

Report No. 1966

April 1970

INTERFACE MESSAGE PROCESSORS FOR
THE ARPA COMPUTER NETWORK

QUARTERLY TECHNICAL REPORT NO. 5

1 January 1970 to 31 March 1970

SUBMITTED TO:

Advanced Research Projects Agency
Washington, D.C. 20301
Attn: Dr. L.G. Roberts

This research was supported by the Advanced Research Projects Agency of the Department of Defense under Contract No. DAHC-15-69-C-0179.

TABLE OF CONTENTS

	page
SECTION 1. INTRODUCTION	1
2. SOFTWARE DEVELOPMENT	2
3. HARDWARE DEVELOPMENT	4
A. Retrofits.	4
B. Modem Simulator.	5
4. NETWORK TESTING.	7
5. PHONE LINE TEST PROGRAM.10

1. INTRODUCTION

This report describes several aspects of our progress on the ARPA computer network during the first quarter of 1970:

- Several minor changes and additions have been made to the operational program. In addition we have streamlined the program assembly process. Software activity during the quarter is described in Section 2.

- We have designed a modem simulator that allows six full duplex modem channels to be interconnected in various configurations for testing IMP topologies in the BBN test cell. The design of this simulator and the results of other hardware activity are described in Section 3.

- During the first quarter, we conducted the first of our planned network tests in the field. A description of the test procedure and some of the preliminary results are discussed in Section 4.

- A second phone line test program has been coded to allow the buffer size to be selected in advance. This program is described in Section 5.

2. SOFTWARE DEVELOPMENT

A new version of the operational program was released on March 1. This program incorporates the cease-on-link mechanism that resulted from the Utah meeting and subsequent discussions with the Hosts. We explained this mechanism in the previous Quarterly Technical Report (No. 4). [A Host may send a cease-on-link message to its IMP at any time. Its IMP will then specially mark the next RFNM it returns on that link and deliver that RFNM as a cease-on-link message to the other Host. The link will then be unblocked.] It is for the Hosts to agree to heed this convention and regulate the flow of traffic.

The program has been modified to allow a maximum of five modems and a maximum of three Hosts in its standard configuration; should a further increase be required, the program can be modified to handle four Hosts.

~~During the last quarter, we had our first opportunity to make the multiple Host logic work.~~ In addition to requiring minor changes to the main software, a number of changes were made in the statistics programs to reduce the required amount of table space. One such change was to consolidate the statistics of all the Hosts into a single set of Host tables.

The status reporting feature was modified to include two Hosts and five modems. In addition, the status of the sense switches, measurement flags, memory protect switch and the number of free buffers are now reported. We implemented the sense switch 2 feature that enables the use of the parameter change routine.

A major change was made in the method of assembling the operational program. This change has considerably speeded up the assembly process and permits patch free system releases. Previously, we edited the program on the PDP-1, assembled on the DDP-516, and listed the program using the XDS-940 line printer. With this process, about three working days (or one day/night session) were typically required for each assembly. We have now modified an existing assembler on the PDP-1 to assemble 516 code, thus reducing the total assembly time to several hours and in addition, allowing a new system tape to be prepared via remote teletype.

3. HARDWARE DEVELOPMENT

During the last quarter, we received delivery from Honeywell of IMPs No. 5, 6, and 7; these were installed and tested in the BBN test cell. Testing revealed a number of minor problems and, while these problems have, for the most part, been easily corrected, considerable effort has been expended in locating them and in working with Honeywell to have them corrected. As of this date, machines No. 5, 6, and 7 are operating correctly in the test cell.

In late March, AT&T succeeded in providing a 50-kilobit circuit across the country. IMP No. 5 (BBN) was then connected into the network via a temporary tie to the IMP at UCLA. The standard IMP-TST program was run between these two machines and shortly afterwards the operational IMP programs were in communication over this line. In particular, we have encountered no problem in communicating over this line and are able to communicate with each of the other four sites in the network using the IMP teletype. Our preliminary experience indicates that there is no delay noticeable to a person at a teletype in echoing single character messages back and forth across the country.

A. Retrofits

The original four IMPs were delivered with certain known deficiencies, which will be remedied by field retrofits during the next quarter. These include the following:

- 1) The original design of the auto restart, including time-delay relays, had inadequate reliability. An altogether new

design delivered with machine No. 7 has been tested and appears to work satisfactorily. This design will be retrofitted to all machines.

2) Machine No. 5 was the first IMP delivered with the ruggedized control panel. All successive machines have been delivered with the new panel, which appears to be working satisfactorily. Ruggedized panels will be retrofitted to the four IMPs presently in the field.

3) A problem existed with the status lights, causing lights which should have been off to be partially lit. This problem arose from operating the light driver transistors beyond their maximum voltage ratings. A new power supply for these drivers was installed and demonstrated in machine No. 7. This cured the difficulty, and machines No. 5 and 6 have been retrofitted with the new supply. Field machines will also be retrofitted.

B. Modem Simulator

The 50-kilobit room circuit (2 modems) in our test cell provides only one full duplex IMP-to-IMP connection to be made in the standard way. To test more complex configurations (involving more than two IMPs), modem interfaces were connected directly via special cables. This method was awkward and required some temporary rewiring of the machines. To avoid this awkwardness and to facilitate a wider variety of situations for testing the operational program, a "modem simulator" box (that simulates six full-duplex communication circuits) is being designed and constructed. The IMPs will connect directly into this box via their standard modem cables. The box will

provide the interconnection between pairs of send and receive signals and will also provide a common clock signal simulating the modem send and receive clocks required by IMP modem interfaces.

4. NETWORK TESTING

During the last quarter, we conducted the first of several planned network tests on the initial four-node net connecting UCLA, SRI, UCSB, and Utah. The UCLA computer center and the facilities developed by UCLA for processing network measurement information were made available to us for an extended period of time. In addition, UCLA prepared a stand-alone test program for the XDS Sigma-7 computer to help us in our testing. In this and subsequent reports, we will be describing our testing activities, along with any test results of interest. We do not, however, intend these tests to be merely systematic measurements of net capability. Rather, the tests will be used to determine the limits of performance and capabilities under extreme loading conditions, real and simulated. From our initial testing we obtained a clearer picture of the network performance and, at high traffic loads, located a number of program bugs that were easily corrected.

The network can be described as a taut communication system in the sense that each message entering the net proceeds directly to its destination. We observed that messages do not wander about the net, even at high traffic loads. The net easily handles sequences of single-character messages typed at the IMP teletype. In all our testing, the Host was prompt in handling messages and we observed no backup of traffic. We explored the variations in buffer occupancy when the upper limits for reassembly buffers and store and forward buffers were varied. We observed that two conditions appear to cause a large number of IMP buffers to be occupied when the traffic is heavy. One

~~Destination IMP's~~
~~buffers~~
~~and~~
~~Source IMP's~~
~~buffers~~
~~and~~

condition is an insufficient supply of reassembly buffers at the destination IMP. A second condition is the use of multiple links by a Host (for fanning messages), which causes a large part (or all) of the source IMP's store and forward buffers to be filled.

looking
 IMPs
 only

There was no observed buildup of store and forward traffic at intermediate IMPs, except when the number of reassembly buffers at the destination was insufficient to handle the incoming traffic. From our tests, it appears as if a limit of 8 or 16 reassembly buffers is not sufficient to handle heavy traffic to the Host. A limit of 32 reassembly buffers (which is the current maximum) does appear to be sufficient. In our initial tests, the variation of the store and forward limit appeared primarily to affect the maximum rate at which traffic could flow from the Host. The limit on the number of store and forward buffers is currently set to 21.

The most important change made to the IMP program as a result of network testing was the removal of a phase-lock condition between message transmission time and the timeout period. Although the presence of this phenomenon was known about in advance of the testing, the extent to which it affected the round-trip transit time had not been expected. In the original version of the program, a periodically-run timeout program was relied upon to start up, if possible, each process that had not been able to continue running earlier, for example, while waiting for an event to occur. Although this mechanism had been introduced primarily as a safety feature, it had the unplanned effect of causing the round-trip times to be multiples of the timeout period for the shorter single-packet messages. This effect is no longer present in

the new system. The round-trip time now depends only upon message length, the route, the phone line bit rates, and the traffic load.

A bug was discovered and fixed in the multiple Host logic in the Host routine. A multiplication table used in the routing algorithm had not fully assembled. This produced a brief flurry of random routing whenever the quality on the phone lines decreased momentarily. The required table was patched into the program.

In an alternate routing experiment (performed with the routing bug still present), 8000-bit messages were sent from UCLA to SRI on 60 links at the RFSM limited rate. The traffic on the Host line to the IMP was measured to be approximately 80 kilobits/sec for this case, thus making essentially complete use of both circuits from UCLA to SRI. The Host line was tested at about 300 kilobits/sec with different Host message lengths and no timing problems were noticed.

5. PHONE LINE TEST PROGRAM

Our first phone line test program was designed as a simple way to obtain raw data on the occurrence of packets in error on the phone line using short 88-bit packets. This data was printed on the IMP teletype (in octal) and also stored in a histogram. However, certain properties of the recorded data, such as the cumulative packet error rate, were quite inconvenient to compute from the recorded data and were only poorly approximated from the histogram. Furthermore, we had no convenient way to obtain the characteristics of the phone line when larger packets were used.

During the last quarter, a second phone line test program was written. This program permits the selection of packet size, over the range from 88 bits to 9288 bits. The program first counts the number of consecutive good packets (i.e., packets that contain no errors) and then counts the number of consecutive packets in error and so forth. However, the raw data is not recorded. A histogram is kept of runs without error and also of error bursts, both of which are typed out (but not cleared) every ten minutes. In addition, the program computes the ratio of total good packets to bad packets for each ten-minute interval, along with the ratio of cumulative good to cumulative bad since the beginning of the test.

The new program was designed to measure as many as three phone lines simultaneously. The phone line may be looped back, thus giving a two-way test of the line; or another copy of the test program may be loaded into a neighboring IMP to obtain two one-way line tests.

The line test program also measures the propagation delay time to within 0.1 msec accuracy. As part of its initialization procedure, the line test program transmits a message to its neighbors, in which it includes the current time as well as other pieces of information such as the buffer size. The neighboring IMPs, if any, must all have the line test program loaded *and in a waiting state*. When the neighboring IMP receives the initialization message, its line test program starts and returns the received initialization time. When this message returns, the delay is easily calculated.

In addition, the line test program may be stopped whenever a phone line error occurs and, using DDT, the exact location of the bit errors may be determined. This technique is useful with moderate- to long-size packets to help in understanding the structure of error patterns. The program actually checks the data and thus allows the performance of the error check register to be monitored. A separate recording is made if the data is ever incorrect and not detected by the check register. However, this event has never been observed to happen.

DOCUMENT CONTROL DATA - R & D

(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, Massachusetts 02138		2a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED	
		2b. GROUP	
3. REPORT TITLE INTERFACE MESSAGE PROCESSORS FOR THE ARPA COMPUTER NETWORK			
4. DESCRIPTIVE NOTES (Type of report and, inclusive dates) QUARTERLY TECHNICAL REPORT NO. 5, 1 January 1970 to 31 March 1970			
5. AUTHOR(S) (First name, middle initial, last name)			
6. REPORT DATE April 1970		7a. TOTAL NO. OF PAGES 14	7b. NO. OF REFS
8a. CONTRACT OR GRANT NO. DAHC15-69-C-0179		9a. ORIGINATOR'S REPORT NUMBER(S) Advanced Research Projects Agency Washington, D.C. 20301	
b. PROJECT NO.		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	
13. ABSTRACT The basic function of the IMP computer network is to allow large existing time-shared (Host) computers with different system configurations to communicate with each other. Each IMP (Interface Message Processor) computer accepts messages for its Host from other Host computers and transmits messages from its Host to other Hosts. Since there will not always be a direct link between two Hosts that wish to communicate, individual IMPs will, from time to time, perform the function of transferring a message between Hosts that are not directly connected. This then leads to the two basic IMP configurations — interfacing between Host computers and acting as a message switcher in the IMP network. The message switching is performed as a store and forward operation. Each IMP adapts its message routine to the condition of those portions of the IMP network to which it is connected. IMPs regularly measure network performance and report in special messages to the network measurement center. Provision of a tracing capability permits the net operation to be studied comprehensively. An automatic trouble reporting capability detects a variety of network difficulties and reports them to an interested Host. An IMP can throw away packets that it has received but not yet acknowledged, transmitting packets to other IMPs at its own discretion. Self-contained network operation is designed to protect and deliver messages from the source Host to the destination Host.			