the identities of both scokets; the foreign HOST and the link number. Although charges are not being raised against network users it is considered desirable that sufficient information be recorded to allow the identification of all users and their utilization of system resources.

Other elements contained in NCFMAIN includes a subroutine for writing logging records to disk and a HOSTS WE KNOW table. The subroutine is used for recording any unusual events or error conditions detected by the NCP. The HOSTS WE KNOW table is used to maintain an entry for the current status and link utilization of every HOST on the network.

### 4.4.2 NCPXFER

NCPXFER consists of a routine to read messages from the VMCOM interface and a set of command routines. Each message is analyzed and interpreted as a command from NET and the appropriate command routine is then entered. These routines typically generate, refer to, and modify entries in the EV TABLE; add network messages to the IMP's QUEUE, and send reply messages back to the requesting local process.

the second s

USERID (8) DATE WRITTEN (8) FIRST RFC TIME (8) CONFLETED CONNECTION TIME (8) FIRST CLOSE TIME (8) COMPLETED CLOSE TIME (8) LCCAL SOCKET (4) FOREIGN SOCKET (4) FOREIGN HOST (1) FOREIGN LINK (1) FV TABLE STATUS (2) NUMPER OF MESSAGES (4) NUMBEF CF EITS (4)

Fig. 21: Accounting Record Format

TN73-50-I(21)

TN73-50-I(22) USERID (8) ) LOCAL ID (3) ICCAL SOCKET (4) LOCAL TAG (1) \_\_\_\_\_ FOREIGN HOST (1) FOREIGN ID (3) FOREIGN SOCKET (5) ١ FCREIGN TAG (1) LINK (1) STATUS (2) CONNECTION EYTE SIZE (1) HEAD OF BUFFER CHAIN (3) TEST BYTT (1) TAIL OF EUFFER CHAIN (3) ADDRESS OF NEXT RV ENTRY (4) MESSAGE ALLOCATION (2) MESSAGE ALLOCATION (2) MESSAGE FOR RETURN (2) DAIE (8) TIME1 (8) TIME2 (8) NUMBER OF MESSAGES TRANSMITTED OVER CONNECTION (4) NUMBER OF BITS TRANSMITTED OVER CONNECTION (4) EITS FOR RETURN (4)

Fig. 22: RV Table Format

The RV TABLE contains an entry for every connection supported by the NCP and the format for each entry is shown in Fig. 22.

Fields are provided for the USERID of the local process; the local and foreign sockets; the link umber assigned to the connection; the connection byte size and the current message and bit allocations. Accounting fields contain details of the date; the time of connecting; the time of closing and the cumulative number of messages and bits transmitted over the connection. The STATUS field holds a set of flags (Fig. 14), describing the current state of the connection and address fields point to the head and tail of a buffer chain associated with it and the next entry in the table. The MESSAGES FOR RETURN and EITS FOR RETURN fields are used to determine when a new allocate command should be issued to a foreign NCP.

established RV TABLE entries are for control connections, on links zero, to all HOSTs in the HOSTS WE KNOW table, during the initialization phase. This enables a common, internal procedure to be adopted for all network transmissions. Any message for cutput to the network, including HCST-to-HOST control messages, is formatted by the originating routine and added to the chain of buffers supported by the appropriate send connection entry in the RV TABLE. A pointer to this RV TABLE entry is then added to the IMF'S QUEUE, unless one already exists. HOST-to-IMP messages are treated as a special case and are only issued during initialization and shutdown by NCPINIT. NCPIMPO uses the pointers in the IMP'S QUEUE to access the RV TABLE entry and the associated buffers for transmission, controlling flow over the connection by means of the RFNM AWAITED flag.

Messages received from the network are analyzed by NCPIMFI, and all IMF-to-HOST and HOST-to-HOST messages are trapped and acted upon directly by this routine. Process-to-process messages are stripped of their message headers and added to the buffer chain of the appropriate receive connection entry in the RV TAPLE.

The CONNECT and LISTEN command routines of NCPXFER, after validating the request, examine the RV TABLE for a complementary entry set by NCPIMPI from the network, and, if none exists, generate a new one. An STR or RTS ccrtrol message will, in either case, be added to the buffer chain of the appropriate link zero send entry, for transmission to the foreign NCP. When a LISTEN request matches a foreign STR entry, an initial ALL cortrol message is also added to the same buffer chain. The SEND command routine examines the RV TABLE for the appropriate send connection entry and, after ensuring that there is sufficient allocation, adds the interprocess message to its buffer chain, decrementing the allocation fields accordingly. The RECEIVE command routine, after having found the appropriate receive

States States and States and

connection entry, collects the next buffer from the buffer chain and passes it to the requesting local process. It increments the MESSAGES FOR RETURN and BITS FOR RETURN fields by the size of the message and if either value exceeds a given thresholds limit, it chains an ALL control message to the associated link zero send entry buffer chain and resets the return fields.

The other command routines follow similar sequences; the CLOSE routine issues a CIS control message and removes the RV TABLE entry if the foreign CLS has already been received; the ENABLE CONNECT and ENABLE LISTEN routines generate new RV TABLE entries without issuing network messages; the DISABLE and PURGE routines remove entries without network messages; the INTERRUPT routine issues an INS or INR control message; and the STATUS routine clears certain bits in the STATUS field after placing its value, together with the system status, in the reply control message buffer.

In all cases the command routine issues a reply control message to the requesting local process and returns control to the calling routine. This routine returns control to the poller, with its WAITSET flag clear if there are further messages to be read, cr with it set if the last message has been analyzed.

### 4.4.3 NCFIMPI

This process accepts and anlyzes input information from either the real IMP hardware interface or from an artificial software 'IMP'. NCPINIT activates, during the initialization phase, either the hardware device control routines in both NCPIMPO and NCPIMPI, or the corresponding artificial IMF routines, in accordance with a load time parameter. The NCPIMP artificial IMP routine passes the output from the NCP directly to the analysis routine of NCPIMPI via an internal buffer. This mode of operation is to test new sections of the NCP withcut disturbing the rest of the network.

The same analysis routine is used in either case, and separates input into three classes of messages; irregular messages of the IMP-to-HOST PROTOCOL; regular messages on link zero of the HCST-to-HCST protocol; and regular messages on non-zero links which are part of interprocess transmissions.

Most irregular messages are logged, by means of a subroutine in NCPMAIN, and appropriate flags are adjusted in RV TABLE ENTRIES, HOSTS WE KNOW status fields, and system status words. A type 5 irregular message (RFNM) causes the RPNM AWAITED flag to be cleared in the associated RV TABLE entry and the NCPIMPO routine to be flagged for re-activation. A type 7 (destination IMF or HOST dead) causes a flag to be set in a particular entry of the HOSTS WE KNOW table, the RFNM AWAITED flag to be cleared in an HV TABLE entry, and a logging record to be written to disk. Other irregular messages cause similar sequences to be fcllowed.

Host-to-Host control messages, on links zero, are interpreted as one or more control commands and cause the appropriate control command routine to be entered. These routines perform similar tasks and are complementary to the command routines of NCPXFER. The STR and RTS control command routines, after validating the format of the command, examine the RV TABLE for a matching entry set by NCPXFER on behalf of a local process, and, if none exists, generate a new one. When an STR matches a local IISTEN REQUEST, an initial ALL control message is added to the buffer chain of the appropriate link zero send entry in the RV TAELE. When an STR matches an ENABLE LISTEN or an RTS matches an ENABLE CONNECT, an interrupt control message is sent, over the VMCCM interface, to the associated process. The CLS routine clears the RV TABLE entry if a CLOSE command has already been issued by the local process, otherwise, it chains a CLS control command to the appropriate link entry in the RV TABLE and issues an interrupt zero control messages to the local process. INS and INR command routines issue interrupt control messages to the lccal process; the GVB routine returns the required fraction of the allocation by chaining a RET control command to a link zero entry in the RV TABLE; and FST removes all RV TABLE entries associated with the issuing NCP and warns all local processes involved via an interrupt control message. Similar sequences are followed by the other command routines.

If errors are detected in a message from a foreign host, an ERR control message is returned to that host. Errors that may cccur are:

- a. Illegal OP Code encountered in operations deblocking.
- Short parameter space. End of message encountered before all expected parameters.
- c. Bad parameter(s), e.g., two receive (send) sockets in a RTS (STR), link number not 1<L<32, bad socket polarity within command.
- d. Request on a closed (null) socket. A request (other than RIS/SIF) was made for a non-existent socket.
- e. Socket (link) not connected. A request which requires a connected socket (link), (i.e., ALL, INR, transmit) was made for an existing but not connected socket (link).
- f. Eyte size not consistent.

- g. RFC for your socket already issued by me (timing error).
- h. RFC on an existing socket pair.
- i. RFC on a closed socket pair.
- j. Allocation too small for this transmission.

Only the first five errors have been given a network defined error code. Refer to Appendix D for a description of the ERR control message.

Regular messages on non-zero links are stripped of their message headers and added to the buffer chain of the associated receive entry in the RV TABLE. If the BUFFLD flig (Fig. 14) is clear, it is set and an interrupt cortrol message sent to the local user process.

After completing all activity associated with the current input buffer, NCFIMPI issues a new input transfer order to the IMP interface and suspends itself, to await further interrupts, by returning control to the pcller.

### 4.4.4 NCPIMPO

NCPIMPO accesses buffers indirectly pointed to by entries in the IMP'S QUEUE and outputs their contents to either the real IMP hardware interface or the artificial IMP internal buffer. Only messages addressed to (and criginated by) the local NCP are passed by the artificial IMP, since when operation is in this mode, all foreign HOSTs are flagged as being dead in the HOSTS WE KNOW table.

For each entry in the IMP'S QUEUE, the PV TABLE entry is accessed and if the connection has not been closed and the RFNM AWAITED flag is not set, the next buffer in the buffer chain is output to the IMP, and the RFNM AWAITED flag set. As buffers are output, the routine frees the associated storage and when all buffers in a chain have been output the corresponding pointer in the IMP'S QUEUE is removed.

When a connection is fully closed, the closing routine (NCPXFER or NCFIMPI) frees any storage associated with the buffers in the buffer chain, removes the entry from the RV TABLE and flags the pointer in the IMP'S QUEUE as 'closed'. NCPIMFO, on detecting such a pointer, removes it from the IMP'S QUEUE.

After initiating a transfer, NCPIMPO suspends itself by returning control to the poller with its WAITSET and BUFOUT flags set, awaiting reply interrupts from the IMP. On re-activation, after completion of the transfer, NCPIMPO attempts to output further buffers, until either all have been transferred or only buffers associated with RFNM AWAITED flags are left. Under these circumstances, it suspends itself, by returning control to the poller with its WAITSFM and BUFOUT flags clear. It can then only be re-activated by NCPIMPI or NCPXFER setting its BUFOUT flag. This occurs when NCPIMPI clears a RFNM AWAITED flag in the RV TABLE as a result of a network message, or when NCFIMPI or NCPXFER adds a new entry to the IMP'S QUEUE

## 4.4.5 NCEMONIT

The NCP virtual machine can be run disconnected from its control controle which is the normal mode of operation employed. For monitor purposes, however, the virtual machine is re-connected, usually to a visual display unit, and the NCPMONIT process activated by means of an 'external interrupt. This interrupt' is detected by a routine in NCPINIT which cause the typing of a monitor message and the clearing of the WAITSET flag in the NCPMONIT entry of the PORT TABLE.

Once activated, NCFMONIT attempts to read and obey Monitor commands typed on the console keyboard until a BESIN command is entered. This command de-activates the NCPMONIT process until the next 'external' interrupt is received.

In addition to BEGIN, commands are available to perform the following functions:

- a. Disconnect the console from the NCP virtual machine and obey a BEGIN command (DISC).
- b. Set the DRAIN flag in the system status word to deny any further requests for connection (DRAIN). This is normally issued shortly before shutdown.
- c. Issue an echo network control command to a named foreigu HOST, await the echo reply, ch∈ k its validity and type an appropriate message on the console (ECHO). This is useful for determining whether the NCP of a refate HOST is active or not.
- d. Output to the console, the current status of all HOSIS in the HOSIS WE KNOW table (FHOSI).
- f. Issue a NOP network control command to all HOSTS in the HOSTS WE KNOW table and obey a DISC MONITOR CCMMAND (NOP). This is useful in that it obtains, from the network replies, the current status of all HOSTS.
- g. Issue an FST network control commaid to all HOSTs in the HOSTS WF 'NOW TABLE AND OBEY A DISC monitor command (RESET) This is issued whenever the NCP is re-established fter a shutdown.
- h. List, either on the console or to a named disk

the second state of the second

file, all the entries in the RV TABLE (RV).

- i. List, on the console, all the entries in the RV TABLE associated with a named socket (SOCKET).
- i. Output to, either the console cr a named disk file, a record of the current utilization of the NCP (STATCS).
- k. Shut the NCP down cleanly by closing all current connections, notifying all local users and warning the IMP (SHUTDOWN). The final part of this process is performed by NCFINIT.
- List, on the console, all the entries in the RV TABLE associated with a named local virteral machine (USER).

#### 4.4.6 NCPINIT and Error Recovery

NCPINIT consisted of a set of routines which are called by NCPMAIN for initialization, by NCPMONIT for shutdown, and by all processes for error recovery.

The initialization routine reads and analyzes lead time parameters which determine whether the real or artificial IMP routines are to be established in NCPIMPI and NCPIMPO, and whether or not NCPMONIT is to be activated at lead time. Additional parameters can be 'stacked' in the console input buffer by this routine, for later interpretation, by NCPMONIT. This mechanism is normally used to issue a RESET monitor command when the system first becomes active.

After analysis of the load time parameters, the interrupt handler in NCPMAIN is established and all centrol link entries are set in the RV TAPLE. Selection of the appropriate device control routines in NCPIMPI and NCPIMPO is then made and if the artificial IMP has been specified, all foreign HOSTs are flagged as dead in the HOSTS WE KNOW TABLE. If the real IMP is specified, the hardware interface is activated and HOST-to-IMP pretocol NOP messages are sent to the IMP until a satisfactory reply has been received. A HOST-to-HOST PPETOCOL ECO control command is then sent, by NCPINIT, to itself, and when this is received correctly the IMP interface is considered to be established.

Before returning control to NCPMAIN, the program interrupt PSW is set to the address of the error recovery routine in NCPINIT and the external interrupt trap to the address of the external interrupt handler.

When an external interrupt is received, the routine in NCPINIT determines whether or not a console is connected to the NCP virtual machine and, if one is connected, activates NCPMCNIT. If no console is connected, a SHUTDOWN monitor message is stacked in the console input

The second s

A STATE OF THE OWNER OF THE OWNER OF

buffer before the activation of NCPMONIT. The system can, therefore, be closed down simply by sending an external interrupt to the disconnected NCP virtual machine.

The shutdown routine in NCPINIT issues an interrupt control message to all local users, a HOST-to-IMP protocol 'HOST going down' message to the IMP, and closes the hardware interface. It then clears the interrupt handler and returns control to CMS with a return code set. When the shutdown is a direct result of a SHUTDOWN monitor command, all network connections are closed by NCPMOINT in the normal way, and control returned to CMS with a return code of zero.

In the case of an involuntary crash, caused by a program interrupt from CP, the program interrupt PSW is set to the address of a final error recovery routine; a copy of all registers and core storage is dumped to the line printer; the event logged; and a clean close down attempted by transferring control to the shutdown routine with an identifying return code set. If a second program interrupt occurs, control is returned directly to CMS with a different return code set. The other trocesses call NCPINIT whenever they detect an irrecoverable error condition by deliberately causing a program interrupt, after logging the event.

The NCP virtual machine, after having been automatically established by CP, executes a 'profile' of CMS commands by means of the EXEC facility (Section 3). These commands cause the Network Control Program to be loaded from disk and to be entered at its start address. On control being returned to CMS, further commands in the profile cause a warning message to be sent to the CP console operator and either the re-loading of the NCP or the 'logging off' of its virtual machine according to the value of the return code and the number of times re-loading has been attempted.

## 4.5 The Treatment of Eurors

It is considered that the major requirement of error recovery for the network system is to maintain effective connections, on behalf of users, for as long as possible and to notify them whenever a crash is imminent. A secondary requirement is to record those events which lead up to a crash, in order that remedial action can be taken in future versions of the system.

Three classes of error are detected; device errors associated with the IMP and VMCOM interfaces; system errors caused by faults in the program logic or the exhaustion of facilities; and protocol errors in the use of any of the software interfaces. A fourth class of error, outside the control of the virtual machine, is concerned with real machine failures. Detection and recovery procedures for this class are the

## responsibility of CP.

IMP hardware errors, occurring during the initialization phase, result in an identifying message being typed on the NCF virtual machine console, a particular return code set and control returned to CMS. The profile, controlling the NCP virtual machine, attempts to re-initialize the system a number of times and if the failure persists notifies the CP operator so that the IMP can be investigated.

Device errors cccurring after the initialization phase has been completed invariably cause a system crash. The type of error is then recorded in the NCP disk log and the crash routine in NCPINIT is entered. The crash rcutine attempts to issue an interrupt control message to all local users and a 'closing down' message to the IMP after sending a complete core dump to the (virtual) line printer. Any device errors detected by the routine are ignored.

System errors (Fig. 23) in any part of the Network Control Program also results in the event being logged and the crash routine in NCFINIT entered.

System errors are caused by such events as the failure to obtain a block of free store for new table or queue entries; the detection of an inconsistent entry or structure in one of the tables; the receipt of an interrupt from a device not in the POFT TABLE and the detection by CF of an attempt to perform an illegal operation.

Protocol errors in the use of the software interfaces (Fig. 24) never result directly in a system crash.

Most IMP-to-HOST messages are logged and any format errors from the IMP are ignored. Checks are carried out on the validity of all requests from foreign NCPs and any inconsistencies or errors in format are logged and an appropriate error message despatched to the originator. The validity of all requests by NFT is checked and the status field in the reply control message set accordingly (Fig. 14). The NFT routines tests the validity of requests from a process and replies with error codes as previously listed.

#### 4.6 Concluding Remarks

This report has described the approach adopted by one site in implementing its network software. Other sites have adopted different approaches consistent with their own local requiremen's. The NCP developed by WHITE (Ref. 2<sup>°</sup>) for the IBM 360/75 is accessed from a batch-mode job under the NVT operating system. Some sites directly modified their operating systems in order to incorporate the network facilities.

the second state of the second

and the second second

'TN73	3-50-1(23)		
CODE	ADD INFO	PEASON	CRASH
0	Header and	Code 11 (ERR) message detected	
	12 bytes	from forNCP (NCPIMPI)	No
1	Header	Code 1 message detected from IMP (NCPIMPI)	No
2	Header	Code 3 message detected from IMP (NCPIMPI)	No
3	Header	Code 8 message detected from IMP (NCPIMFI)	No
4	Header	Code 9 message detected from IMP (NCPIMPT)	No
5	Header	We get a reset reply	No
6		SIO IMP IN INOPERABLE (NCPIMPI)	Yes
7		SIO IMP IN BUSY (NCPIMPI)	Yes
8		UE ON CC=1 IMP IN (NCPIMPI)	Yes
9		Other error CC=1 IMP IN (NCPIMPI)	Yes
10		Unexpected Bits IMP IN (NCPIMPI)	Yes
11		UE IMP IN (NCPIMPI)	Yes
12	* * * * * * *	Not DE on IMP IN (NCPIMPI)	Yes
13		Insufficient free store (NCPIMPI)	Yes
14	~~~~~	Our IMP d∈ad (NCPIMPI)	soon
15		IMPS link table full (NCPIMPI)	No
16		NOP From NCP	No
17			
18			
19	Header and	We sent error message to NCP	
~ ^	16 bytes	(NCPIMEI)	No
20	Header	We get a reset (NCPIMEI)	NO
21	Ucoden	System error in CLS (NCPIMPI)	Yes
22	Header	NEW NOST SENT TO US FOR IMPRUSE DEAD (NOTIMPI)	NO
23	Header	NOD from THD	NO
25	neauer	STO IMP OUT Increable (NCRIMPO)	NO
26		SIO IMF OUT PREV (NCTIMPO)	Ies Voc
27	*****	HF ON CC=1 IMPOHT (NCDIMPO)	Vac
2.8		Other error $CC=1$ IMPOUT (NCPIMPO)	Yes
29		Unexpected bits IMPOUT (NCPIMPO)	Yes
30		UE IMP OUT IMPOUT (NCPIMPO)	Yes
31		Interrupt Handler error NCPMAIN	Yes
32	XX/XX/XX	Date/time up	No
33	~ = = = = =	Card reader error (NCPXFER)	Yes
34	* * * * * *	Insufficient free store (NCPXFER)	Yes
35		Software Table error (NCPXFER)	Yes

Fig. 23: NCP Error Codes

51



Fig. 24: Major Interfaces in the System

Subsequent to the original development of the network software, a number of changes have been made to the network. New HOSTs have been connected, bringing the total now to thirty-six, and a terminal IMP (TIP) has been introduced (ORNSTEIN, Ref. 13) to allow terminal users without a HOST to access the network. Some changes have been implemented in the HOST-to-IMP protocol (MCKENZIE, Ref. 11) and proposals made for a different HOST-to-HOST protocol (WALDEN, Ref. 19). Further work has been carried out on the process-to-process protocols (CFOCKEE, Ref. 7), on the design of the network (MCQUILLAN, Ref. 12), and on its performance measurement (COLE. Ref. 5). Future development and utilization has been discussed ty ROBERTS (Ref. 16).

Some of these changes necessitated modifications to the NCF. In the case of new HOSTS, this simply requires the addition of new entries to the HOSTS WE KNOW table, but protocol changes require some modifications to the routines themselves. It is believed that the development of the NCP in a virtual machine and its modular structure has assisted the modification process.

#### Appendiz A

## IMP Interface

The IMP is connected to the 360 multiplexer channel via a duplex interface unit. The 360 BEAD command may be issued only to the input segment of the controller (address 60) and the 360 WRITE command may be issued only to the output segment of the controller (address 61).

The controller will not allow its Hest Master Ready relay to close unless the CONTFOL-ON command has been issued. This relay closure is tested by the IMP at all times. The 360 program must issue a CONTROL-ON to initiate IMF operation and must maintain such a command at all times. When terminating IMF operation the CENTROL-OFF command should be issued. This command renders the controller inert to all IMP service requests and no "Attention" interrupts will be sent to the 360.

The CONTROL-CN command the CONTROL-OFF command may be issued only to the read address. The NO-OF TEST I/O, and HALT I/O commands may be issued to eit address. All other commands, or an improper combination, will then result in the command being rejected during initial selection. Unit Check (Bit 6) will appear in the initial status.

After the 36C program has issued a CCNTROL-ON command the controller goes into a ready state. This condition will allow the input segment to receive data from the IMP. Should the IMP send a message to the input segment, the first 16 bits are shifted in. If no READ command had been issued by the 360 program an "Attention" interrupt is sent to the 360. The 360 program is expected to respond by issuing a READ command.

The status response in the CSW during initial selection may take the following format:

- a. All zero for PEAD, WRITE, or TIO, indicating that the command was successfully initiated.
- b. Unit Check, indicating an improper command. The command will have been rejected.
- c. Unit Exception, indicating that the IMP Master Ready relay is not closed.
- d. Device End/Channel End will be given in response to CONTROL-ON, CONTRCL-OFF, and NO-OP. This is also considered as final status for these commands.
- e. Status Modifier/Control Unit End/Busy will result if the controller is "Busy" at the time the command is issued.

Final Status will be given at the completion of all operations. It will consist of Device End/Channel End (DE/CE) for both READ and WRITE. Certain other conditions may be indicated, as follows:

- a. Unit Exception will be given if the last data transfer in, or out, was interrupted by the InP dropping its Master Ready signal, even ucmentarily. Unit Exception indicates that the data received, or sent, may be in error.
- b. A Device End/Channel End will be given in response to Halt I/O if a command was indeed in operation at the time the Halt was issued.
- c. Final status may include an Attention Bit along with DE/CE. This indicates that a subsequent READ must be issued to obtain data already within the input segment.

#### Appendix B

#### HOST-to-IMP Messages

A HOST-to-IMP message consists of not more than 8096 bits of information of which the first 32 bits are known as the "leader".

	TN73-50	)-I(B.1)					
1	4	1 4	8	8	8	variable	
1	FIAGS	MESSAGE TYPE	CESTINATION HOST	LINK NUMBER	NOT ASSIGNED	DATA	
1	<	<b></b>	32-Bit Leader		>	L • • •    	
I							

Fig. B.1: HOST-to-IMP Message Format

The leader is always present and defines the type of message, its destination and the link number, if any, with which it is to be associated. A description of each hOSI-to-IMP message type follows:

a. Regular Message (Message Type )).

Regular messages are forwarded through the network and presented to a destination HOST, specified by the DESTINATION HOST field in the leader, as a regular IMP-to-HOST message. Fegular messages are the medium for all HOST-to-HOST communications, the DATA field being formatted according to the HOST-to-HOST conventions. Having issued a regular message on a particular link, that link is said to be 'blocked' and further messages cannot be sent until a PFNM or other 'unblocking' message has been received.

b. Error in leader (Message Type 1).

If the NCF receives an IMF-to-HOST message with an invalid leader, it can respond with a Type 1 KOST-to-IMP message to advise the IMP of the format error.

c. HOST Going Down (Message Type 2).

The NCP should transmit a Type 2 HOST-to-IMP message about 10 seconds before voluntarily closing down. The IMP would then mark the HOST dead and advise the rest of the network accordingly. The 10-second delay is to allow time for the notification to reach all other IMPs in the network and for messages already in transit to be processed.

d. No Operation (Message Type 4).

The IMP always discards this message which is intended for use during initialization. A number of Type 4 HOST-to-IMP messages are to be issued by the NCP, on re-activiation after a shutdown, to advise the IMP that the HOST is alive. The IMP will then relay this information to other IMPs in the network.

e. Regular Message for Discard (Message Type 5).

This message has all the properties of a Type 0 HOST-to-IMP regular message except that it is discarded by the destination IMP, before it reaches the destination HOST, and a Type 9 IMP-to-HOST message, rather than RFNM, is returned to the originating HOST to unblock the link. The mussage can be used for test purposes.

f. Error not in Leader (Message Type 8).

If the NCP receives an IMP-to-HOST message which has a valid leader but is otherwise in error with regard to the HOST-to-IMP protocol, it can respond with a Type 8 HOST-to-IMP message. The NCP might respond with such a message upon receiving an IMP-to-HOST message of excessive length.

g. All other HOST-to-IMP message types are undefined.

## Appendix C

## IMP-to-HOST Messages

A IMF-to-HOST message consists of not more than 8096 bits of information of which the first 32 bits are known as the 'leader'.

	TN73-5	6C-I(C1)				
1	4	4	8	8	8	variatle
	FLAGS	MESSAGE TYPE	SOURCE HOST	LINK NUMEFF	NO T ASSIGNED	DATA
1		1	<b>I</b>	l		L • • •
1	<	32-1	Bit Lead	1er	>	l i

Fig. C.1: IMP-to-HCST Message Format

The leader is always present and defines the type of message, its source and link number, if any, with which it is associated. A description of each IMF-to-HOST message type follows:

a. Regular Message (Message Type 0).

Type C IMP-to-HOST message, received by a destination HOST, originates as Type ? HOST-to-IMP messages in the source HOST. All fields are preserved except that what originates as the destination HOST identifier in the HOST-to-IMP leader is converted to the source HOST identifier in the IME-to-HOST leader.

b. Error in Leader (Message Type 1).

The IMP issues a Type 1 IMP-to-HOST message on detecting an error in the leader of a previous HOST-to-IMP message. With the exception of the MESSAGE TYPE field, no significance is to be placed on the contents of the Type 1 leader.

c. IMP Going Down (Message Type 2).

A Type 2 IMP-to-HOST message is issued by the IMP about 5 to 10 seconds before it voluntarily closes down. The delay is to allow notification of its impending shut down to propagate through the network and for messages already in transit to be processed. With the exception of the MESSAGE TYPE field, no significance is to be placed on the contents of the Type 2 leader. d. Blocked Link (Message Type 3).

The IMP will issue a Type 3 IMF-to-HOST message whenever the NCP attempts to transmit a regular HOST-to-IMP message on a blocked link. With the exception of the MESSAGE TYPE field the contents of the leader will be identical to that of the discarded HOST-to-IMP message.

e. No Operation (Message Type 4).

The NCP is to discard all Type 4 messages from the IMP. These are intended for use during the initialization phase only.

f. Ready For Next Message (Message Type 5).

This message is used to notify the NCP that the first packet of a HOST-to-IMP regular message has been successfully passed is the destination HOST and that the link associated it is now unblocked. The link is identified by the contents of the SOURCE HOST and LINK NUMBER fields in the leader of the Type 5 message.

g. Link Table Full (Message Type 6).

The IMP is capable of supporting at most 63 active transmit and 63 active receive links at any one time. A link is defined to be active if the NCP has issued a HOST-to-IMP regular message (or a Type 5 regular message for discard) over the link within a period of about 30 seconds. If the NCP attempts to exceed this limit the IMP will respond with a Type 6 message whose leader is identical, with the exception of the MESSAGE TYPE field, to that of the rejected HOST-to-IMP regular message.

h. Destination IMP or HOST Dead (Message Type 7).

This message is issued by the IMP if the NCP attempts to transmit a HOST-to-IMP regular message (or a Type 5 regular message for discard) to a HOST or IMP which is known to be dead. The leader of the Type 7 message is identical, with the exception of the MESSAGE TYPE field, to that of the rejected HOST-to-IMP regular message. i. Error not in Leader (Message Type 8).

If the IMP receives a HOST-to-IMP message which has a valid leader but is otherwise in error with regard to the HOST-to-IMP protocol, it will respond with a Type 8 IMP-to-HOST message. This message has a leader which is identical, with the exception of the MESSAGE TYPE field, to that of the rejected HOST-to-IMF regular message. The IMP, for example, will issue a Type 8 message on receiving a HOST-to-IMP message of excessive length.

j. Incomplete Transmission (Message Type 9)

A Type 9 IMF-to-HCS<sup>-</sup> message is issued to notify the NCP that the last HOST-to-IMP message sent over a particular link failed in transmission for scme reason. The failure can be attributed to the unaccounted loss of some portion of the message; the loss of the associated RFNM; a message of excessive length or an excessive time required for re-assembly at its destination. A Type 9 IMP-to-HOST message is always returned in response a Type 5 HOST-to-IMP message, instead of a FPNM. Its receipt unblocks the link and its leader is identical, with the exception of the MESSAGE TYPE field, to that of the rejected HoST-to-IMP message.

κ. All other IMP-to-HOST messake types are undefined.

## Appendit D

## The HOST-to-FC.F Protocol

All HOST-to-HOST communications are effected by regular HOST-to-IMP and IMP-to-HCST messages. Control messages are associated with link 0 and interprocess messages with non-zero links. The format of all HOST-to-HOST messages is as shown in Pig. 5.

Link Assignments:

link numbers are assigned as follows:

a.	Link O	- Control Link.
b.	Links 2-71	- Interprocess Communications.
c.	Links 1, 72-190	_ Reserved.
đ.	Link 191	- Network Measurement.
e.	Links 192-255	- Private Experimental Use.

## <u>Control Messages:</u>

Messages sent over the control link have the same format as other HOST-to-HOST messages. The connection byte size is to be 8 bits and the number of bytes is not to exceed 120. Control messages are always to consist of an integral number of control commands.

# Control Commands:

Each control command begins with an 8-bit 'op-code', which defines the function of the command, followed by a set, in some cases null, of parameters which gualifies the function. A description of each control command follows:

a. NOP - No Operation.

Article

90И |

a manufacture and the second state

The NOP Command can be issued at any time and is to be discarded by the receiving NCP. It is useful for formatting control messages and for determining, from the IMP-to-HOST reply, whether a foreign NCP is active or not.

And Land Derry William Strategy And

b. RTS - Request for Connection, Receiver to Sender.

1	1		1			
1	RTS	receive	socket	send	socket	link
I				L		L !

The RTS command is used to establish a connection and is sent from the HOST containing the receive socket to the HCST containing the send socket. The link number to be used by the interprocess connection is to be in the range 2-71 and is assigned by the issuing NCP. There is no prescribed lifetime for an RTS command which can be gueued by the receiving NCP for an aribitrarily long period of time. The command could be cancelled by either NCP issuing a CLS (section C.3.d) and the conrection is established when either the ETS matches a previously issued STR (section C.3.c) or when an STF is subsequently issued which matches the FTS.

c. SIR - Request for Connection, Sender to Receiver.

		-						-		
L		1			1			ł		1
I	STR	1	send	sccket	1	receive	socket	I	size	1
L		1			1			1		Ł

The SIR command is used to establish a connection and is sent from the HOST containing the send socket to the HOST containing the receive socket. The connection byte size is to be in the range 1-255 and is assigned by the issuing NCP. There is no prescribed lifetime for an STR command which can be gueued by the receiving NCP for an arbitrarily long period of time. The command can be cancelled by either NCP issuing a CIS (section C.3.d) and the connection is established when either the STR matches a previously issued PTS (section C.3.b) or when an PTS is subsequently issued which matches the STR.

d. CLS - Close.

| CLS | my sccket | your sccket |

The CLS command is used to terminate a connection or to cancel a request for connection. The 'my socket' field contains the socket local to the issuing NCP and the 'your socket' field contains the socket local to the receiving NCP. There is no prescribed lifetime for a CLS command but it is required that an NCP reply with a matching CLS 'as guickly as possible' in order that the sockets can be released for further assignments

A CONTRACTOR OF THE REAL PROPERTY OF

e. ALL - Allocate.

										_
1		1		1			1			1
I	ALL	1	link	1	message	space	1	bit	space	1
1		1_		1			1			

The ALL command is sent from a receiving NCP to a sending NCP to increase the allocation counters of the sending NCP. Allocate commands can be issued at any time while a connection is fully established. Interprocess messages can only be transmitted over a connection within the specified allocation.

f. GVB - Give Back.

١		ł		1						1
I	GVB	Ì	link	Ì	message	fraction	İ	bit	fracti	oni
I		1.		_1.			1.			!

This command is sent from a receiving NCP to a sending NCP to r quest that the sending NCP return all or part of its message and bit allocation for the connection. The fraction fields specify what portion, in 128ths, of each allocation is to be returned. The command can only be issued while a connection is fully established and the recipient NCP is to reply with a RET command (section C.3.g).

7. RET - Return Allocation.

		-		_						
I				1						
İ	RET	Í.	link	1	message	space	1	bit	space	1
i		1_		1.			1.			_ I

This command is sent from a sending NCP to a receiving NCF either voluntarily or in response to a GVE command (section C.3.f). The command can only be issued while a connection is fully established.

h. INR . Interrupt by Receiver.

		-			
1		1			I
ł	INF	1	lir	i k	۱
1		1_			I

A REAL PROPERTY OF THE PARTY OF THE PARTY OF
i. INS - Interrupt by Sender.

INS | link

The INS command can be sent from a sending NCP to a receiving NCP when the sending process wishes to interrupt the receiving process. The command can only be issued while a connection is fully established.

j. ECO - Echo Request.

| ECO | data |

The ECO command is used only for test purposes. An NCP could issue an ECC command at any time to another NCP and would expect to receive an EBP command (section C.3.k) in reply. The data field can be any bit pattern convenient to the sending NCP and the pattern is to be repeated in the data field of the ERP command.

k. EPP - Echo Reply.

ERP | data |

The ERP command is to be issued by an NCP whenever it receives an echo request (section C.3.j). The data field of the ERP command is to be identical to the data field of the incoming ECP command.

1. EPR - Error Detected.

EPR data

An NCP is permitted to issue an EPP command whenever it detectes a HOST-to-HOST protocol error by another NCP. The code field specifies the type of error and the data field provides additional information about the error. The following codes have been defined:

or designed and which descent on the

The other in the second second second as

A REAL PROPERTY AND A REAL

- Code 0 Undefined. An error for which nc code exists has been detected. The data field is to be formatted according to the requirements of the originating NCP.
- Code 1 Illegal Op-code. An illegal op-code has been detected in a control message. The data field is to contain a copy of the command in error.
- Code 2 Short Farameter Space. The end of a control message has been encountered before all the required parameters have been found. The data field is to contain a copy of the command in error.
- Code 3 Bad Parameters. Incorrect parameters have been found in a control command. The data field is to contain a copy of the command in error.
- Code 4 Request on a Non-Existent Socket. A request, other than an ETS or STR, has been received associated with a link or socket for which no request for connection has been issued in either direction. The data field is to contain a cory of the command in error.
- Code 5 Socket or Link Not Connected. A request, other than an RTS or STR, has been received for a connection not fully established or for a link not yet assigned. The data field is to contain a copy of the cormand in error.

m. RST - Reset.

RST

The RST command is used by one NCP to inform another that all information concerning existing connections between the two should be purged from the tables of the NCP receiving the RST. The sending NCP can expect to receive an RRP command (section C.3.n) in reply.

n. RRP - Reset Feply.

This command is to be issued by an NCP in reply to an RST command (section C.3.a).

o. All other op-codes are undefined.

#### Appendix E

## CP Support for a Virtual Terminal Device

Mcdifications to CP have been made to support a 'software terminal', i.e., one driven by a virtual machine rather than by a real device such as a 2741. CP performs, via software, functions usually done by SIO's to a real terminal device. The purpose is to provide terminal-like access to CP for users wishing to log on to the system through the APPA network. The LOGGER virtual machine runs a program which interfaces between the network and CP virtual terminals.

Two new multiplexer devices (TYPNCPT and TYPNCPD) have been defined to support virtual terminals. An NCP terminal block (NCPT) has been added to the chain of real device blocks (MRDEFICKs) and a user logged on to such a terminal appears to CF like any other user except at the console (PDCONS and WETCONS) level or lower. RDCONS and WETCONS detect I/O to NCPTs and create CCW packages identical to those made up for 1052's.

Virtual terminal devices (NCFD's) are defined, through the directory, as part of the LCGGEF virtual machine. I/O instructions from the LOGGEF to an NCFD are interpreted according to a protocol for communicating with an NCPT. Initially, an addressed NCPD must be associated with an availabl. NCFT so that a user may log on. Subsequently, information is transferred between the NCPD and an NCPT to satisfy CP read and write creations.

Initial connection is established by a HIO from the LOGGER virtual machine to one of its available NCP devices. This HIO causes CP to search for a free NCP terminal (in the chain of MRDEBLCKs) and to establish a connection between the two. All subsequent I/O operations concern the transfer of data accross this connection until the link is broken.

A connection once established may be broken in one of two ways: by the user typing 'LOGO' or by the LOGGER virtual machine issuing a TIC to its NCP device. The second case is analogous to a user turning power off on a 2741 terminal, and may occur because the LOGGER crashed and came tack up, or because a network link was closed. The first case is detected in the same way as a logoff from any other terminal device is detected. The LOGGER is informed of the broken connection by means of a unit check in his CSW.

A read-type command from the virtual machine to its NCP device is issued to pick up data from CP writes. The data in CP's write buffer is moved into the area beginning at the virtual machine's read address, and the read command is terminated with CE and DE. NCP support follows the protocol for a 1052 in that a 'locked keyboard' system is maintained. Thus, the virtual machine controlling NCP devices should always keep a read-type command active on these devices.

When CP puts up a read to an NCF terminal, the virtual machine is signalled to send data by a unit exception on the NCP device. This sequence is analogous to the write circle C/read circle D needed to unlock the keyboard of a 2741. On receipt of a unit exception, the LOGGER virtual machine should respond by issuing a write-type command to the NCF device.

The I/O operations defined for an NCP device, and the responses to them, are as follows.

TIO: used as a power off signal when a network connection has been closed.

CC=0 means a user was on and has been logged off. CC=1 means no connection existed.

HIO: used as an attention signal. If no connection between the addressed NCP device and an NCP terminal exists, this routine will attempt to find an available NCP terminal and establish a connection. In this case,

CC=1 means a connection was established. CC=2 means there is no available NCP terminal.

If a connection exists, HIO acts like an attention interrupt.

- CC=0 means an ATTN has been simulated, and an interrupt is pending.
- CC=1 means ATTN has been simulated, but the NCP device had no active command, so no interrupt will be received.
- SIO: used to transmit and recieve data. Only read and write are defined.

CC=0 means OF, I/O started.

- CC=1 means CSW stored. Meaningful bits are busy, attention (means warning is coming, and a read should be issued), unit check (means no link to an NCPT exists), unit exception (means CP wants to read, and user should write to his NCPE), and program check (means a read was issued when a write was required, or vice versa).
- CC=2 means device busy.

Interrupts received after a successful SIO may be:

CE+DE: successful termination CE+DE+UE: occurs only on a read-type command, and CF+DE+UF+SM: means that CP is asking for a userid. It allows the LCGGEP to monitor his user's logon synchronously with CP. ATTN: cocurs asynchronously when CP is waiting for the LCGGEP to send data, i.e. when his keyboard is unlocked. It means a warning-type message wants to blast through.

CE+DE+UC: means user has logged off.

## Appendix F

# CP Virtual Machine Communications Facility

## Purrese:

The VMCOM facility allows the virtual user to send data to another virtual user logged on to the CF system, or receive data transmitted from another virtual user.

## Fcrmat:

The VMCOM diagnose instruction is defined as follows:



# where:

RT	-	Pointer to a Control Block
F2	-	nct used
CODE	-	(x'CO 50') used to designate the VMCOM request

# Control Block:

<u>Byte</u>	<u>Data</u>
C	USERID (8)
8	Reguest flag (1)
9	Data buffer address (3)
12	Length of data buffer (2)
14	Number data bytes read (2)

Request Plag Description

.....

x'01'	Write data to user's stack (wait)
x 02 1	Fead data from own data stack
x 1031	Purge data from own data stack (all)
x"C4"	Furge data sent to particular using
x'05'	Test for any data in own data stack
x'06'	Test for data sent to particular upper
x*07*	Write data to user's stack (no wait)
x' 90 '	Lost messaça flag

## <u>Usage</u>:

The receiver of a data write will receive a device end interrupt on virtual device x'fd' when some data has been placed on his data stack. The receiver should then initiate a read request to get the data into his virtual machine. For the writer of data, there is a limit of one data transfer between a virtual machine pair. If this limit is exceeded, a condition code 2 will be reflected on the data write attempt, and the receiver's control block will have the request flag byte set to bl1'1xxxxxx' to indicate data lost. Lata writes will write the whole data buffer out, data reads will read up to the buffer length and then indicate data buffer cverflow. Null writes and null reads are allowed to communicate interrupts between virtual users. A write with no wait will be allowed for class 'C' users only.

## Ending Conditions:

<u>Condition Code</u> <u>Description</u>

	successful completion
CC1	Unsuccessful completion
	Control Block not full word aligned
	Control Block not within virtual core
	Lata buffer not within virtual core
	Invalid request flag specified
	Data buffer > 4096 bytes
	Not a class 'C' user, request invalid
CC2	Lata buffer overflow on read request
	Unprocessed data in receiver's data stack
CC3	Uper not on system
	Unable to reflect interrupt to receiver
	No data in stack on read request

## Appendix G

## The NET Routine

Purpose:

.

Ic communicate with the NCP to issue NCP control commands or to receive a message which has been buffered by the NCP

## <u>Calling Seguence:</u>

The user makes a subroutine call of the form:

LA	1,NETPLIST
L	15,=A (NET)
EALR	14,15
LTR	15, 15
BNZ	15, EPROR

## where NETPLIST is of the form:

CI4	'Control Code'
XI.3	"My ID"
XL1	'Ky Tag'
XL1	'For Host'
XL3	'For ID'
XLI	'For Tag'
XL1	'Flag'
Н	'Eit count'
AL4	(Euffer Address)
F	'Status Reply'
F	'Eits Left'
Н	"Messages left"
X	'Eyte Size'
X	'Spare'
AL4	(Address of interrupt stack)

## and where Control Codes are:

Connect
Listen
Send
Receive
Close
Request Status of Socket
Sent interrupt
returns net status. Useful after wait.
Enable connect
Enable listen
Disable
Purge

a substitute of the second second second second second second second second second second second second second

10-10 h

## <u>lsage</u>:

A send buffer must contain an 80-byte header followed by up to 1000 bytes of data. The NET routine followed by 8C-byte header to build an NCP control message. Similarly, 7. receive buffer must be 1080 bytes long. The first 80 bytes will be used for the NCP control message and the rest of the buffer will be used to store the text of any message which was held by the NCP.

Before any call to NET is made, the interrupt stack printer should be cleared. On return from a call to NET the interrupt stack pointer should be checked. If it is zero, no interrupts were received. A non-zero pointer indicates that an interrupt was received and the interrupt stack control block should be analyzed and then returned to free storage. The Interrupt Stack Control Block is shown in Fig. G.1 and the Interrupt Codes are listed in Fig. G.2.

#### EILOIS:

Control is always returned from the NET package with a code in general purpose register 15. The return codes are as follows:

- 0 Command accepted by NCP. Tag status given in reply.
  1 Command rejected by NCF. Tag status given in reply indicates reason for rejection.
  2 NETPLIST control code invalid. Status reply not significant.
- 3 VMCOM failure.
- 4 Nothing there (NCP Special Case).
- 5 Unexpected message received by NET .
- 6 NCP error.
- Reply does not match socket number issued.  $\epsilon$  VMLOM read error.
- 9 VMCOM send error.
- 10 Nc more free store for NET.
- 11 NCP just died.
- 12 Zero bit count.



Fig. G.1: Interrupt Stack Control Block

## CODE MEANING TN73-50-1(G.2) 0 Interrupt received. Could be SND or RCV socket. Status must be issued to clear this bit. 1 His listen issued on our SND socket. After the connection has been completed. 2 Receive buffer now loaded after issuing our listen (His connect implied). His CLS received before we have issued ours. 3 4 Enabled connect matched by his connect. 5 Enabled listen matched by his connect. 6 NCP Died. 7 Foreign NCP Reset (RST received). Fig. G.2: Interrupt Codes

73

A STREET NO.

#### Appendix H

#### NET User Macros

#### General:

A link block must be defined for each link required. This block contains local and foreign socket numbers together with status and flag areas. The NETFLIST macro is a DSECT which specifies the format. Link blocks are defined by NETSOCK or NETSOCKV macros, and refer either to send or receive links. The user's network id must be obtained with the NETID macro which obtains a user number and places it in the appropriate portion of the socket number.

The link must then be activated via a NETOPEN macro.

Transfers are achieved via NETREAD and NETWRITE macros and the link may be closed via the NETCLOSE macro.

Link status may be obtained using NETSTAT and network interrupts may be sent using NETINT. NETNOP obtains the current flag state without accessing the NCP.

## Macros:

NETSOCK (SA, EB, EC)

Defines a link block. It requires a label which is used by all other macros to identify the link. It has three parameters:

 $\mathcal{E}A$  - a 1-tyte hex value of the local tag  $\mathcal{E}B$  - a 1-byte hex value of the foreign host  $\mathcal{E}C$  - a 4-byte hex value of the foreign socket

NETSOCKV (SA, SB, SC)

Is an alternative to NETSOCK. It also defines a link block and requires a label for link identification purposes. It has three parameters.

&& - a label for the 1-byte local tag &B - a latel for the 1-byte foreign host &C - a latel for the 4-byte foreign sccket

Before using the link specified by NETSOCKV values must be moved into the locations named by GA, BB and BC. For example:

MVI TAG,X'C1' MVI FHOST,X'OA' MVC FSOC(4),=X'COOCOOCO1' . . LINK: NETSOCKV TAG,FRCST,FSOC

and the second state of the se

The following macros require the link block lawel as the first parameter  $(\mathcal{E}A)$ :

- NETID Obtains user's network id and places it in the link block
- NETOPEN Assigns a link and establishes the connection.
- NETOFENE Used after a NETENA to open a connection
- NETREAD Transfers into the specified buffers any outstanding data in the link (max 1000 bytes)
- NEIWRITE Transfers from specified buffers to network. (max 1000 bytes)
- NEISTAY Obtains network status of link
- NET. Issues appropriate network interrupt
- NEIN P Obtains current interrupt state
- NEICLOSE Closes the connection
- NETENA Defires a local socket for which a foreign host can initiate a connection
- NETDISA Releases a local socket which was enabled
- NETPURGE Causes the entry in the NCP RV table to be deleted for a closed connection which has not been closed by the foreign host

The NETREAD and NETWRITE macros have two other parameters. &B is the address of a full word containing the buffer address. &C is the address of a full word containing the fit count.

A NETWAIT may be issued to wait for an NCP interrupt. When the wait completes, a NE' NCF may be issued and the current interrupt(s) will be placed in the interrupt stack.

#### Feturn Codes:

'eturn codes in GPR15 are:

- 0 Operation completed satisfactorily Status field significant
- 1 Operation denied by network. Status field significant
- 2 User made logical error. Status not significant
- 3 System Pailure. Status not significant

A REAL PROPERTY AND A REAL

#### Appendix I

#### The MLCTST Frogram

## Furpose:

This program is used to issue CCW's to the MLC/IMP. The program has a one-character command language. It is useful to issue these commands on one line separated by the logical break character (\*).

<u>Requests</u>:

- I (INPUT) To set the current I/O device to the IMP's input device
- C (OUTPUT) To set the current I/O device to the IMP's cutput device.
- N (ON) To turn the IMP on. This request must be issued on the IMP's input device.
- F (OFF) To turn the IMP off. This request must also be issued on the IMP's input device.
- F (READ) To issue a read SIO. If a read is already outstanding, a 7=2 (busy) results. When the read completes, the data read will be in an input buffer which can be displayed or dumped. A read on the output device causes a CC=1 with the CSW stored indicating a program check.
- W (WRITE) To issue a write SIO. The data written is set up to be a: IMP NOP (a 4-byte leader only). A write on the input device causes a CC=1 with the CSW stored indicating a program check.
- H (HIO) To issue a HIC instruction to the current device.
- T (TIO) To issue a TIO instruction to ""e current device
- X (NOP) To issue a NOP CCW to the current device.
- V (invalid) To issue an invalid CCW operation code to the current device.
- D (DISPLAY) \_ To dump the input buffer which contains the data last read. CP is used to display the storage area.
- P (PAUSE) To issue a CMS WAIT on the input device to wait for an interrupt.

AND REAL PROPERTY.

Ç (QUIT) - To leave the MLCTST program.

State Barrister and and the

## Appendix J

## The TRINET Program

Purrose:

To test an NCP using the NET routine.

## Sequests:

A MARCINE MARTIN

	CON	1s	fh	fs	bs
	LIS	ls	£h	fs	ts
	SND	ls	nes	sag	je
	RCV	ls			
	INT	ls			
	CLS	ls			
	STA	ls			
	ENC	ls			
	ENL	ls			
	DIS	ls			
	PUR	ls			
	END				
where:					

## <u>Usage</u>:

Requests are issued from the terminal to perform a basic NCP function. For each request issued, the RV table entry is returned if one exists. The second two status bytes are printed indicating the status of the socket connection.

Whenever an interrupt is received, a line is printed giving the type of interrupt. When an RCV is issued, the message received, if any, is also printed. If an NCP communications message is lost (VMCOM interrupt lost), this fact is also printed. The user should then issue a status on each of the active sockets to determine the interrupt which was lost.

## Appendix K

#### The TELNET LCGGER/SERVEF

#### ICP Connection:

The TELNET LOGCER/SERVER follows the ICP protocol for making a pair of connections. The LOGGER is initially enabled for a connection on socket X'00000001'. When an RFC is received for this socket a pair of sockets will be chosen for the TELNET connections. If the maximum number of TELNET users which can be served are active, the initial connection is refused. Currently, three TELNET users can be served.

## TELNET LOGGER:

Television of the same

At the stime . I have

After the ICP connections have been setup, the LOGGER expects a TELNET data type code, a string of network ASCII characters, or a null line (just CR-LF) to indicate whether its operation should be in ASCII or in EBCDIC character codes. ASCII is assumed unless the first bype received is the TELNET EBCDIC data type code (X'A2). When something has been received, the message:

#### Lincoln Laboratory CP/CMS Online

will be transmitted by the ICGGEF. For example, if ASCII operation is desired a null line (just CR-LF) transmitted on the send socket will cause the welcoming message to be sent in ASCII. The CP login procedure can then begin. If communications is desired to be carried on with EBCDIC character codes, the first byte transmitted shculd be the TELNET data type code for EECDIC (X\*A2\*). Thereafter all communications will be in the code originally used.

The CP login procedure expects the user to enter:

#### LOGIN userid

where the userid specifies the desired virtual machine. CP then replies with:

#### ENTER PASSWORD:

followed by the EBCDIC code for byp.ss (X\*24\*) which is mapped into the TELNET code hide-your-input.

The user should then enter a password. Fasswords entered from the network may be different from those entered from a local terminal. The LOGGER maps network passwords into a corresponding CP password. Thus, access to an account can only be made from the network if a network password, together with a CP password and userid, is entered into a file which is read by the LOGGER. If a userid entered from the network is not in

the LOGGER FILE (or if the network password does not match the one included in the file for the specified userid) the IOGGER passes an invalid userid (or password) to CP. The CP response for an invalid userid or password is then sent to the network user.

After a password is received by CP, CP transmits the EBCDIC code for restore (X\*14\*) which is mapped into the TELNET control noecho.

#### <u>**TELNET SERVER:**</u>

Since the CP/CMS systom operates with EFCDIC codes, ASCII codes must Le translated into EBCDIC before being sent to a virtual machine. Figure K.1 gives the ASCII codes and their EBCDIC mapping. When the ASCII sequence CF-LF is received, it is mapped into the EBCDIC code NL. Whenever the TEINET control NOP is included in an input string, it is mapped into an EBCDIC idle (X'17') and then removed from the string. Thus, if TELNET NOP codes are included between a CR and LF, they are removed before the CR-LF is mapped into the EBCDIC NL.

The TELNET control hide-your-input is mapped into the EECDIC code for bypass (X'24') and the TELNET control echo is mapped into the EBCDIC control for restore (X'14'). If the TELNET control echo is received, the SERVER should send the control noecho but this feature has not yet been implemented. Instead, the TELNET control echo is mapped into the EBCDIC code X'23'. If the TELNET break is received, it is interpreted as an attention signal and the appropriate action is taken by CP or CMS.

CP/CMS is a line at a time system and expects all input to consist of lines ending with a NL code. Characters received are buffered until the newline code is received.

Since CP/CMS is also a half duplex system, characters are only examined when the system is expecting input. If the system is not expecting input, a network interrupt is required to cause the SERVER to process received characters. On receipt of a network interrupt, characters received before the TELNET data mark is received are examined and discarded, except that if a TELNET break code is found, the appropriate CP,/CMS interrupt action is stimulated.

On output, EECDIC codes are mapped into network ASCII if a mapping exists; otherwise, the codes are mapped into the TELNET control NOP. A NL code is mapped into CR-LF. The EECEIC code for uypass maps into the TELNET control hide-your-input and the EBCDIC code for restore maps into the TELNET control noecho. Also, the code X'23' maps into the TELNET control echo and the code X'38' maps into the TELNET control break.

Since CP/CMS is a line at a time, half duplex system the TELNET control break is transmitted as an end of message signal and also as an input prompt code. If characters were output without a NL, the break, as an end of message code, indicates to the user TELNET operating on a line at a time mode that the characters previously transmitted should be printed without waiting for the end of line sequence. If the user TELNET is also operating in a half duplex mode, the break as an input prompt indicates that the system is ready for input.

If input had been anticipated and sent by a full duplex user TELNET, the TELNET SERVER will have that input available for immediate processing. Thus, in the case of a full duplex user TELNET the break as a prompt should be ignored.

Though CP/CMS operates in a half duplex mode, it supports half duplex terminals with the reverse break feature allowing the system to abort an input mode in order to transmit a priority output message. In this situation, the TELNET SERVER transmits a TELNET SYNC. A half duplex user TELNET should interpret this by aborting the input mode, i.e., revoking a previous TELNET break which was interpreted as an input prompt.

No codes in the output character stream can cause the TELNEI data mark to be transmitted.

LCGCUI:

When a user logs out from his virtual machine, CP passes the equivalent of a line disconnect to the LOGGER. The LOGGER then closes the TELNET send and receive sockets.

internation of alles

TN73-	-50-I(K.1	.)			
ASCII	ASCI	ASCII	SYMBCLS	EECDIC	FECDIC
DEC	OCT	HEX		HEX	DEC
c	c	(00)	NUL	(00)	00
1	1	(01)	SOH	(01)	01
2	2	(02)	STX	(0.2)	02
3	3	(03)	ETX	(03)	03
4	4	(04)	EOT	(37)	55
5	5	(05)	ENO	(2D)	45
6	6	(06)	ACK	(2E)	46
7	7	(07)	BEL	(21)	47
8	10	(08)	BS	(16)	22
C,	11	(09)	HT	(05)	05
10	12	(OA)	LF	(25)	37
11	13	(CB)	TV	(0B)	11
12	14	(CC)	FF	(0C)	12
13	15	(CD)	CR	() D)	13
14	16	(OE)	SO	(9E)	14
15	17	(OF)	SI	(0F)	15
16	20	(10)	DLE	(10)	16
17	21	(11)	DC 1	(11)	17
18	22	(12)	DC2	(12)	18
19	23	(13)	DC 3	(13)	19
20	24	(14)	DC 4	(3C)	60
21	25	(15)	NAK	(3D)	51
22	26	(16)	SYN	(32)	50
23	27	(17)	ETB	(26)	38
24	30	(18)	CAN	(18)	24
25	31	(19)	ΞM	(19)	25
26	32	(1A)	SUB	(3F)	63
27	33	(18)	CTL	(27)	39
28	34	(10)	FS	(1C)	28
29	35	(TD)	GS	(1D)	29
30	30	(12)	FS	(1E)	30
21	37	(12)	05	(15)	15
Fig.	K.1:	LOGGER AS	CII, EBODIC	Code Mar	ping

TN73-	50-I(K.1	1)			
ASCII DEC	ASCI OCT	I ASCII HEX	SYMBOLS	EBCDIC HEX	EBCDIC DEC
32	40	(20)	SP	(40)	64
33	41	(21)	1	(5A)	90
34	42	(22)	21	(7F)	127
35	43	(23)	*	(7 E)	123
36	44	(24)	\$	(5B)	91
37	45	(25)	<b>%</b>	(60)	108
38	46	(26)	3	(50)	80
39	47	(27)	•	(7D)	124
40	50	(28)	(	(4 E)	77
41	51	(29)	ý	(5 D)	93
42	52	(2A)	*	(50)	92
43	53	(2B)	+	(4 E)	78
44	54	(2C)	,	(6 D)	109
45	55	(2D)	-	(* - )	96
46	56	(2E)	•	<i>i</i> 4 y	75
47	57	(2F)	1	1611	97
48	60	(30)	Ó	(F0)	240
49	61	(31)	1	(F1)	241
50	62	(32)	2	(F2)	242
51	63	(23)	3	(F3)	243
52	64	(34)	4	(F4)	244
53	65	(35)	5	(F 5)	245
54	66	(36)	6	(F6)	246
55	67	(37)	7	(F7)	247
56	70	(38)	8	(F 8)	248
57	71	(39)	ġ	(F9)	249
58	72	(3A)	•	(7))	122
59	73	(3B)	:	(5E)	94
50	74	(3C)	<	(4C)	76
61	75	(3D)	2	(7E)	126
62	76	(3E)	>	(6E)	110
63	77	(3F)	?	(6F)	111
Fig.	K.1:	LOGGER AS	CII/EBCDIC tinued)	Code Mar	ping

12/24

T'N7	<b>3-</b> 50-1 (K	.1)				
ASCII	ASC I	I ASCII	SYMBOLS	EPCDIC	FPCT	Tr
<b>EEC</b>	OCT	HEX		HEX	DEC	•
64	100	(40)	æ	(7C)	124	
65	101	(41)	Α	(01)	193	
66	102	(42)	В	(C2)	194	
67	103	(43)	С	ic3j	195	
68	104	(44)	D	(C4)	196	
69	105	(45)	Ē	(05)	197	
70	106	(46)	F	icsi	198	
71	107	(47)	G	(C7)	199	
72	1 10	(48)	Н	(08)	200	
73	111	(40)	I	(00)	201	
74	112	(4A)	J	(01)	209	
75	113	(4B)	K	(D2)	210	
76	114	(4C)	L	(D3)	211	
77	1 15	(4D)	М	(D4)	212	
78	1 16	(4E)	N	(D5)	213	
79	1 17	(4F)	с	(D6)	214	
80	120	(50)	P	(07)	215	
81	121	(51)	ç	(28)	216	
82	122	(52)	R	(D9)	217	
83	123	(53)	S	(E2)	226	
84	124	(54)	Т	(E3)	227	
85	125	(55)	U	(E4)	228	
86	126	(56)	V	(E5)	220	
8 <b>7</b>	127	(57)	W	(E6)	230	
98	130	(58)	8	(E7)	231	
89	131	(59)	Y	(E8)	232	
90	132	(5A)	2	(E9)	233	
91	133	(5E)	ſ		177	
92	134	(5C)	R	(4A)	74	(BACK-STASH)
93	135	(5D)	1	(BD)	189	(brea STRON)
94	136	(55)	-	(71)	113	(ראקאס)
95	137	(5F)	-	(3D)	109	(
Fig.	K.1:	LOGGER ASC	II/EBCDIC	Code Map	ping	

TN73-	50-I(K.1)					
ASCI DEC	I ASCI OCT	I ASCII HEX	SYMBCLS	EBCDIC HEX	EPCDI DEC	c
96 97 98 99 100 101 102 103 104 105 106 107 108 109 110 111 112 113 114 115 116 117 118 119 120 121 122 123 124 125 126 127	140 141 142 143 144 145 146 147 150 151 152 153 154 155 156 157 160 161 162 163 164 165 166 167 170 171 172 173 174 175 176	(60) (61) (62) (63) (64) (65) (66) (67) (68) (69) (68) (60) (61) (62) (62) (62) (62) (62) (62) (62) (62) (62) (62) (62) (62) (62) (62) (62) (62) (62) (62) (70) (71) (72) (73) (74) (75) (76) (77) (78) (77) (78) (72) (77) (78) (77) (78) (77) (77) (78) (77) $(77)($	a b c d e f f f h i j k l m n o P g r s t u v w x y z { ! ] D EL	(7 9) (9 1) (8 2) (8 3) (8 4) (8 5) (8 6) (8 7) (8 8) (9 1) (9 2) (9 3) (9 4) (9 2) (9 3) (9 4) (9 5) (9 6) (9 7) (9 8) (9 9) (A 2) (A 3) (A 4) (A 5) (A 6) (A 7) (A 8) (A 5) (A 6) (A 7) (A 8) (A 9) (8 B) (4 F) (9 B) (5 F) (0 7)	$\begin{array}{c} 121\\ 129\\ 130\\ 131\\ 132\\ 133\\ 134\\ 135\\ 136\\ 137\\ 145\\ 146\\ 147\\ 148\\ 150\\ 151\\ 152\\ 163\\ 166\\ 167\\ 168\\ 169\\ 139\\ 155\\ 95\\ 7\end{array}$	(GRAVE) (BAR/OR) (TILDE/NOT)
A SCII DEC	ASCII OCT	ASCII HEX	TELNET CONTROLS	EBC DI HEX	C 70B DE	CDIC
128 129 130 131 132 133 Fie	100 101 102 103 104 105 g, K.1:	(80) DA2 (81) BR1 (82) NC1 (83) NC1 (83) NC1 (83) NC1 (85) HI1 LOGGER AS	TA-MARK EAK ECHC HO DE-YCUR-IN SCII/EECDI	(80) (38) (17) (14) (23) PUT (24) C Code Map	12 5 2 3 3 9 9	9 6 3 IDLE 0 RESTORE 5 6 BYPASS
		(Cor	ntinued)			

A MARCH STREET

## Appendix L

#### The User TELNET

## Purrese:

To access another terminal-oriented system on the ARPA network.

## Fcrnat:

## TELNET host <tag> RESUME EBCDIC HALFLUP <u>1 OPEN ASCII FULLDUP</u>

- host either the hexadecimal code for a foreign network service site or a standard mnemonic for a foreign site. See Figure 1.
- tag the identifier for the local connections to the network. The tag is used together with the address of the virtual machine descriptor table (UTABLE) to form local socket numbers which are used in the network protocol.
- RESUME used to reactivate communications with a foreign site after having previously left the TELNET command leaving the connections open.
- EBCEIC to communicate with EBCDIC codes. The default is network ASCII.
- HALFDUP to operate under a half duplex protocol, i.e. with a locked keyboard.

The EECDIC HALFDUP the protocol assumes that the TELNET break code (circle C) will be received to indicate when the keyboard should be locked for input.

In ASCII HALFDUP the keyboard will lock after a line of input and will unlock after one or more lines have been received for output. An external interrupt will also unlock a locked keyboard.

The default is full duplex where the keyboard is alwasy unlocked for input. A null line is required to temporarily lock the keyboard in order to receive output.

#### <u>Usage</u>:

THE CASE AND AND AND A DECKS

A number of hosts on the ARPA network provide TELNET service. A Network Virtual Terminal (NVT) has been specified so that using sites can write one TELNET program which maps a local terminal into the NVT to access any serving site on the network. Once communication has been established between a using site and a serving site, keyed input is sent to the serving system and output from the serving site, when received, is typed on the local terminal.

The NVT protoccl requires that the keyboard be capable of entering all of the 128 ASCII codes together with a number of other TEINET control codes. To support an NVT with an IBM 2741 terminal, it is necessary to adapt a control convention for entering codes which are not associated with single keys on the keyboard. In addition, since CP/CMS processes input from a 2741 on a line at a time terminated with a newline, a means must be established for entering a sequence of characters for transmission which is not terminated with a newline code.

When TELNET is ' ... 'ated the message

#### FRIER CONTFOL CHARACTEF

is typed. A non-blank character should then be entered which defines the character which, in combination with another character, will be used to enter codes not associated with single keys. The control character is al~c used for other special control functions as described below.

and an and the second second second second second second second second second second second second second second

## <u>Ccdes</u>:

The NVI usually requires that characters be transmitted in an eight-bit ASCII code. Since the TELNET command is written to process EBCDIC codes ASCII codes received are translated into EPCDIC and characters to be transmitted are translated into ASCII before being sent to a serving site. Figure 2 gives the complete definition of EBCDIC indicating the EPCDIC controls and EFCTIC graphics. Figure 3 gives the codes for the ASCII controls and graphics. The complete mapping between 8-bit EFCDIC codes and 8-bit network ASCII codes is shown in Figure 4. The EBCDIC newline code (NL) is mapped into the ASCII codes for the pair of characters CR-IF.

The fcllwcing ASCII/EBCDIC mapping is used for the non-FECDIC graphics:

ASCII EBCDIC

TILDE (7E) = (A1) NOT BAR (7C) = (6A) OF BACK SIASH (5C) = (6A) OF CARAT (5F) = (71)GRAVE (60) = (79)LEFT BRACE (7B) = (8B)RIGHT BRACE (7D) = (9B)LEFT BRACKET (5B) = (AD)RIGHT ERACKET (5D) = (ED)

The ASCII control DC3 (X'13') maps to the EBCDIC control TM (X'13'). The ASCII control NUL (X'00') is sent to the terminal as the EBCDIC code for NULL (X'00') and is not mapped into an IDLE (X'17').

The TELNET control hide-your-input is mapped into the EECDIC code for bypass (print supress) and the TELNET control noeche is mapped into the EECDIC code for restore (print restore). If the TELNET control for eche is received, a message is printed and it is mapped inte an IDIE. Similiarly, if the TELNET control for break is received, a message is printed and it is mapped into an IDLE unless operation is in EBCDIC HALFDUP mode in which case the break is used to indicate that any received characters should be printed and the keyboard unlocked for input. If a data mark or an interrupt is received, no action is taken except to print a message to notify the user of this occurance.

Irput:

When the control character is entered, the following character is mapped into a different code than that which it is normally mapped into, except when the following character is a space or a character not defined to have a meaning when preceded by the control character. Figure 5 gives the mapping of the characters on a 2741 keyboard when preceded by a control character. The following 2741 keyboard characters do not have a different meaning when preceded by the control character.

```
$ # * % &
+ - =
. . : :
! | ? ¢
SPACE
BACKSPACE
TAB
```

When a character is mapped into its control code, the control character is mapped into the code for IDLE. If the control character is entered as the last character before the newline key is entered, the sequence of characters entered is transmitted without the newline code. That is, the newline code is not transmitted when it is proceded by the control character.

When the 2741 keyboard is unlocked for input, characters received cannot be typed until the keyboard is locked again. After a line is entered, received characters can then be typed. When operating in full duplex or ASCII half duplex, a null line entered will allow received characters to be typed but will not cause the new line code to be transmitted. To cause a null line, i.e., just the new line code to be transmitted, the control character should be entered as the only character in the input line. In EBCDIC HALFDUP a null line entered will cause a null line to be transmitted.

## <u>Cutrut</u>:

ASCII cutput received from the NVT is converted into EBCDIC with the sequences CR-LF converted into IDLF-NL. The EBCDIC characters are then sent to the terminal. Note that not all 128 ASCII codes when converted to EBCDIC will print on a 2741. Of the 95 ASCII graphics and the 8 ASCII controls which are defined for the NVT printer, the following are not visible or audible:

> CARAT GRAVE EACK SLASH LEFT BRACE RIGHT BRACE LEFT BRACKET RIGHT BRACKET ASCII CONTROL BELL (PEL) ASCII CONTROL VERTICAL TAB (HT) ASCII CONTROL FORM FEED (FF) ASCII CONTROL CARRIAGE FETUEN: (CF)

Figure 6 shows how the FBCDIC codes from X\*40\* through X\*FF will appear on a 2741 terminal. Figure 7 shows how the EBCDIC codes will appear when printed with a PN train on the offline printer and Figure 8 shows how these codes appear when printed with a TN train.

#### <u>Controls</u>:

If the first character in an input line is the control character and the next character is a space, the rest of the line is interpreted as a TELNET control command. A control command consists of a control word and parameters separated by spaces. Controls are defined which permit TELNET controls to be transmitted to the serving site, allow input to come from a file or output to go to a file, allow CMS functions or transient commands to be issued, redefine the control character or TELNET mode, close connections or leave the TELNET command with the connections still open, as well as controls to support a reader, punch, and printer with PJS operation. The controls are described below.

#### CONTROL X

Where x is the new control character

CIOSE

To close all connections and guit

**TIUQ** 

To leave TELNET

EECDIC

To go into transparent mode, i.e., no translation ASCII

To translate input and output to network ASCII BREAK

To send the TEINET break code

SYNC

To send the TEINET data mark code and an interrupt

To send a TEINET break and a SYNC

HIDE-YOUR-INPUT

To send the TEINET hide you input code.

NCECHC

Tc send the TEINET noecho code

ECHC

To send the TEINET echo code.

CMS command arg1...argN

To issue CMS core resident function or transient command.

INPUT fn ft \* TERMIN \* \*

> To get input from a file If fn is defaulted, input is reset to come from the terminal. If fn is \* file input resumes after the last line read. After an EOP, the next line read will be the first line of the file.

> An external interrupt while input is coming from a file will cause the line number of the next line to be read from the file to be typed and input to be reset to come from the terminal.

OUTPUT fn OFF TERM INPUT INCUT \* ON NOTERM NOINPUT OUTPUT

> To write output to the file 'fn TERMOUT'. If fn is defaulted, output is reset to go to the terminal. If fn is \*, file OUTPUT is resumed with the same options as were last used.

For Output to the Terminal:

If the last character is a CR, a line with just the control character is typed on the next line (with a NL)

If the last character is not a NL or a CR, the line is typed without a NL (i.e., with TYPE).

For Output to a File:

If just a NL is in the line, just the control character is sent to the file.

If the last CHAP is not NL or CR, the control character is added after the last character, except if 130 characters must be sent to the file.

If the last CHAR is a CR, it is included in the file.

OFF causes all output to be discarded.

CN is the default, and causes output to go to the terminal.

IERM causes output to also go to the terminal.

NOTERM is the default, and causes output to go the the file but not to the terminal.

OUTPUT is the default and causes just terminal output to be put to the file 'FN termout'.

INOUT causes both terminal input and terminal output to be put to the cutput file.

INPUT causes terminal input but not output to be put to the cutput file.

NOINPUL is defaulted and causes input to not go to the file.

ALL LOOK STREET

PURGE

To purge all output currently received by the NCP. This control is not currently implemented. BEADEF fn ft \* READER Ic send a job to the RJS system at UCLA'S CCN. If fn and ft are defaulted, input will come from the card reader. FFINTEF fn ft \* PRINTER To receive printer output from the PJS system at UCLA's CCN. To receive punch output from the RJS system at UCLA's CCN. If fn and ft are defaulted, output goes to the printer. PUNCH fn ft \* PUNCE If fn and ft the defaulted, output goes to the punch.

93

and and the head the

TN73-5	0-I(L.1)					
HOSI	SITE	MACHINE	SYSTEM	HO S DEC	ST NUM OCT	BER HEX
NMC ARC UCSE UTAH MULTICS SDC HARV LL CASE CMU ILLIAC AMES CCN SRI	UCLA SRI UCSE UTAH MIT SIC HARVARD LI CASE CMU AMES AMES UCLA SEL-AT	SIGNA-7 PDP-10 360/75 PDP-10 H-645 370/155 PDP-10 360/67 PDP-10 B-6500 360/67 360/91 PDP-10	SEX NIC OS/MVT TENEX MULTICS ADEPT 4S72 CP/CMS 10/50 TOPS-10 ? TSS/360 OS/MVT TENEX	1 2 3 4 6 8 9 10 13 14 15 16 5 5	1 2 3 4 6 10 11 12 15 16 17 18 101	01 02 03 04 06 08 09 04 00 00 00 00 00 00 00 00 00 00 00 00
EENA DMCG RAND TX2	BEN MIT RAND-FCC LL	PDP-10 PDP-10 PDP-10 TX-2	TENEX TENEX ITS TENEX	69 70 71	105 106 107	42 45 46 47
EENE MITAI	BEN MIT	PDP-10 PDP-10	TENEX	133 134	205	48 85 86

-----

Fig. 1.1: Serving Hosts on the ARPA Network

94

T. . B. . .....

TN73-50-I(L.2)	
----------------	--

-

	01	0	0	C 0	0 0	0 1	0		0 1	0		1	1	1		1	1	1	1 1	1	
	23	C	1	0	1	0 0	1		0	1		0	1	0		1	0	1	°,	1	
4567 0000	+	+ NUL	DLE	 [][]		SP	+ 3			+ 	+-		+	+	- 4   	0	•\ 	•	+	0	+
000 1		SOH	DC1	SOS			+   +		/		-+	a	1 3	0	• •   •	1	A	J		1	1
0010		STX	DC 2	FS	SYN							b	k	l s	1	2	В	K	S	2	I
0011	Ì	ETX	TM									с	1	t		3	C	L	Ī	3	1
0100	l	PF	RES	PYP	PN					1	1	đ	1 m	l u	,		D	1.1	U	1 4	1
0101	Ì	PT	NL	LF	RS					1	-+-   -+	е	n	7		୍ଟ	E	N	V	5	Ì
0 110		LC	ES	ETB	UC .		1					f	0	K		6	P	0	W	16	1
(111	1	DEL	IL	ESC	EOT		!					g	P	X		7	G	P	X	7	1
1000	I	GE I	CAN				1					h	l g	і у		8	H	0	Y	8	I
1001	1	FLF	EM				1					i	[	2		9	I	R	12	19	1
1010	l	SMMI	сс	SM		¢	!			1 :			1	1							I
1011		VT I	CU 1	CU2	1043		1 \$	5	,	#		{	}	L						1	ļ
1 100	1	PF	IFS		DC4	<	1 *		8	6		<		r		1 1		1		!	1
1 10 1	ļ	CR I	IGS	ENŲ	INAK	(	1)					(	)	] [		]					1
1110		S0	IPS	ACK		+	;		>	=		+	±	,   ≥		.¥.			1		1
1111	1	SI	IUS	PEL	SUB		-	•	?	"	-+	+	•	•	- 4			 +	 +===	 	1

1011213141516171 Code Structure

Fig. 1..2: The EBCDIC Code

TN	7350-	I(L <b>.3</b> )	1_						
" 224	8 0 7 0 6 0 5 0	0000	C C 1 C	0 0 1 1	0 1 0 0	0 1 0 1	0 1 1 0	C 1 1 1	
0000	+	LIDLE	+~	+	-+	-+	• • • • •	+	4
0001	+   SO I	ELDC 1	+	+	-+   A	-+ I Q	-+   a	+   g	   
0010	IST	KIDC2	+	1 2	-+   B	-+   E	+	+	+
0011	+	IDC 3	+   #	1 3	C	-+   S	+   c	+	
0100	EOI	IDC4	I S	1 4	D	-+   T	+	++   t	•
0 1/) 1	+	INAK	+	+	+	-+   U	+	+-«-+	•
0110	ACK	ISYN	.3	16	I P	V	+	++   ▼	
0111	BE I	+	+	+	G	-+   W	+	t+   w	
1000	BS		(	8	I H	I X	h	+   X	
1001	HT +	EM	)	+	+	+   Y	+	+ У [	
1010	LF	SUB	*	1:	1.2	1 2	I j	z	
1011	+	++  ESC  ++	+	+	+   K	+   [	+	+	
1100	FF	IFS I		<	L	1	1	+	
1101	ICR	ics		+   = 	+	+	+ m	3	
1110	150	IRS	• 1	>	I N	1 (	n	+	
1111	ISI	US	/	?	0 	++	0	DEL	
	4							+	
	1 8 1	7 1	6	5	4	++	2 1	+	
	+	++	+ Cod	e 51	+	++	+	+	
			u		L L L L	CULC			

Fig. L.S: The USASCII Code

TN73	-50-I(L.	4)			
ASCII	ASCII	ASCII	SYMBOLS	EBCDI	C EBCDIC
DEC	OCT	HEX		HEX	DEC
0	0	(00)	NUI	(00)	00
1	1	(01)	SOH	(01)	01
2	2	(02)	STX	(02)	02
3	3	(03)	ETX	(03)	<b>n</b> 3
4	4	(04)	EOT	(37)	55
5	5	(05)	ENÇ	(2D)	45
6	6	(06)	ACK	(2E)	46
7	7	(07)	PEL	(2F)	47
9	10	(08)	ES	(16)	22
9	11	(09)	HT	(05)	05
10	12	(CA)	LF	(25)	37
11	13	(OB)	VT	(0 B)	11
12	14	icci	FP	(00)	12
13	15	(C D)	CR	(0D)	13
14	16	(OE)	<b>°</b> 0	(0E)	14
15	17	(CF)	: T	(CP)	15
16	20	(10)	DLE	(10)	16
17	21	(11)	DC 1	(11)	17
18	22	(12)	DC2	(12)	18
19	23	(13)	DC 3	(13)	19
20	24	(14)	DC4	(3C)	60
21	25	(15)	NAK	(3D)	61
22	26	(16)	SYN	(32)	50
23	27	(17)	ETB	(26)	38
24	30	(18)	CAN	(18)	24
25	31	(19)	EM	(19)	25
26	32	(1A)	SUB	(3F)	63
27	33	(18)	CTI	(27)	39
28	34	(10)	FS	(10)	28
29	35	(10)	GS	(1D)	29
30	36	(1E)	FS	(1E)	30
31	37	(1F)	05	(1F)	31
Fig.	L.4:	TEINET AS	CII/EECIDC	Code !	Mapping

.

TN73	-50-I(L.	4)			
ASCII	ASCII	ASCII	SYMBOLS	EBCDIC	EBCDIC
DEC	OCT	HEX		HEX	DEC
32	40	(20)	SP	(40)	64
33	41	(21)	1	(5A)	90
34	42	(22)	**	(7F)	127
35	43	(23)		(7 B)	123
36	44	(24)	\$	(5B)	91
37	45	(25)	4	(6C)	108
38	46	(26)	3	(50)	90
39	47	(27)	•	(7D)	124
40	50	(28)	(	(4D)	77
41	51	(29)	j	(5D)	93
42	52	(2A)	*	(50)	92
43	53	(2B)	+	(4 E)	78
44	54	(2C)		(6 D)	109
45	55	(2D)	_	(6 ())	96
46	56	(2E)	•	(4B)	75
47	57	(2F)	1	(61)	97
24 R	60	(30)	Ċ	(F 0)	240
49	61	(31)	1	(F1)	241
50	62	(32)	2	(F2)	242
51	63	(33)	3	(P3)	243
52	64	(34)	4	(F4)	244
53	65	(35)	5	(F5)	245
54	66	(36)	E	(F6)	246
55	67	(37)	7	(F7)	247
56	70	(38)	3	(F 8)	248
57	71	(39)	9	(F°)	249
58	72	(3A)	:	(7A)	122
59	73	(3B)	;	(5E)	94
60	74	(3C)	<	(4)	76
61	75	(3D)	-	(7E)	126
62	76	(3E)	>	(6E)	1 10
63	77	(3F)	?	(6F)	111
Fig.	L.4:	TEINET AS	CII/EBCIDC ontinued)	Code Maj	cping

TN73	3-50-i(L	.4)				
ASCII	ASCI	I ASCII	SYMBOLS	EBCD	IC EBCD	IC
	001			пел	DEC	
64	100	(40)	â	(7C)	124	
65	101	(41)	A	(c1)	193	
66	102	(42)	В	$i(C_2)$	194	
67	103	(43)	С	(C3)	195	
68	104	(44)	D	(C4)	196	
69	105	(45)	E	(C5)	197	
70	106	(46)	F	(C6)	198	
71	107	(47)	G	(07)	199	
72	110	(48)	Н	(08)	200	
73	111	(49)	I	ic 9j	201	
74	112	(4A)	J	(D1)	209	
75	113	(4B)	F	(D2)	2 10	
76	114	(4C)	. L	(D3)	211	
77	115	(4D)	Ľ.	(D4)	212	
78	116	(4E)	N	(D5)	213	
79	117	(4F)	0	(D6)	214	
80	120	(50)	Р	(D7)	215	
81	121	(51)	Ç	(D8)	216	
82	122	(52)	R	(D9)	217	
83	123	(53)	S	(E2)	226	
84	124	(54)	T	(E3)	227	
85	125	(55)	U	(E4)	2 2 8	
86	125	(56)	V	(E5)	229	
87	127	(57)	W	(E6)	230	
88	130	(58)	8	(E7)	231	
89	131	(59)	Y	(BE)	232	
90	132	(5A)	2	(E9)	233	
91	133	(5B)	[	(AD)	173	
92	134	(" C)	R.	(4A)	74	(BACK-SLASH)
93	135	(5D)	]	(BD)	189	
94	136	(5E)		(71)	113	(CARAT)
95	137	(5F)	-	(6D)	109	
Fig.	L.4:	TELNET ASC	CII/EBCIDC	Code	Mapping	
		100	ontinued)			
TN73-	50-I(L.4)					
--	---	--	--	--	--	--
ASCI: DEC	I ASCI 007	I ASCII HEX	SYMBOLS	2BCDIC HEX	E BC DI DEC	С
96 97 98 99 100 101 102 103 104 105 106 107 108 109 110 111 112 113 114 115 116 117 118 119 120 121 123 124 125 126	$\begin{array}{c} 149\\ 141\\ 142\\ 143\\ 144\\ 145\\ 146\\ 146\\ 151\\ 152\\ 153\\ 155\\ 155\\ 156\\ 161\\ 162\\ 165\\ 167\\ 177\\ 172\\ 177\\ 177\\ 177\\ 177\\ 177\\ 17$	(60) (61) (62) (63) (64) (65) (66) (66) (68) (68) (62) (62) (62) (62) (62) (62) (62) (62) (62) (62) (72) (73) (72) (73) (74) (75) (76) (77) (79) (77) (72) (77)	a b c d e f g h i j k 1 m n o F g r s t u v w x y z {   } r E I	(79) (81) (82) (83) (84) (85) (86) (87) (87) (89) (91) (92) (93) (94) (93) (94) (95) (96) (97) (98) (97) (89) (89) (89) (89) (89) (81) (93) (94) (93) (94) (95) (96) (97) (88) (87) (89) (97) (87) (87) (87) (97) (87) (87) (97) (87) (87) (97) (97) (87) (87) (97) (97) (87) (87) (97) (97) (97) (97) (97) (97) (87) (97) (97) (97) (97) (97) (97) (97) (9	$\begin{array}{c} 121\\ 129\\ 130\\ 131\\ 132\\ 133\\ 134\\ 135\\ 136\\ 137\\ 145\\ 146\\ 147\\ 148\\ 149\\ 151\\ 152\\ 163\\ 164\\ 165\\ 166\\ 167\\ 168\\ 169\\ 139\\ 155\\ 95\\ 7\end{array}$	(GRAVE) (DAR/OR) (TILDE/NOT)
ASCII DEC	ASCII OCT	ASCII HEX	TELNET COMIROLS	EBCT1 HEX	C E P DE	CDIC C
128 129 130 131 132 133	100 101 102 103 104 105	(8C) DAT (81) BRE (82) NOF (83) NOE (84) ECH (85) HID	A – MARK AK CHO IO V VOUF – INH	(80) (38) (17) (14) (23) PUT (24)	12 5 2 3 3	8 6 3 IDLE 0 RFSTORE 5 6 BYPASS
Fig	g. 1.4:	TEINFT AS	CTI/EBCIDC ontinued)	Code Map	ping	

190

A SALAN

TN73-50-I(L.5)

	FBCD IC	EBCDIC	ASCII
CINT	(4A) = ESC	(27)	(1B)
CTL < CTL > CTL ( CTL ) CTL / CTL / CTL '	(4C) = IEFT PFACKE (6E) = RIGHT BRACK (4D) = IEFT PFACE (5D) = RIGHT FFACE (61) = BACK SLASH (7F) = CARAT (7D) = GRAVE	T (AD) ET (BD) (8B) (9B) (4A) (71) (79)	(5B) (5D) (7B) (7D) (5E) (60)
CTL 6 CTL 7 CTL 8 CTL 9 CTL _	(F6) = FS (F7) = GS (F9) = RS (F9) = US (6D) = US	(1C) (1D) (1E) (1F) (1F)	(1C) (1D) (1E) (1F) (1F)
CTL -	(5F) = DEL	(07)	(7F)
CTL D CTL D CTL C CTL D CTL D CTL F CTL G CTL H CTL I CTL J CTL K CTL L CTL N CTL N CTL O	(7C) = NUI (C1) = SOH (C2) = STX (C3) = ETX (C4) = FOT (C5) = FNQ (C6) = ACR (C7) = BEL (C8) = BS (C9) = HT (D1) = LF (D2) = VT (D3) = FF (D4) = CR (D5) = SO (D5) = SI	(00) (01) (02) (03) (2D) (2E) (2E) (2F) (16) (05) (25) (0B) (0C) (0D) (0C) (0D) (0F)	(0J) (01) (02) (03) (04) (05) (05) (05) (05) (05) (05) (05) (05
CTL P CTL Q CTL P CTL S CTL T CTL U CTL V CTL V CTL Y CTL Z	(D7) = DLE (D8) = DC1 (D9) = DC2 (E2) = DC3 (E3) = DC4 (E4) = NAK (E5) = SYN (E6) = ETB (E7) = CAN (E8) = EM (E9) = SUB	(10) (11) (12) (13) (3C) (3D) (32) (26) (18) (19) (3F)	(10) (11) (12) (13) (14) (15) (16) (17) (18) (19) (1))

Fig. L.5: 2741 Keyboard Control Character Mappings

A REAL PROPERTY AND A REAL

TN73 50-I(L.5)

## FECDIC EBCDIC ASCII

CTL CTL CTL CTL CTL	1 2 3 4 5	(F1) (F2) (F3) (F4) (F5)	 BREAN NOP NO EC ECHO HIDE	YOU YOU	(38) (17) (14) (23) INPUT (24)	(81) (82) (83) (84) (85)	• •	CIRCLE C IDLE RESTORE BYPASS

DATA MARK (80) CANNOT BE ENTERED FROM THE KEYBOARD

THE POLLOWING 2741 KEYBOARD CHARACTERS DO NOT HAVE A MEANING AS A CONTROL:

> \$ # # % & > - = . . : : ! | ? & SPACE EACKSPACE TAB

Fig. L.5: 2741 Keyboard Control Character Mappings (Continued)

TN73-50-I(L.6) 0 1 2 3 4 5 6 7 8 9 A B C D E P • y X 4 • • • • • • • # • • < • ( • + • ] • • • ٠ 5 6 7 8 • .a.b.c.d.e.f.g.h.i. . . . 9 • .j.k.l.m.n.o.p.q.r. . . A • • .s.t.u.7.w.x.y.z. . . В . . . . . . . . . . . . . . . С • • A.B.C.D.E.F.G.H.I. . . • • D • .J.K.L.M.N.O.P.Q.R. . . . . Ε • • • S.T.M.V.W.X.Y.Z. • • ٠ ٠ .0.1.2.3.4.5.6.7.8.9. . . F Hex Code X'xy' for Characters •• y 0 1 2 3 4 5 6 7 8 9 XX 06 07 · · · \$ · · · < · ( · + · | · 08 . . . . . . . . . . .1.8.\*.).;........ 09 10 . • 11 .>.?. . . . 12 • • •:• # • @ • \* • = • " • • a • 13 .b.c.d.e.f.g.h.i. . . 14 •••••j.k.l.m.n. 15 .o.p.q.r. . . . . . • . .s.t.u.v.w.x.y.z. 16 17 • • • • • • • • 18 . . . . . . . . 19 • • • • A.B.C.D.E.F.G. 20 • H• I• • • • • • • • J• .K.L.M.N.O.P.Q.R. . . 21 22 23 24 .0.1.2.3.4.5.6 7.8.9. 25 . . . . . . . Decimal Code D'xxy' for Characters X'05' = D'005' Horizontal Tab HT LC X'06' = D'006' Lower Case RES X'14' = D'02'' Print Restore NL X'15' = D'01 ' New Line ES  $X^{*}16^{*} = D^{*}02$ Back Space X\*17\* = D\*023\* IL Idle BYP X'24' = D'036' Frint Bypass X'25' = D'037' LP Line Feed UC X'36' = D'054' Upper Case Hex Code X'xy' and Decimal Code D'xxy' Control Codes

Fig. L.ó: 2741 Character and Control Codes

-	TN7	3_	50		(1		7)	]															
L	1.147		-	- 1				<u> </u>	-		_												
•y x	0	ĩ		2	3	4		5	6		7	ł	8	9		A	B	С	D	E	E	F	
0123456789kBCDEF	3. 0. 0. 3. 0. 0. 0.	• AJ/1 AJ/1 AJ/1 AJ/1 AJ/1 AJ/1 AJ/1 AJ/1		BK52BK52BK52BK52	CLTSCLTSCLTSCLIS	.D.M. .U.4. .D.M. .U.4. .D.M. .U.4. .D.M. .U.4.U	· · · · · · · · · · · X	ENVSENVSENVSENV5	FOW6FOW6FOU6	d	GPX7GTX7GPX7GPX7 e			IRZ9IRZ9IRZ9IRZ9 x	· · · · · · · · · · · · · · · · · · ·	: : : :	***************************************	くれるのくれらのくれたのくれるの				1 ~ ~ 1 ~ . 1 ~ . 1 ~ .	
		•• x x	у		0	1		2	3		4	ı	5	6		7	8	9					
		00120345 00120345 00111234 1111111 111222222 222222222222222			MIY2DES UV8BVO- 4-HK+HO:	A • N 7 2 3 * G J & V ? 9 C ( P / • 5 " I L ) X 1 #			C(P/,5"IL)X1#ELRT 7A.N [23-		D+QS%6 M:Y2@F& U>8BVO- 4=			TO UNBRO - 4=HK=WO:D+QS%6		GJ\$V?9C(P/.5"IL)X1#E RT 7		II)X1#Elix JA.N Z3+GJSV?9					

Decimal Code D'xxy'

Fig. L.7: Code for Characters on a PN Print Train

A. Walking and

de villes

A STATE OF THE STA

Τ	TN7	3-5	50-	I (1	8	)										
•y x	Û	1	2	3	4	5	6	7	8	9	A	В	С	D	E	F
0 1 2	•	• •	• •	•		•	•	•	•	•		•	•	• •	•	•
3 4 5	• •						•	•	•		e.	4 4	. < . . * .		, <b>+</b> .	•
6 7 8	• -	·/·			6		f		• •	•			1 3 <	· · · ·	> . = . +	?.
9 A B	- -	j 0	5.2	1 t	U.		C . W.	9. P. X.	q. y.	Г. Z.				)	±. ≥.	τ・ Β •
C D E	•	. A . . J .	B K S	C. L.	D H U	E. N. V.	F. 0. W.	G. P. X.	Н. Q Y.	I. R. Z.			•	• •		•

Hex Code X\*xy\*

• • - y	0123456789
XX	
00	
01	
02	
63	
04	
05	
06	
07	
08	
09	.1.5.*.).:/.
10	
11	.>.?
12	• • • • • • • • • • • • • • • • • • •
13	.b.c.d.e.f.q.h.if.
14	.<. (.+.+j.k.l.w.n.
15	.C.P.q.r
16	º.s.t.u.v.w.x.v.z.
17	· · · · · · · · · · · · · · · · · · ·
18	4.5.6.7.8.9
19	
20	•H•I••••••••J•
21	.K.L.M.N.O.P.C.R
22	••••••••••••••••••••••••••••••••••••••
23	.W.X.Y.Z
24	.0. 1.2.3.4.5.6.7.8.9.
25	

Decimal Code D'xxy'

There is a second state of a second

Fig. L.8: Code for Characters on a PN Print Train

## References

- BOLT, BERANEK & NEWMAN Inc. Report Number 1763, January 1969, AD-682 905.
- BOL1, BERANFK & NEWMAN Inc. Interface Message Processor, Specifications for the Interconnection of a HOST and an IMP. Report Number 1322, February 1971 Revision.
- BRYAN, P.P. IMP Interface Specification. Private Communication, December 1970.
- 4. CARR, C., CFJCKER, S. & CERF, V. HOST-HOST Communication Protocol II the ARPA Network. AFIPS Confelence Proceedings, 1970 Spring Joint Computer Conference, Atlancic City N.J., 569-597.
- COLE, G.D. Ferformance Measure ents on the ASPA Network. IFEE Transactions on Communications, Vol COM-20, No 3, June 1972, 630-636.
- CROCKER, S. L. HOST/HOST Protocol for the AHPA Network.
  Document Number 1, 1971 Fevision.
- CROCKER, S., HEAFNER, J., METCALFE, P. & POSTEL, J. Function Oriented Protocols for the ARPA Network. AFIPS Conference Proceedings, 1972 Spring Joint Computer Conference, 271-280.
- HEART, F., KAHN, R., ORNSTEIN, S., CPOWTHER, W. & WALDEN, I. The Interface Message Processor for the ARPA Computer Network. AFIPS Conference Proceedings, 1970 Spring Joint Computer Conference, Atlantic City N.J., 551-567.
- 9. 1DM CP-67/CMS User's Guide. Cambridge Scientific Centre, IBM Data Processing Division, 320-2015, July 1969 Revision.
- KIRSTEIN, P.T. On the Status of the ARPA Network, Summer 1972. INDPA Note 275, University of London Institute of Computer Science, November 1972.
- MCKENZIE, A. Froposed Change in IMF-to-HOST Protocol. Network Working Group, Request for Comment 312, NIC 9342, 22 March 1972.
- MCQUILLAN, J., CFOWTHER, W., COSELL, B., WALDEN, D. & HEART, F. Improvements in the Design and Performance of the ARPA Network. AFIPS Conference Proceedings, 1972 Fall Joint Computer Conference, 741-754.

- ORNSTEIN, S., HEART, F., CROWTHER, W., PISING, H., RUSSEL, S. & MICHEL, A. The Terminal IMP for the APPA Computer Network. AFIPS Conference Proceedings, 1972 Spring Joint Computer Conference, 243-254.
- 14. O'SULLIVAN, T.C. TELNET Protocol. Network Working Group, Request for Comment 158, NIC 6768, 19 May 1371.
- FOSTEL, J. A Proferred Official Connection Protocol. Network Working Group, Request for Comment 165, NIC 6779, 25 May 1971.
- 16. ROBERTS, L.G. & Forward Look. Journal of the Armed Forces Communications and Electronics Association, Vol XXV, No 12, August 1971, 77-81.
- FOFERTS, U.C. & WESSLER, B.D. Computer Network Development to Achieve Pesource Sharing. APJPS Conference Proceedings, 1970 Spring Joint Computer Conference, 543-549.
- SPOONER, C.P. A Software Architecture for the 70\*s: Part 1 - The General Approach. Foftware Practice and Experience, Vol 1, 1971, 5-37.
- WALDEN, D.C. A System for Interprocess Communication in a Resource Sharing Network. Communications of the ACN, Vol 15, No 4, April 1972, 221-230.
- 20. WHITE, J.E. An NCP for the ARPA Network, Computer Research Laboratory, UCSE, California, 21 December 7C.