

THE NETWORK CONTROL CENTER FOR THE ARPA NETWORK

A. A. McKenzie, B. P. Cosell, J. M. McQuillan, M. J. Thrope
Bolt Beranek and Newman, Inc.
Cambridge, Massachusetts, U.S.A.

Abstract

The ARPA Network allows dissimilar, geographically separated computers (Hosts) to communicate with each other by connecting each Host into the network through an Interface Message Processor (IMP); the IMPs themselves form a subnetwork that can be thought of as a distributed computation system. To detect failures in this system each IMP automatically and periodically examines itself and its environment and reports the results to the Network Control Center (NCC), at Bolt Beranek and Newman Inc., for action. The NCC computer, like any other Host, can itself fail without affecting network integrity; further, the NCC central processor can easily be replaced, in case of failure, by any standard IMP.

The present paper briefly describes the NCC hardware; discusses such software issues as NCC-related routines in the IMPs, data-collection and interpretation mechanisms, line status determination, IMP status and program reloading, and Host and line throughput; details NCC operations (manning, problem-handling procedures, track record); and summarizes overall NCC experiences and future plans.

I. Introduction

Almost four years ago the Advanced Research Projects Agency of the Department of Defense (ARPA) began the implementation of a new type of computer network. The ARPA Network provides a capability for geographically separated computers, called *Hosts*, to communicate with each other via common-carrier circuits. The Host computers typically differ from one another in type, speed, word length, operating system, etc. Each Host is connected into the network through a small local computer called an *Interface Message Processor (IMP)*; each IMP is connected to several other IMPs via wideband communication lines. The IMPs, all of which are virtually identical, are programmed to store and forward messages to their neighbor IMPs based on address information contained in each message.

In a typical network operation a Host passes a message, including a destination address, to its local IMP. The message is passed from IMP to IMP through the network until it finally arrives at the destination Host. An important aspect of this operation is that the path the message will traverse is not determined in advance; rather, an IMP forwards each message on the path it determines to be best, based on its current estimate of

local network delay. Since the path choices are determined dynamically, IMPs can take account of circuit or computer loading (or failures) in an attempt to insure prompt delivery of each message.

In three years the network has expanded from 4 to over 25 IMPs and is still growing. Early work on the ARPA Network is described in some detail in a set of papers presented at the 1970 Spring Joint Computer Conference¹⁻⁵. Additional work is described in a paper presented at the 1972 SJCC⁶.

An interesting aspect of the IMP subnetwork (i.e., the set of IMPs and communication lines) is that it can be considered a distributed computation system. Each IMP performs its own tasks relatively independently of its neighbor IMPs; nevertheless all IMPs are cooperating to achieve a single goal — reliable Host-Host communication — and in some cases, for example, the dynamic path selection mentioned above, each IMP cooperates with its neighbors in making reliable delay estimates for various path choices.

In any distributed computation system it is likely to be difficult to detect component failures quickly; the difficulty is increased in the IMP subnetwork by the wide geographic separation of components. For this reason we chose at the outset to incorporate automatic reporting functions in the IMPs as an aid to failure diagnosis. Each IMP is programmed to examine itself and its environment periodically and to report the results of these examinations to a central mediating agent. This agent has the function of collecting the (possibly conflicting) IMP reports, determining the most likely actual state of the network and, in the case of failures, initiating repair activity. The mediation function is performed by the Network Control Center (NCC) located at Bolt Beranek and Newman Inc. (BBN) in Cambridge, Mass. The mediating agent is the NCC computer, which is attached, as a normal Host, to the BBN IMP. It should be noted that although the NCC computer is an *important* component of the network it is not an *essential* component; as with any other Host it can fail without disturbing overall network integrity.

The NCC computer is concerned primarily with the detection of line failures and IMP failures. In addition, the NCC computer monitors the volumes of Host traffic and line traffic; these are parameters which can give advance warning of network elements whose capacity may need to be increased and which can be used for site usage

accounting. Finally, the NCC computer keeps track of other data, such as switch settings and buffer usage, for each IMP; these data are frequently helpful in diagnosing IMP failures.

The remainder of this paper describes the operation of the Network Control Center. Section II describes the NCC hardware located at BBN and Section III provides details of the overall software operation. Section IV discusses the manual procedures followed by NCC operators and technical staff in diagnosing and correcting network malfunctions. In Section V we have provided typical summaries of the types of information collected at the NCC in recent months, and mention some anticipated changes in NCC operation.

II. NCC Hardware

The central site NCC hardware consists of two packages, a central processor with 12K of 16-bit memory, a real-time clock and a "special Host interface", and a special set of hardware designed specifically for NCC functions. The current CPU is a Honeywell 316 computer; this choice provides two important advantages. First, the "special Host interface" required for connection to the network is exactly identical to the "standard Host interface" already designed as part of the IMP, thus reducing the implementation cost. Second, because all special hardware has been kept modular and external to the CPU package, if the NCC computer goes down for an extended period, it can be replaced by any standard IMP (Model 316, Model 516, or Terminal IMP). This is significant because we frequently have several IMPs on site in preparation for field delivery; thus the potential for substitution of the NCC machine is of practical value.

The special NCC hardware consists of two dial-up line controllers, a half-duplex Teletype I/O interface, and hardware associated with a panel of 32 display lights, a programmable audible alarm, and 16 control switches. All of this equipment is housed in a separate cabinet along with the required power supplies. When necessary, it can be simply connected to the I/O bus of an alternate CPU.

Two Model 35 ASR Teletypes handle most of the input and output functions. One, attached through the Teletype interface in the special hardware package, is dedicated to a print-only logging function while the other, the NCC computer's standard console Teletype, serves both as a report printer and as the primary source of operator input. Input can also be provided through the 16 control switches, and other output is given via the 32 display lights and the alarm. The dial-up line controllers are reserved for possible future use. The external I/O equipment is duplicated at nearby locations for the convenience of NCC personnel.

III. Software Operation

The IMP subnetwork consists of three principal classes of components: 1) a collection of wide-band common carrier data lines, 2) a set of IMP processors, 3) IMP system software. Network per-

formance can be affected by failures in any of the components in each of these classes. Therefore, in conjunction with our construction of the network, we had to develop procedures for quickly detecting and repairing component failures within any of these classes. In this section we will describe the software used to assist in detecting such failures.

NCC-Related Software in the IMPs

A basic assumption, which underlies the NCC effort, is that the most effective way of detecting failures is to have each IMP periodically compile a report on the status of its local environment and forward this report through the network to a mediating agent, the NCC. This agent has the task of collecting and integrating the reports from all of the IMPs to build up a global picture of the current state of the network. The data generation within each IMP is performed by two routines: a timing routine which controls the periodic execution of the report routine, and the report routine itself.

The timing routine used is the IMP's statistics mechanism. This mechanism establishes a network-wide synchronized clock which it uses to coordinate the execution of a set of self-measurement (statistics) routines which have been incorporated into the IMP. The bulk of the statistics routines are concerned with factors such as measuring IMP bandwidth capacity and storage utilization, etc. One of the statistics routines, however, is the "Trouble Reports" routine, which provides data to the Network Control Center.

The Trouble Reports routine, when initiated by the timing routine, interrogates various parts of the IMP system to determine which lines are alive, which Hosts are up, etc. It formats that information into a message which is forwarded to the NCC's data collection mechanism. Since space is at a premium in the IMP system, the routine does no pre-processing of the information; it is merely collected and forwarded.

In addition to the statistics and reporting packages, each IMP contains a small debugging package, DDT. DDT is a simple command interpreter capable of such functions as examining and modifying a memory word, clearing a block of memory, searching memory for a particular stored value, etc. DDT is structured so that it can be driven remotely through the network, returning any responses back through the network. The remote use of DDT is important to many NCC operations.

Each IMP contains several routines which perform such NCC-related actions as "looping" data sets and line interfaces, testing Host interfaces, etc. ("Looping" is the interconnection of circuit elements such that all transmissions from an IMP are returned to that IMP rather than being sent to the IMP at the other end of the line.) DDT is used to initiate and terminate these routines by modifying words in IMP memory which contain their parameters, including an enable/disable bit. For example, one routine monitors a word which, when changed to a line interface number, loops the appropriate interface. This particular ability is

vital to isolating the malfunction when a line goes down, so that we know whether to notify telephone company personnel to fix the line, or to notify Honeywell field engineering to repair the interface.

NCC Development

While the data generation scheme and (in large part) the data actually collected have remained invariant during the development of the NCC, the data collection/interpretation mechanisms at BBN have undergone steady evolution. In the first versions, while the network was small, the data were sent as ASCII text which was typed out on the BBN IMP console Teletype; personnel at BBN periodically scanned the typescript to determine if anything noteworthy was happening in the network. Since the collection was being done on a Teletype, a low bandwidth device, space within the message was at a premium; however, since a person was required to read it and make sense out of it, the format had to be intelligible. The only way to balance these factors was to omit the collection of much interesting data.

As the network became larger and more reliable, the proportion of status messages which said anything other than "everything's still OK here" decreased, thus making the location of the messages which required action on our part more difficult. The scheme we developed to make the location of critical messages somewhat easier consisted of having each IMP: 1) send us a status message every 15 minutes and 2) examine its status every minute and send an additional message at that time if it detected a *change* in status. Since these routines were being driven by the synchronized clock of the statistics package, the effect of this scheme was that every 15 minutes we would receive a block of "checkin" reports, one from each IMP; interspersed between these "checkin" blocks on the typescript there would be an occasional "change" report.

This setup functioned tolerably for some time, but eventually several factors combined to make it unwieldy. First, the number of IMPs in the network was constantly increasing, so that the amount of typescript which had to be scanned in order to determine what was happening in the network became overwhelming. Second, outside organizations became increasingly interested in receiving monthly reports on IMP and line performance; the prompt and accurate compilation of these reports by hand became more and more difficult. Third, there was pressure to take statistics on line usage and Host traffic in order to obtain advance warning of network elements whose usage was approaching saturation and to investigate accounting algorithms for network usage. All of these factors led us to install a Host on the BBN IMP which is dedicated to monitoring network performance and doing much of the bookkeeping required for our reports.

With a separate Host dedicated to monitoring the network, we were able to abandon ASCII text format in favor of binary format, and to expand the reports to include more internal status information as well as statistics on Host traffic and

line usage. We were also able to increase the frequency of reporting to once a minute for the "checkin", and to send "change" reports as soon as changes are detected. We also worked, and are still working, on the knotty problem of formalizing the heuristics which are used to integrate the (often conflicting) reports from the individual IMPs. The following paragraphs discuss several of the problem areas of greatest interest.

Line Status

For its own routing purposes, each IMP is continually measuring the quality of each of its data lines. Every half-second it sends a thousand-bit status message on each line and expects to receive a similar message from its neighbors. Each status message includes the number of the IMP which originated it. When an IMP receives a status message from one of its neighbors correctly, it marks its next status message to that neighbor with an acknowledge bit. Thus an IMP's receipt of a status message with the acknowledge bit set indicates that the line is in good condition. Conversely, whenever a half-second interval elapses and the IMP does not receive a status message with an acknowledgment of its own previous message, it counts an error on that line.

In conjunction with this acknowledgment scheme, an important system debugging feature is the ability to "loop" lines for test purposes. Each line is nominally a pair of independent one-way circuits, one in each direction. "Looping" is the interconnection of these circuits such that one end is disconnected and the other end receives its own transmissions. A line can be looped in one of three places: either inside the IMP's line interface or at the local data set (under program control), or at the remote data set (manually). The IMP system, by checking the origin of status messages, can detect looped lines.

Using its line error count and detection of looped lines, an IMP can make a simple usable/unusable decision, for its own purposes, for each of its lines. A line can, however, be "network unusable" for a variety of reasons (the IMP at the other end is down, the interface on the local IMP is broken, the line itself is broken, etc.) and at the NCC we must be able to distinguish amongst them in order to initiate the appropriate repair procedure. Therefore we supplement the IMP's report of whether it thinks the line is up, down, or looped with the IMP number of the IMP on the other end of the line, the total number of status messages sent on the line, and the total number of status messages received on the line (whether their acknowledge bit was set or not). The NCC takes the 3-way division from the IMP (up, down, or looped) and incorporates into it a 2-way division (status messages coming in or not) to form a 5-way breakdown of line status as seen by the IMP at one end of a line: up, down without errors (unusable but with status messages without errors being received), down with errors, looped, and "no information" (the IMP has not reported to the NCC recently). Every minute, for each line in the network, the NCC takes the latest status for

each end of the line and determines the state of the line according to the decision rules shown in Table 1. Whenever any line's state changes, a message is printed in the log.

The IMPs are essentially synchronized with regard to the generation of status messages; furthermore, status messages constitute a known constant traffic load on each line. Therefore, for lines whose state is declared up, a measure of line quality is given by the fraction formed by dividing the number of status messages correctly received by the number of status messages sent, since only line errors (detected by checksum hardware) will cause status messages to be incorrectly received. This fraction is printed in the log whenever the numerator differs from the denominator by more than one and the fraction is neither zero nor one. Thus we are alerted to line failures before the lines become completely unusable.

Since the IMPs have been designed to infer the network's topology dynamically, they are not directly concerned with the common carrier data lines; rather, they are interested only in which portions of the network they can access through a particular line interface. NCC personnel, however, must deal with the actual lines. A report from an IMP that the line connected to interface

2 has become unusable is not useful unless we can determine which line is actually connected to that interface. Toward this end, the NCC maintains a connectivity table which contains, for each line in the network, the IMP numbers for the IMPs at each end and the interface numbers that that line should be connected to. The NCC types a message in the log whenever it determines that a line has been moved from its nominal interface or when a report for a line not contained in the connectivity table is received.

IMP Status and Program Reloading

The NCC is faced with a difficult problem in attempting to determine that an IMP is no longer functioning. Since a broken IMP can't send us a message indicating that it's broken, we must infer this condition from the absence of its "check-in" messages. In the past, this decision was made after a scan of the typescript and the observation that the IMP had not checked in "for a while". The current NCC system declares an IMP dead when it has not reported for three minutes. Because of the effects of problems like network partitioning, this is an inadequate test for actually determining whether the IMP is up or down, but it does alert our personnel to the need for further diagnosis. For example, all lines to the IMP may

STATUS FROM HIGH NUMBER IMP

		UP	DOWN NO ERRORS	DOWN WITH ERRORS	LOOPED	NO INFORMATION
STATUS FROM LOW NUMBER IMP	UP	UP	IN LIMBO	IN LIMBO	IN LIMBO	UP
	DOWN NO ERRORS	IN LIMBO	DOWN	DOWN ON HIGH END	LOOPED ON HIGH END	UNKNOWN
	DOWN WITH ERRORS	IN LIMBO	DOWN ON LOW END	DOWN	LOOPED ON HIGH END	UNKNOWN
	LOOPED	IN LIMBO	LOOPED ON LOW END	LOOPED ON LOW END	LOOPED ON BOTH ENDS	LOOPED ON LOW END
	NO INFORMATION	UP	UNKNOWN	UNKNOWN	LOOPED ON HIGH END	UNKNOWN

The terms are defined as follows:

- HIGH NUMBER IMP - IMP with higher network address
- LOW NUMBER IMP - IMP with lower network address
- UP - The line is usable for both IMPs.
- DOWN - The line is unusable for both IMPs.
- DOWN ON ONE END - One IMP can transmit to the other, but not vice versa.
- LOOPED ON ONE END - The line is looped as seen by one IMP, but not as seen by the other.
- LOOPED ON BOTH ENDS - The line is looped as seen by each IMP.
- IN LIMBO - Conflicting reports from the two IMPs.
- UNKNOWN - Insufficient data to make a decision on the line's state.

TABLE 1: DECISION RULES FOR LINE STATE

be down, rather than the IMP being down.

In the rare case of network-wide failures it is often difficult to determine which IMP triggered the network failure, much less what caused that IMP to fail. Nevertheless, personnel at the NCC must attempt to make these determinations. To assist them, each status report that an IMP sends to the NCC contains a snapshot of the IMP's environment. The snapshot information is used to determine if the IMP is experiencing a transient or getting into some kind of trouble. This information includes the version number of the program running in the IMP, the storage utilization and the amount of free storage left in the IMP, the state of the sense switches and the memory protect switch (to detect unauthorized tampering), a list of the statistics programs which are enabled, a list of which Hosts are up, and an indication of whether tracing is enabled. The NCC logs any change in reported status and, in the event of a network failure, we attempt to correlate environmental data for individual IMPs with the network failure as a whole.

Since all IMPs run the same program, we built a small "bootstrap" routine into the IMP which, when initiated, sends out a request for a core image on a line selected either by parameter or at random. When any IMP receives such a request it returns a copy of its entire (running) program as a single message. The bootstrap routine then checks incoming messages on the selected line for correct length and checksum; when the core image is successfully received it is initialized and started. If an incorrect core image is received, the bootstrap routine sends another request. This facility provides a quick and easy way to obtain a fresh copy of the IMP system. Since the bootstrap resides in the protected memory sector, and thus is nearly always intact, site personnel are almost never required to handle IMP system paper tapes when an IMP requires reloading.

Program reloading can be initiated remotely by commanding DDT to execute a transfer to the bootstrap. Without the remote reloading ability, the only way to distribute a revision of the system would be to mail out paper tapes of the new program to each site, and then schedule a time, with personnel available at each site, to load and start the new version. With the remote reloading ability, however, we merely load the new version into the BBN IMP, direct BBN's neighbors to reload from us, then direct their neighbors to reload from them, and so on until the new version is propagated through the entire network. In fact, propagation of a new program release can be accomplished by one person in a few minutes, rather than requiring a month of planning and several hours of work by a nationwide "team". Also, since the procedure doesn't require assistance from site personnel, it can be scheduled to occur at a time when network usage is extremely low (typically very early morning, a time when site assistance would be most difficult to arrange), thus minimizing the loss of network availability.

Host and Line Throughput

With the change from ASCII to binary reporting and the consequent easing of bandwidth limitations on the NCC we have been able to take initial steps toward building an accounting facility for network usage. Toward this end, the IMP measures the amount of use each Host has made of the network in eight categories. The eight categories are the combinations of the following parameters: transmissions from and to the Host, inter-site and intra-site transmissions, and packet and message traffic. Thus, the eight categories are:

1. inter-site messages sent
2. inter-site messages received
3. inter-site packets sent
4. inter-site packets received
5. intra-site messages sent
6. intra-site messages received
7. intra-site packets sent
8. intra-site packets received

The IMP counts data transmissions only; control messages and RFNMs (destination-to-source message acknowledgments) are not included. The NCC tabulates Host traffic data from all the IMPs in the network. At the end of each hour it copies this table into a second table and then clears the first to obtain a clean "snapshot" for the hour, which is then printed on the report Teletype. This table is also added into a daily table which is printed at midnight every day.

In order to be able to better predict when lines may become overloaded, we also keep track of the line utilization in the network. The IMP measures the line throughput by counting the number of successfully acknowledged packets. The NCC accumulates these line throughput data for each line and types them out with the Host throughput at the end of each hour, and at the end of the day.

Visual and Audible Alarms

Although the computerization of the NCC virtually eliminated extraneous typescript, it was still desirable to free the NCC personnel from having to regularly check the typescript to determine whether action was required. We therefore attached a set of lights and an audible alarm to the NCC. The NCC maintains two sets of "virtual light" display information: which IMPs are alive and which lines are functioning. The NCC staff can select either of these sets for output in the physical display lights. This provides for a quick visual survey of the state of the network.

Whenever a line breaks, or an IMP stops working, the alarm is sounded and the virtual light for that IMP or line is flashed. This minimizes the time for the NCC personnel to notice and take some corrective action, while at the same time freeing them from having to watch the lights or log to achieve this rapid response to network failures.

IV. NCC Operation

The Network Control Center staff consists of five computer operators and several regular BBN

technical staff members. The operators are familiar with the operation of the NCC machine and, to a certain extent, with the diagnosis and resolution of network problems. The technical staff members are both hardware and software specialists, most of whom have participated in the design and implementation of the network from the outset.

As the network developed and the role of the NCC increased and became more clearly defined, a fairly comprehensive scheme for manning evolved. It became clear very early that 9 to 5 coverage with informal arrangements to contact staff members at home was insufficient. A dedicated Network Control Center telephone line was installed, with computer operators acting in a monitoring capacity to direct inquiries and problems to available staff members. This has become the single contact telephone for Host site personnel, the telephone company and Honeywell field engineering.

The present NCC program, with its detailed log, allows the operators to assume first order responsibility for network operation. Operators now man the NCC 24 hours a day, 168 hours a week; technical staff coverage is normally close to 50 hours a week, with additional "at-home" availability of key personnel. Routine chores are handled by the operators and only more complicated situations are referred to staff members. After hours, the operators attempt to contact specific staff members at home in the event that a problem arises or a phone call is received which the operators cannot properly field. In specific rare cases, such as attempting to pin down an obscure hardware malfunction, a problem may be preserved (i.e., not fixed) until a staff member can investigate. Even outside regular working hours, most problems are resolved, or at least under control, within a few hours.

There are several different means of handling problems, depending on severity and type. For routine controlled situations (such as IMP preventative maintenance and scheduled repair, scheduled Host testing, and scheduled line test and repair) the NCC operators coordinate the activities of the Honeywell field engineering, Host site, and telephone company personnel involved. We have established the policy that the state of an IMP or a line is not to be intentionally modified without first seeking the permission of the NCC. We insist upon strict adherence to this policy in order to prevent a deferrable outage from occurring during an unscheduled failure and thereby jeopardizing network integrity.

The alarm calls attention to IMP or line outages. The display lights, in conjunction with the log and the ability to obtain a quick printout of network status, usually make it fairly easy for the operators to determine what has failed. In the case of an indicated IMP failure, the operator on duty calls the IMP site, verifies that the IMP has failed, gathers some rudimentary information as to the type of failure, and enlists the aid of site personnel to bring the machine back on the network. If this is not possible, technical staff members are called in to investigate further. If a hard-

ware malfunction is indicated, Honeywell field engineering is alerted to repair the problem.

At present, IMP maintenance and repair are carried out under contract by Honeywell field engineering. Coverage is prime shift with guaranteed 2-hour response time. When circumstances warrant, however, the NCC will request extended coverage for repair or for standby backup. Most repairs are completed by the end of the day they are reported.

In the case of an indicated line failure, the operator performs a series of checks to confirm that the line has actually failed. This is necessary since some IMP failures appear in the log as line failures (the converse is also true). Diagnosis is performed from the NCC by using IMP DDT to test the terminal equipment. If a line failure has isolated a site from the NCC, the operator will contact site personnel and direct them in performing the tests for him. When a line problem has been confirmed, the operator notifies the appropriate telephone company office, frequently supplying considerable detail.

Each line is maintained and tested from a private line office at one end. Manned around the clock, these offices are equipped with test facilities for finding and repairing line problems. Unless there is a manpower or access problem related to local facilities, line failures usually are corrected within a few hours of the initial report. Maximum repair time is normally about a day.

For many NCC activities the cooperation of the sites is essential. Site personnel aid the NCC in the diagnosis of a variety of problems, help in recovering from IMP failures, and take local responsibility for the IMP. Their assistance is particularly useful when investigating obscure hardware and software malfunctions.

Our relationship with the organizations involved in network maintenance has been good. Honeywell field engineering, telephone company, and site personnel have a high regard for the conclusions reached in our problem analysis. This believability has been fostered by a good track record; in at least 75% of all failure reports to Honeywell or the telephone company, an actual problem has been detected. Line problems and many IMP problems are usually fully diagnosed and dispatched to the appropriate maintenance group within half an hour. Some more subtle IMP problems, however, occasionally require gathering data over a number of failures before a conclusion can be drawn.

V. Experience and Future Plans

A great deal of additional work is done with the NCC's log and summary reports in order to produce monthly reports on network status and usage for ARPA and other interested parties. Since the NCC machine has no secondary storage capability we are unable to accumulate monthly summary information on that machine; instead the daily summaries and log information must be used as input to manual preparation of Host traffic reports and IMP Down

Month	Average Line Outage	Average IMP Down	# of Nodes	Average Host Inter-site Output (packets/day)
September 1971	.59%	3.27%	18	51,386
October	1.66%	1.77%	18	95,930
November	1.65%	5.50%	18	116,515
December	3.21%	3.95%	19	107,896
January 1972	1.02%	1.92%	19	172,037
February	1.23%	2.73%	19	224,668
March	1.36%	4.00%	23	240,144
April	.88%	2.86%	25	362,064
May	1.11%	2.57%	25	505,639
June	.41%	.97%	29	807,164

TABLE 2: SUMMARY OF NETWORK OPERATION

and Line Outage summaries. A certain amount of judgment is used in the latter two summaries; several outages which the technical staff feels are due to a single cause are normally combined into a single (longer) reported outage. Table 2 provides summary information for an actual ten-month period of network operation.

Until early this year the Host traffic summaries were produced manually from the NCC's hard-copy summary reports. The NCC now punches a paper tape (on the report Teletype) of all daily summary information and this is used as input to computer programs which produce the reports more rapidly and accurately. Eventually, when experience indicates that several of the network's service Hosts are reliably up around the clock, we expect to have the NCC transmit all of the summary information through the network for storage and later manipulation. This will enable us to more easily provide answers to interesting questions such as:

- What are the peak hours of network use and what is the peak-to-average traffic ratio?
- What percentage of network traffic do single-packet messages constitute, and how does this percentage vary from Host to Host?
- What is the ratio of weekday use to weekend use?
- What percentage of line capacity is used during peak hours, on the average, and during weekends?

Although the data needed to answer these questions are available now, the data manipulation required constitutes a prohibitive manual burden. Thus, the installation of an NCC computer lifted one bandwidth limitation only to reveal another. In an attempt to deal with this new problem, we are planning to experiment with an additional Host which was recently added at BBN. This is the machine which is currently being used "off line" to process the paper tapes mentioned above.

Another change which is under consideration is automated single-point reporting of line problems. The NCC program, after appropriate automated line testing, could report confirmed line

failures directly to a Teletype at some telephone company central location via one of the dial-up line controllers. Telephone company personnel would then direct this report to the appropriate office for test and repair.

Finally, certain of the NCC command options will be made available to other organizations (such as the ARPA office) via one of the dial-up line controllers. This will be primarily to allow access to information on the overall state of the network, particularly the up/down status of the IMPs and lines.

Acknowledgment

This work was sponsored by the Advanced Research Projects Agency under Contract No. DAHCl5-69-C-0179.

References

1. "Computer Network Development to Achieve Resource Sharing," L.G. Roberts and B.D. Wessler, Proc. 1970 SJCC.
2. "The Interface Message Processor for the ARPA Computer Network," F.E. Heart, R.E. Kahn, S.M. Ornstein, W.R. Crowther, and D.C. Walden, Proc. 1970 SJCC.
3. "Analytic and Simulation Methods in Computer Network Design," L. Kleinrock, Proc. 1970 SJCC.
4. "Topological Considerations in the Design of the ARPA Computer Network," H. Frank, I.T. Frisch and W. Chou, Proc. 1970 SJCC.
5. "Host-Host Communication Protocol in the ARPA Network," C.S. Carr, S.D. Crocker, and V.G. Cerf, Proc. 1970 SJCC.
6. "The Terminal IMP for the ARPA Computer Network," S.M. Ornstein, F.E. Heart, W.R. Crowther, H.K. Rising, S.B. Russell, and A. Michel, Proc. 1972 SJCC.