

Report No. 2003

July 1970

TABLE OF CONTENTS

INTERFACE MESSAGE PROCESSORS FOR  
THE ARPA COMPUTER NETWORK

QUARTERLY TECHNICAL REPORT NO. 6  
1 April 1970 to 30 June 1970

Submitted to:

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

*Impersonal  
collapse  
abruptly!*

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

Section 1. INTRODUCTION.....1  
2. SOFTWARE DEVELOPMENT.....3  
3. HARDWARE.....5  
4. NETWORK TESTING.....7  
5. NETWORK CONTROL CENTER.....11

## 1. INTRODUCTION

This Quarterly Technical Report No. 6 describes several aspects of our progress on the ARPA Computer Network during the second quarter of 1970. During this period, we have installed IMPs Nos. 6, 7, 8, and 9, at MIT, Rand, SDC, and Harvard, respectively. AT&T has not yet provided circuits to Harvard, but the eight-node subnet is operational. The temporary circuit connecting BBN and UCLA was taken down and replaced by a circuit between BBN and Rand. A circuit that is listed as temporary was installed between MIT and Utah.

We issued a fourth revision of the operational IMP program (IMPSYS 23) and called back the March 1 system. In addition, we constructed a second simulation program to study the effects on throughput of routing and anti-congestion algorithms. Our software activity is described in Section 2.

Our primary hardware effort, described in Section 3, was devoted to testing, debugging, and field installation of IMPs. A slight modification was made to the terminating network in the distant Host driver. In addition, a potential hazard in the standard Host/IMP interface was located and corrected.

We conducted the second in a series of network tests designed to study the performance of the subnet under extreme loading conditions. Our efforts were considerably aided by the helpful

cooperation of several Hosts, including UCLA, Rand and SRI. Various aspects of this testing are described in Section 4.

A Network Control Center has been established and is operational at BBN. Each IMP in the net reports to the Network Control Center every fifteen minutes or whenever important status changes occur. Our experience in operating the control center is described in Section 5.

During this quarter, a revised version of the Host Specification (BBN Report No. 1822) and the initial IMP Operating Manual (BBN Report No. 1877) were revised. As normal updating occurred in the program and hardware, revised pages to the Operating Manual were issued to the Hosts.

## 2. SOFTWARE DEVELOPMENT

A fourth version of the operational program (IMPSYS 23) was released during the last quarter. We have gathered considerable experience with this version, particularly during a period of intensive testing on the west coast, and are preparing a new and improved model for the next quarter. The areas of routing and congestion are particularly troublesome, largely because there has not been enough natural traffic in the net to exercise our routing and anti-congestion algorithms. We generated artificial traffic producing massive congestion aimed at exploring the weaknesses of our algorithms, and are considering schemes to bolster those algorithms.

Because past line performance is a poor predictor of future line performance, we have modified the routing algorithm by deleting the use of line quality in estimating the delay to a neighbor. Although errors occur in bursts on the phone lines, there is no assurance that an estimate based in recent performance of the line will apply to the next interval.

During this report period, the Trouble/Status reports to the Network Control Center have also been somewhat expanded.

We have written a program that can simulate network activity. The purpose of this simulator is to test routing and anti-congestion schemes in real time and to measure their effect on

the code/  
walk  
tests of  
SR1

throughput. For this purpose, it is imperative that the simulation be able to run for reasonable lengths of simulated time in order to let singular cases occur. With the current ARPA net geometry and one Host generating as much traffic as possible, the simulator runs about one simulated minute per real minute, or at real time. This speed is quite adequate, particularly since it is programmed for the IMP prototype, which can be dedicated to simulation tasks for hours at a time.

The simulator permits any reasonable number of IMPs in the network, interconnected in a user selectable way, with one Host per IMP. Each Host may send a fixed length message at a fixed repetition rate (or RFNM limited) over a fixed number of links. Each fixed item may be set before the run.

The simulator is quite realistic, modeling IMP internal program delays, retransmission, our actual routing algorithm, and the asynchronous nature of the IMP clocks. It has one major difference from the real network in that it does not allow for the delay encountered in the processing of one message due to the interruption of that processing to process other messages. This delay is typically small (say, 1% of the total), provided the IMP throughput is not approaching capacity. In the current network, it is impossible to push an IMP beyond its computational capacity, so we do not consider this a strong limitation of the simulator.

### 3. HARDWARE

The primary hardware effort in this quarter has been devoted to insuring on-time machine delivery from Honeywell, to thorough IMP testing and debugging in the BBN test cell, and to field installation. Comparatively few technical problems have occurred with the hardware. A minor problem in the Distant Host interface was uncovered and corrected, as was also a potential hazard in the standard Host/IMP interface.

Discussions with Lincoln Laboratory designers uncovered the minor problem in the Distant Host interface. Specifically, the written specifications for the Honeywell driver card were not being met by their circuits. A circuit analysis indicated that an unbalanced differential signal would be produced and that the signal amplitude was not always sufficient. These problems were cured by deleting the termination network in the receiver and adding external resistors to the driver. A revision to the Host Specification (BBN Report No. 1822) that reflects these changes will be issued during the next quarter.

It has become necessary to change the connector type used for the Distant Host cable. The connector type originally specified has been preempted by a high-priority government project. This "mid-stream" change to the IMP, although a trivial matter, has been a considerable nuisance.

A potential hazard was located in the standard Host/IMP interface. A short delay was originally designed into that interface so that, following the receipt of the There's-Your-Bit signal from the Host, the bit would be accepted after a brief delay to allow for possible differences in transmission time between the Data line and the There's-Your-Bit line from the Host. These delays can arise from differences in driver and receiver speeds and from differences in wire lengths within the Host/IMP cable. Unfortunately, due to the nature of the Honeywell flip-flops in the interface, the delay does not always accomplish its purpose; in fact, if a change occurs on the data line to the IMP after the There's-Your-Bit signal arrives, ~~an incorrect data bit may be taken in.~~ This flaw is easily corrected by employing an unused delay in one of the packages of the interface. This simple modification involving a very few wires has been implemented at BBN and will be retrofitted to all machines by Honeywell.

(IMP/IMP  
 broke  
 - the errors  
 in the  
 sig ->)

The issue of retrofits has been a matter of concern for some time now, as Honeywell has fallen somewhat behind the original schedule. Although very little retrofitting has actually occurred to date, the plans to retrofit are now congealed. In addition to the retrofit addition of interfaces required to expand particular IMP configurations, the three are: replacement of non-ruggedized control panels with ruggedized panels at the first four IMP sites, a fix to reduce light driver noise in the first four sites, and a change to the auto-restart mechanism in the first six sites. These retrofits will be completed during the next quarter.

## 4. NETWORK TESTING

A second series of intensive network tests was performed over a three-week period on an eight-node subnet. These tests were designed to uncover residual system bugs, and to measure certain limits of system performance under extreme loading conditions. Approximately nine minor bugs were located and fixed. We were unable to push the IMP to its computational capacity, because the Host was unable to generate traffic fast enough.

*cloud  
together  
Hosts*

One interesting phenomenon observed during network testing was that in an IMP with steady state through traffic running at line capacity, there is no natural tendency for the store and forward queue length to become smaller once it has become large. In fact, by forcing a temporary imbalance between input and output, we were able to set this queue length to any value we wished and it remained there.

We concluded the alternate routing experiment on a triangular net that was begun at the end of our last testing session. Eight-packet messages were sent from one Host, say Host A, to another Host, say Host B, on 16 links at the RFNM limited rate. The following results were obtained as the limit on the maximum number of store and forward buffers (S/F) was varied. For  $S/F \leq 4$ , all traffic went out the direct path; for  $S/F = 5, 6, \text{ or } 7$ , the traffic

appeared to alternate between the two paths; there was little noticeable overlapping use in time of both paths until the S/F limit reached 7, after which the degree of overlapping began to increase. With enough room in reassembly for four messages (34 buffers), the throughput increased essentially monotonically until the S/F limit reached about  $30_{10}$ . There was no further gain in throughput with  $S/F > 30_{10}$ . With enough room in reassembly for 1 message (10 buffers), the maximum throughput occurred when the S/F limit was set to  $19_{10}$ . The limiting value of 85,000 bits/sec that was observed experimentally is about two times the limiting value when the data travels along a single path. The maximum peak occurs because conflicting demands for the minimum amount of reassembly cause increased packet retransmissions to occur and hence increase the time to reassemble certain messages.

Two curves of throughput vs. store and forward limit are shown in Figure 1, for a triangular net with half-second routing update. The upper curve is for a reassembly limit of  $42_8$  at the destination IMP (SRI); the lower curve, a reassembly limit of  $12_{10}$ .

An experiment was performed to determine whether any substantial interaction effect occurred when heavy Host traffic and heavy store and forward traffic are present in a single IMP. A stand-alone program in the UCLA Sigma-7 was used to generate traffic from the Host to itself via the IMP. We then arranged for RAND and UCSB to exchange message generator traffic in both directions through UCLA as an intermediate node. Finally, the

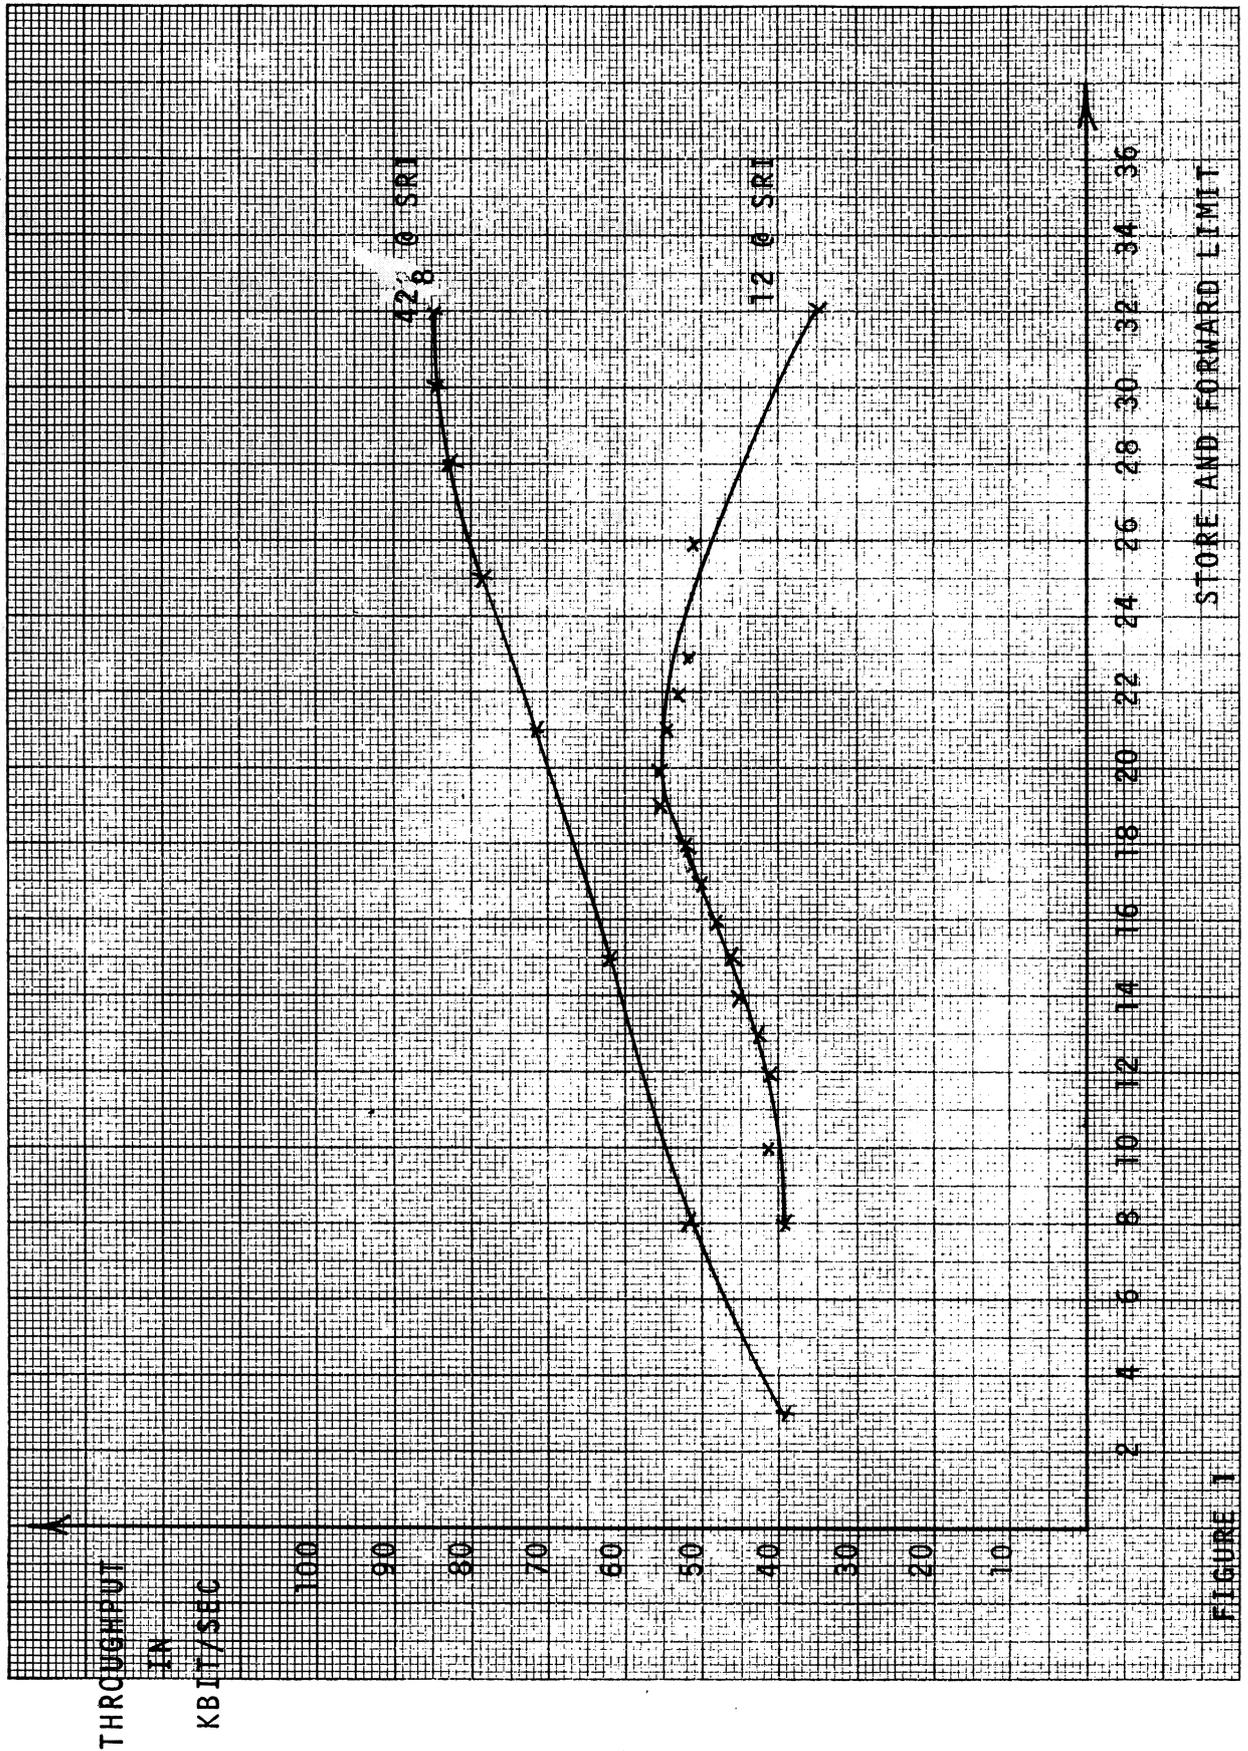


FIGURE 1

Host traffic and the store and forward traffic were allowed to occur simultaneously. No interaction effect was noted; i.e., there was no change in the number of Host messages and the number of store and forward packets when both types of traffic were present.

We experimented with the Host line set at 500 kilobits/sec again using the UCLA stand-alone program in the Sigma-7. The maximum rate at which the dedicated Sigma-7 could be made to handle messages (through its accumulator channel) was about 250 kilobits/sec and therefore the IMP was not able to be subjected to the planned maximum load.

## 5. NETWORK CONTROL CENTER

During the last quarter the network status changed from experimental to operational. A Network Control Center (NCC) was established at BBN in Cambridge to monitor the network, to take responsibility for attending to network trouble, and for handling coordination of network activity.

Status messages are automatically sent from each IMP to the NCC every 15 minutes, and more often when important changes occur. The messages are currently printed out in their entirety on a teletype connected to the BBN IMP, where they are monitored by BBN personnel during business hours. In the event of trouble, personnel at the Host site or the telephone company are called upon to help diagnose the problem and to select an appropriate plan to alleviate the problem. The status reports are also used to tabulate network up time.

An effort is currently underway to automate the process of reducing and tabulating the status data. When this effort is complete, we plan to share our line information with the telephone company via a low-speed connection and are prepared to have them accept a larger share of responsibility in maintaining their circuits. The NCC has the capability of determining the status and quality of each line in the network and the telephone company frequently calls the NCC at BBN to inquire about the status of lines.

The NCC also serves the function of coordination when, for some reason, an IMP must be taken down. This action is sometimes necessary for the site personnel to test or modify their interface (hardware and software) or for Honeywell to repair or retrofit IMP equipment.

The NCC has distributed its telephone number to the sites and to the telephone company, with instructions to call any time they have network problems to report, or want to coordinate some network activity. The NCC telephone is manned from 8 a.m. to midnight eastern time, Monday through Friday, and 8 a.m. to 4 p.m. Saturday. Site and telephone company personnel have been very cooperative.