

AD/A-002 700

INTERFACE MESSAGE PROCESSORS FOR THE
ARPA COMPUTER NETWORK

Frank E. Heart

Bolt Beranek and Newman, Incorporated

Prepared for:

Advanced Research Projects Agency

April 1974

DISTRIBUTED BY:

NTIS

National Technical Information Service
U. S. DEPARTMENT OF COMMERCE

UNCLASSIFIED

Security Classification

DOCUMENT CONTROL DATA - R & D

AD/A. 00 2700

(Security classification of title, body of abstract and indexing annotation must be entered when the overall report is classified)

1. ORIGINATING ACTIVITY (Corporate author) Bolt Beranek and Newman Inc. 50 Moulton Street Cambridge, Mass. 02138		2a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED	
		2b. GROUP	
3. REPORT TITLE QUARTERLY TECHNICAL REPORT NO. 5. INTERFACE MESSAGE PROCESSORS			
4. DESCRIPTIVE NOTES (Type of report and inclusive dates) 1 January 1974 to 31 March 1974			
5. AUTHOR(S) (First name, middle initial, last name) Bolt Beranek and Newman Inc.			
6. REPORT DATE April 1974		7a. TOTAL NO. OF PAGES 25	7b. NO. OF REFS
8a. CONTRACT OR GRANT NO. F08606-73-C-0027		9a. ORIGINATOR'S REPORT NUMBER(S) Report No. 2816	
b. PROJECT NO. 2351		9b. OTHER REPORT NO(S) (Any other numbers that may be assigned this report)	
c.			
d.			
10. DISTRIBUTION STATEMENT			
11. SUPPLEMENTARY NOTES		12. SPONSORING MILITARY ACTIVITY Advanced Research Projects Agency Arlington, Virginia 22209	
13. ABSTRACT The ARPA computer network provides a communication medium which allows dissimilar computers (Hosts) to interchange information. Each Host is connected to an Interface Message Processor (IMP), and IMPs are interconnected by leased common carrier circuits. There is frequently no direct circuit between two communicating Hosts, and the intermediate IMPs store and forward the information. IMPs regularly exchange information which is used to adapt routing to changing network conditions. IMPs also report a variety of parameters to a Network Control Center, which coordinates diagnosis and repair of malfunctions. The Terminal IMP (TIP) permits the direct attachment of 63 character-oriented terminals. The Satellite IMP (SIMP) will allow multi-station use of a single earth satellite channel. A High Speed Modular IMP (HSMIMP) is under development; one goal of this effort is to increase IMP performance by an order of magnitude. Specialized mini-Hosts under development will provide for: connection of remote batch terminals; simulation of a leased point-to-point circuit; encrypted Host communication.			

Reproduced by
NATIONAL TECHNICAL
INFORMATION SERVICE
US Department of Commerce
Springfield, VA 22151

DD FORM 1473 (PAGE 1)

S/N 0101-807-6811

UNCLASSIFIED

Security Classification

A-1140

UNCLASSIFIED

Security Classification

14 KEY WORDS	LINK A		LINK B		LINK C	
	ROLE	WT	ROLE	WT	ROLE	WT
Computers and Communication						
Store and Forward Communication						
ARPA Computer Network						
Interface Message Processor						
IMP						
Terminal IMP						
TIP						
Satellite IMP						
SIMP						
Honeywell DDP-516						
Honeywell H-316						
Multi-Line Controller						
MLC						
Network Control Center						
NCC						
Host Protocol						
High Speed Modular IMP						
HSMIMP						
Lockheed SUE						
RJE mini-Host						
Private Line Interface						
PLI						
Modem Substitute						
Pluribus						

DD FORM 1473 (BACK)

S/N 0101-807-6821

UNCLASSIFIED

Security Classification

A-31401

Report No. 2816

Bolt Beranek and Newman Inc.

INTERFACE MESSAGE PROCESSORS FOR
THE ARPA COMPUTER NETWORK

QUARTERLY TECHNICAL REPORT NO. 5
1 January 1974 to 31 March 1974

Submitted to:

IMP Program Manager
Range Measurements Lab.
Building 981
Patrick Air Force Base
Cocoa Beach, Florida 32925

This research was supported by the Advanced Research Projects
Agency of the Department of Defense and monitored by the Range
Measurements Laboratory under Contract No. F08606-73-C-0027.

TABLE OF CONTENTS

	page
1. OVERVIEW	1
1.1 The High Speed IMP	3
1.2 The Remote Job Entry Mini-Host	5
1.3 The Private Line Interface	6
1.4 The Modem Substitute	8
2. NETWORK RELIABILITY	10
2.1 Problems of Size	10
2.2 Hardware Improvements	13
3. HOST-LEVEL RELIABILITY IMPROVEMENTS	16

1. OVERVIEW

This Quarterly Technical Report, Number 5, describes aspects of our work on the ARPA Computer Network under Contract No. F08606-73-C-0027 during the first quarter of 1974. (Work performed from 1969 through 1972 under Contract No. DAHC-15-69-C-0179 has been reported in an earlier series of Quarterly Technical Reports, numbered 1-16.)

A major effort during this quarter has been to improve the actual and apparent reliability of the network. Sections 2 and 3 of the Quarterly Technical Report describe our work in this area.

During this quarter, we delivered one TIP to Kirtland Air Force Base (New Mexico). In addition, the spare IMP/TIP was reconstituted; it is now housed in three "low-boy" cabinets and resides on shipping pallets for flexible and easy shipment. The first low-boy contains a basic 316 IMP with two modems, one Host, and one Very Distant Host interface, in addition to the mainframe and 32 kilowords of memory. The second cabinet contains a Multi-line Controller and 32 Line Interface Units. The third cabinet contains a third modem interface. This configuration will give us flexibility in case of serious IMP failures.

Development of the Satellite IMP program has continued during this quarter. Although it still lacks some of the features expected in the operational system (e.g., packet tracing and statistics), the current version of the program realizes a Slotted ALOHA protocol for store and forward traffic. In order to guarantee the propagation of routing and to decrease contention for the channel, two slots out of 64 are reserved for the trans-

mission of routing and I-Heard-You packets by each IMP. This routing is now broadcast to all the nodes at once rather than pairwise by IMP. This version also has acknowledgment by slot rather than by the usual IMP-to-IMP mechanism. It also has a facility for reloading memory in pieces 61 words long rather than in one 24,000 word burst.

Our interactions with COMSAT during this quarter have concentrated on the political problems of locating a Satellite IMP in a COMSAT ground station. Although this problem is not yet solved, some progress is being made.

Last quarter, we modified the NCC program to permit privileged users on BBN TENEX or the PDP-1 to examine the contents of memory locations. During this quarter, several programs were developed at BBN TENEX which use this facility to maintain files on network throughput and node and line availability. We intend to use this capability to generate monthly summaries using BBN TENEX rather than the PDP-1.

Our participation in the network Users Interest Group (USING) during this quarter included attendance at a meeting at SRI in January. Members of the IMP group are participating in USING subcommittees to define the User community and its needs, to develop an information management system for the network, and to design a common command language; we are also contributing to the reports on feedback mechanisms, to service center definition, and to the specification of NETED, a simple network-wide line editor.

We presented three papers this quarter at the Seventh International Conference on System Sciences at the University of Hawaii

in January. These were: "Design Considerations for Routing Algorithms in Computer Networks", "The Satellite IMP for the ARPA Network", and "The BBN Multiprocessor". The material in the latter two papers was reported in our Quarterly Technical Report No. 4. In addition, we revised BBN Report No. 1822, "Specifications for the Interconnection of a Host and an IMP", and distributed the revisions to the network.

1.1 The High Speed IMP

During this quarter we have had some difficulties and some successes in the High Speed IMP project. We have continued to struggle with Lockheed design problems, particularly in the memories. These have compounded our own problems with the bus coupler design. Nonetheless, we have recently frozen the bus coupler design and are starting to manufacture them for the production machines. The prototype High Speed IMP has run for considerable periods error-free with all seven processor busses, both memory busses, and a single I/O bus. We recently replaced all the private processor memories with updated versions, and in the same period added the second I/O bus and more sophisticated test programs to run I/O devices to and from memory. We are in the process of debugging the resultant problems. We have also added power protection features to the bus couplers. When power is going to go down on a memory or I/O bus, a special interrupt is sent to all processor busses warning them. They have a short period (2-1/2 ms) to extract valuable parameters, code, etc. from memory before it dies utterly. We have also gone to considerable lengths to try to insure that as power on a bus dies, the bus does not send extraneous harmful signals into the rest of the system through its bus couplers. In addition, we have arranged modifications to the bus controllers that (under switch control)

(a) provide for a "hung bus" watchdog timer that resets the bus in the face of total inactivity for one second, and (b) re-enable the 60-cycle interrupt in the event of a bus reset (without this change, the interrupt is disabled) to allow for automatic restart after power up.

The operational program now has its full complement of fake Hosts. This includes a console-terminal handling package which is available both as a Host that communicates with the network and as an output device for the display program. The display program monitors assigned locations in memory and maintains an updated display of them on the console terminal (Infoton Vistar) screen. Another fake Host now complete is a debugging package, DDT, that is usable both in stand-alone operation and as a network Host for examining and modifying locations and controlling (in a limited sense) the machine. It was necessary to develop means for addressing both common (20-bit address) memory and individual local memories, and to address and control individual processors.

The IMP program now employs a dynamic initialization technique that looks for all possible devices, and if and when found, initializes them. This allows dynamic reconfigurations such as adding and deleting Hosts, modems, and other I/O devices.

All modems, Hosts, and some internal processes have associated with them a block of parameters fully describing their context. These blocks are now monitored to detect illegal states, as well as legal states that are "hung". The code itself is checksummed (in both local and common memories) to detect smashed code.

The system variables are periodically examined for self-consistency. In the event of a large failure (e.g., core overwritten), these variables will be reconstructed. This particular code is being tested by randomly zeroing memory locations with DDT.

At this point many of the failure detection mechanisms halt the machine when they find an error. New code is now being debugged to recover from these errors. Part of the recovery mechanism uses the 60-cycle interrupt which is issued to every processor. Each processor, through a communications region in common memory, looks to see if all the other processors are also running. Bad processors are thus detected when they fail to participate in the community action.

Finally, the name "Pluribus" was chosen for the line of machines used to implement the High Speed IMP, which itself is now called the Pluribus IMP.


1.2 The Remote Job Entry Mini-Host

The Remote Job Entry (RJE) mini-Host system has continued to progress during this quarter. The hardware has been thoroughly checked out and has been performing well throughout the quarter. The current configuration consists of a single-processor Pluribus system with 16K core memory with several BBN-designed cards, including a standard Host-to-IMP interface and a Synchronous Line Interface (SLI) for the synchronous modems typically used by IBM model 2780 RJE terminals. After some initial difficulties, the hardware now runs successfully both with test programs and with the RJE mini-Host software. The major part of the software has been written and largely debugged. The system now includes an initial implementation of the Host-to-IMP and the Host-to-Host

network protocols as well as the control structure and IBM 2780 binary synchronous protocol modules mentioned in our Quarterly Technical Report No. 3. The software provides for flexibility in a data structure which can be tailored to various RJE equipment configurations, and in a buffering scheme that can respond dynamically to the actual use of the system. The RJE mini-Host has been successfully connected to the network via the standard Host-to-IMP interface (BBN Report No. 1822); by the end of the quarter, the system was able to communicate across the network back to itself, to open connections and pass data between itself and a TIP terminal, and to establish connections to special sockets such as ECHO and Date-Time at BBN TENEX via the Initial Connection Protocol. On the device side, the RJE mini-Host has been able to accept data input from the IBM 2780 card reader, and to output data selectively to the printer or card punch.

1.3 The Private Line Interface

During this quarter our effort with the Private Line Interface (PLI) has been directed primarily at the problem of modifying its specification to make it suitable for installation in a secure environment. There have been two main problem areas: (1) whether or not the Key Generator (KG) had to be in series between the secure Host and the network, thus providing complete Red to Black (unencrypted to encrypted data) separation; and (2) the extent to which TEMPEST (leakage of secure data by radiation) considerations had to be handled within the PLI rather than counting on a shielded environment. To obtain closure on these issues we have met extensively with representatives of NSA and ARPA.



All cables and cabinets fully shielded

Figure 1 PLI Configuration.

We now believe we have resolved the above two issues. In both cases the most conservative approach is to be taken; namely, the KG is to be in series providing complete Red to Black separation (although retaining a 16-bit Black to Red data path over which necessary control messages can be sent), and the PLI is to be self-contained with regard to TEMPEST considerations. Figure 1 illustrates the conservative design now being adopted.

We hope to construct the PLI in such a manner that for non-KG applications, the Red and Black portions of the PLI collapse naturally into one unit. This possibility is facilitated by the use of the modular Pluribus hardware and by basing the software design on many of the same multi-programming techniques we have developed while constructing the Pluribus IMP.

Despite the modification of the PLI specification, we had already proceeded with the procurement of hardware as originally specified and approved by ARPA. The hardware we now have has been undergoing extensive reliability testing, and several existing diagnostic programs have been upgraded to make them more suitable for testing. Several specific types of failures have been detected and their causes are being determined and corrected.

1.4 The Modem Substitute

During this quarter we completed development of a new method for connecting terminals to the TIP. This device, called the Modem Substitute, is designed to provide data communication in the gap between Very Distant Terminals which connect directly to the TIP using 25-conductor cable. Both the TIP Line Interface Unit (LIU) and commercially available terminals follow the conventions of EIA specification RS-232. These high-impedance,

single-ended, large-magnitude signals do not lend themselves to reliable transmission over distances greater than about 25 meters. For connection to terminals within the same building or campus complex, modems have provided the only alternative. In addition to high cost, conventional modems have the unfortunate property of limiting asynchronous data speeds to 300 Bps in most applications, or 1200 to 1800 Bps under special circumstances.

In situations where multi-pair cables can be installed or leased, the Modem Substitute allows reliable transmission speeds of at least 9600 Bps over distances estimated at one kilometer or better. The RS-232 voltages are changed into low-impedance, balanced signals by the box located at the user's terminal, and a new breed of LIU card in the TIP accepts these signals. Only two twisted pairs are needed, although we recommend three in order to provide the equivalent of a Carrier Detect Signal to the LIU.

Breadboards have been constructed and tested, and the documentation has been completed. We stand ready to produce both the LIU cards and the free-standing units when there is further expression of interest. The complete configuration for one terminal will cost about \$330, which is considerably less than the cost of modems, DAA's, and telephone lines, the potential speed increase being an added bonus.

2. NETWORK RELIABILITY

We have always been concerned that the communications service provided by the network should be as reliable as possible within the constraints of the research and development character of the project. In the past this concern has been manifested in efforts such as the round-the-clock staffing of the Network Control Center, modifications to the IMP software to report an increasing number of "exception" conditions, and the great attention paid to reliable and "fail-soft" operation in the design of the High Speed Modular IMP. We have also engaged in a continuous review of IMP failures, with a view toward uncovering design deficiencies in both the Honeywell hardware and the IMP software. During the past six months, and particularly during the first quarter of 1974, we have further increased our efforts to improve reliability as reported below.

2.1 Problems of Size

The network now includes 45 operational IMPs and TIPs plus, usually, the BBN TESTIP. The size of the network tends to exaggerate the effects of low-probability failures. For example, a hardware design problem which would result in a failure on a single machine once a month will probably occur somewhere in the network once every 16 hours. Even an event which has a once-a-year probability in a single machine has almost a once-a-week probability in the network. Similarly, if new IMP software has an obscure bug, the probability of this bug showing up in the network in a single day is twice the probability of the bug showing up in the BBN test cell operating in three machines for a week, even ignoring those bugs which are purely related to the size and complexity of the network (dynamic routing, for example, might contain such bugs).

In addition, the number of IMP variations continues to increase: there are 516's and 316's; IMPs and TIPs; Local, Distant, and Very Distant Hosts; circuits running at 7.2, 9.6, 50, and 230 kilobits per second; surface and satellite circuits. Each of these "variables" tends to apply to a machine essentially independently of the values of the other "variables", yet there is some probability that a given combination will interact on the hardware or software in some unexpected way. (For example, a problem first *observed* at the Ames TIP was due to the interactions of four circuits, one a satellite circuit and two others running at 230 Kbs, with the software in the Ames IMP.)

Finally, site problems (which have always occurred) are magnified in their effect as the average network connectivity has gradually decreased. For example, we are averaging between one and two site power failures per day. There are now several network areas where there are 5 or 6 IMPs strung out in a row; almost any pair of "simultaneous" failures in such a string will isolate some machines between the failing sites. If the failures are in the IMP software it can quickly be re-initialized; if they are in the hardware we can frequently bypass the machine; however, if they are in the site environment (e.g. power) we are not usually able to influence the duration of the isolation. Occasionally the environmental problems are rather bizarre; during the month of February one machine was damaged by an internal fire (ISI), power at a second was lost due to computer-room flooding (Rutgers), and a third machine was turned off when the room was filled with dust from building construction (Moffett).

In an effort to deal with some of these problems we have made a number of changes to the operation of the Network Control Center, and a number of recommendations to ARPA for network changes (for example, additional circuits, limited backup power at some or all sites). Early in the quarter we added a full-time engineer to the NCC staff to assume responsibility for field hardware maintenance. The added manpower has enabled us to speed the diagnosis of difficult "crisis-type" problems, deal more effectively with the Honeywell field maintenance offices, and engage in additional long-term analysis of the causes of problems and the effectiveness of our response. This in turn has led to a noticeable decrease in the mean time to repair problems, and hence an increase in average IMP "availability."

A second major change has been a formalization of our checkout procedures prior to IMP software releases or delivery of new machines. This is primarily because of the wide variety of possible IMP configurations, as mentioned above, and the necessity for checking as many as possible due to the potential for deleterious interactions among various components in a single system. There are now enough tests to be performed that we desired to have the operators, rather than the system programmers, perform the testing; this led to the necessity for some formalization of the testing procedures. New systems, hardware or software, now undergo about one week of this testing once they are deemed to be "finished".

During the first quarter there were two IMP software releases; both dealt primarily with making the software more resilient in the face of hardware problems rather than with the many new features which are planned for the network. These changes include the addition of a "Master Clear" routine to the program loading

process, a change in the circuit "alive/dead" logic for circuits to singly-connected sites (to make the IMPs more tolerant of circuit problems), software checks on packet length to deal with intermittent I/O hardware problems, addition of an IMP core dump facility, and periodic software checks of the registers used for the priority interrupt mask. In addition, a lockup condition theorized by the Network Measurement Center is in the process of being fixed; the cause of an actual lockup, which occurred in December of 1973, was discovered and fixed.

2.2 Hardware Improvements

Since 1969 our engineers have been engaged in continuous interaction with Honeywell, reporting on "systemic" problems in the design of the 516's and 316's, and keeping track of Honeywell engineering changes to insure that the machines used in the Network incorporate all desirable changes. During late 1973 and most of the first quarter of 1974 we carried out a complete field survey of the IMPs and TIPs. During this survey several engineering changes were installed and a number of wiring changes were made. In addition, a BBN-designed option was installed as part of a long-range plan to have each machine specify its complement of options in a set of hardware registers accessible to the software. Finally, our new maintenance manager visited both the Los Angeles and San Francisco Honeywell Field Service Offices to discuss the operation of the network and his own new role within the Network Control Center in an effort to insure the best possible service in these two areas where there are heavy concentrations of IMPs and TIPs.

During September of 1973 we had a serious failure in the TIP installed at the Department of Commerce, Boulder (Colorado).

After repeated Honeywell efforts to repair the machine in the field proved unsuccessful, we decided in October to send the spare TI² to this site. Honeywell continued to work on the original machine at the site, and continued to be unsuccessful. By early November the original machine had been returned to BBN and by December it was returned to the Honeywell factory. This machine was replaced by Honeywell in late January of this year, once again giving the NCC a "spare" capability by mid-February. After finally getting the original DOCB machine into operating condition, Honeywell used it as a "test vehicle" to study complaints we had made several times about the 316 power supplies.

Although we have yet to (and may never) receive a formal report from Honeywell on the results of this testing, we have learned informally that the results tend to bear out our complaints. Basically, it was verified that under some conditions of failure or noise spikes on the primary power, the power-fail interrupt occurs but the auto-restart does not activate the program when primary power recovers. In addition, the power supplies in option drawers of the 316 may power down before the mainframe power supply. This may make it impossible for the power-fail auto-restart mechanism to operate correctly; in fact the mainframe supply may never notice the failure and thus make no attempt to restart. The classic method of dealing with this type of problem is to "desensitize" the option drawer supplies (a much more desirable method is to "coordinate" the supplies, but supplies designed to permit this are more expensive and therefore not used by Honeywell in the 316), but we understand that the Honeywell test engineers were unable to "desensitize" the supplies sufficiently to guarantee correct operation.

It is our understanding that Honeywell is now investigating the results of these tests, plus the inputs from BBN engineers which inspired the factory testing effort, in an attempt to devise appropriate modifications to the power-fail mechanism. We will continue our interaction to insure that these problems are being dealt with; we have already learned of two engineering changes which may solve most of these problems.

3. HOST-LEVEL RELIABILITY IMPROVEMENTS

During the past quarter we responded to ARPA's admonition that we take increased responsibility for the reliability of a user's connection from the terminal he is using to the Host he is using.

Much of this effort was carried out with close cooperation between the TENEX development group and the TIP development group at BBN. Improvement in the reliability of users' connections between TENEXs and TIPs would have major impact on the appearance of overall network reliability due to the large number and high visibility of TENEXs and TIPs. Despite the emphasis on TIP/TENEX interaction, however, all work done applies equally well to interactions between Hosts of any type.

The remainder of this section gives a sketch of our plan for improving the reliability of connections between TIPs and TENEXs. Major portions of this plan have already been implemented and by the end of this quarter were undergoing final test prior to release throughout the network. Completion of the implementation of the plan is expected in the next quarter. A working document describing the plan in greater detail has been circulated to interested parties throughout the network.

Our plan for improving the reliability of TIP/TENEX connections is concerned with obtaining and maintaining TIP/TENEX connections, gracefully recovering from lost connections, and providing clear messages to the user whenever the state of his connection changes.

When a TIP user attempts to open a connection to TENEX, the TENEX may be down. In this case the user must be provided with helpful information about the extent of the TENEX's unavailability. To facilitate this, we modified the IMP program in this quarter to accept and utilize information from a Host about when the Host will be back up and for what reason it is down. TENEX is to be modified to supply such information before it goes down or, through manual means, after it has gone down. When the TIP user then attempts to connect to the down TENEX, the IMP local to the TENEX returns the information about why and for how long TENEX will be down. The TIP is to be modified to report this information to the user; e.g., "Host unavailable because of hardware maintenance -- expected available Tuesday at 16:30 GMT".

The TIP's logger is presently not reentrant. Thus, no single TIP user can be allowed to tie up the logger for too long at a time; and the TIP therefore enforces a timeout of arbitrary length (about 60 seconds) on logger use. However, a heavily loaded TENEX cannot be guaranteed always to respond within 60 seconds to a TIP login request, and at present TIP users sometimes cannot get connected to a heavily loaded TENEX. To correct this problem, the TIP logger will be made reentrant and the timeout on logger use will be eliminated.

Once a connection is made between a TIP and a TENEX, maintenance of a reliable connection becomes the paramount concern. For instance, the subnetwork may occasionally lose a message being transmitted over the connection. Both TENEX and the TIP have been modified to retransmit such lost messages, loss of which can easily be detected in most cases.

One notorious soft spot in the Host/Host protocol which harms the reliability of connections is the Host/Host protocol incremental allocate mechanism. Low frequency software bugs, intermittent hardware bugs, etc., can lead to the incremental allocates associated with a connection getting out of synchronization. When this happens it usually appears to the user as if the connection just "hung up". A slight addition to the Host/Host protocol to allow connection allocates to be resynchronized has been designed and implemented for both the TIP and TENEX.

TENEX has a number of internal consistency checks (called "bughalts") which occasionally cause TENEX to halt. Frequently, after diagnosis by system personnel, TENEX can be made to proceed without loss from the viewpoint of local users. A mechanism is being provided which allows TENEX to proceed in this case from the point of view of TIP users of TENEX.

The appropriate mechanism entails the following: TENEX will not drop its ready line during a bughalt (from which TENEX can usually proceed successfully), nor will it clear its NCP tables and abort all connections. Instead, after a bughalt TENEX will: discard the message it is currently receiving, as the IMP has returned an Incomplete Transmission to the source for this message; reinitialize the interface to the IMP; and resynchronize, on all connections possible, Host/Host protocol allocate inconsistencies due to lost messages, RFNMs, etc. The latter is done with the same mechanism described above. This pro-

cedure is not guaranteed to save all data -- a tiny bit may be lost -- but this is of secondary importance to maintaining the connection over the TENEX bughalt.

The TIP user must be kept fully informed as TENEX halts and then continues. Therefore, the TIP has been modified to report "Host not responding -- connection suspended" when it senses that TENEX has halted (it does this by properly interpreting messages returned by the destination IMP). When TENEX resumes service after proceeding from a bughalt, the above procedure notifies the TIP that service is restored, and the TIP has been modified to report "Service resumed" to all users of that Host.

On the other hand, the service interruption may not be proceedable and TENEX may have to do a total system reload and restart. In this case TENEX will clear its NCP connection tables and send a Host/Host protocol reset command to all other Hosts. On receiving this reset command, the TIP will report "Host reset -- connection closed" to all users of that Host with suspended connections. The TIP user can then re-login to the TENEX or to some other Host.

Of course, the user may not have the patience to wait for service to resume after a TENEX bughalt. Instead, he may unilaterally close the connection and connect to some other Host. If TENEX is then able to proceed, its NCP will still think its connection to the TIP is good and suitable for use. Thus, we have a connection which the TIP thinks is closed and TENEX thinks is open, a phenomenon known as the "half-closed connection". A procedure for cleanly completing the closing of such a connection has been specified and implemented for the TIP and TENEX.

Since TENEX will maintain connections across service interruptions, the TIP user will be required to take the security procedure of closing his connection ("@C" command to TIP) before abandoning his terminal. This command will guarantee that his connection will not be reestablished on resumption of service. Otherwise, his job would be left at the mercy of anyone who acquires that terminal.