AD-A016 604

A STUDY OF COMMUNICATIONS FOR THE SEISMIC DATA COLLECTION NETWORK

Howard W. Briscoe

Bolt Beranek and Newman, Incorporated

Prepared for:

Air Force Technical Applications Center
Advanced Research Projects Agency

30 June 1975

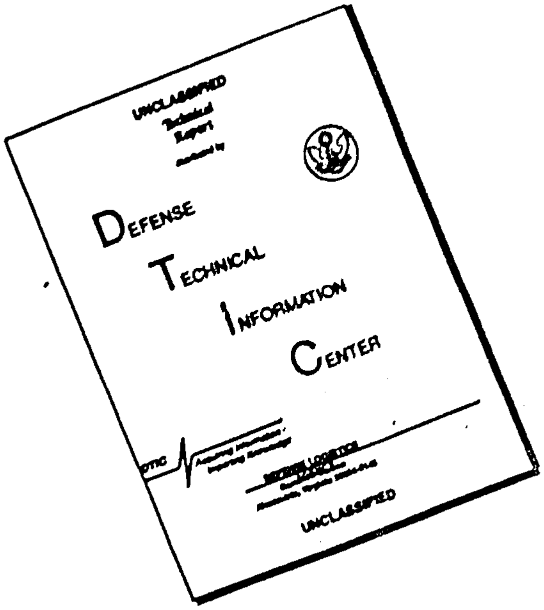

310116

AD-A016604

BOLT BERANEK AND NEWMAN INC

CONSULTING · DEVELOPMENT · RESEARCH

Report No. 3109

FINAL REPORT
A STUDY OF COMMUNICATIONS FOR THE SEISMIC DATA COLLECTION NETWORK

30 June 1975

APPROVED FOR PUBLIC RELEASE.
DISTRIBUTION UNLIMITED.

Prepared For:

VELA Seismological Center
312 Montgomery Street
Alexandria, Virginia 22314

Sponsored by the Advanced Research Projects Agency
ARPA Order No. 2551

FINAL REPORT
A STUDY OF COMMUNICATIONS FOR THE SEISMIC DATA COLLECTION NETWORK

30 June 1975

PRINCIPAL INVESTIGATOR: Howard W. Briscoe
Telephone (617) 491-1850 Ext. 415

This research was supported by the
Advanced Research Projects Agency of the
Department of Defense and was monitored by
AFTAC/VSC, Patrick AFB FL 32925,
under Contract No. F08606-75-C-0017

AFTAC Project Authorization No. VELA T/5701/BETR
Program Code No. 5F10
Effective Data: 1 August 1974
Total Amount of Contract: $70,000.00
Expiration Date: 30 June 1975

Prepared by:

Computer Systems Division
Bolt Beranek and Newman Inc.
Cambridge, Massachusetts 02138

REPORT DOCUMENTATION PAGE		**READ INSTRUCTIONS BEFORE COMPLETING FORM**
1. REPORT NUMBER	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Final Report: A Study of Communications For The Seismic Data Collection Network		5. TYPE OF REPORT & PERIOD COVERED Final: 8/1/74-6/30/75
		6. PERFORMING ORG. REPORT NUMBER BBN Report No. 3109
7. AUTHOR(s) Howard W. Briscoe		8. CONTRACT OR GRANT NUMBER(s) F08606-75-C-0017
9. PERFORMING ORGANIZATION NAME AND ADDRESS Bolt Beranek and Newman Inc. 50 Moulton Street Cambridge, MA 02138		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS Program Code #5F10 Project #VELA T/5701/BETR
11. CONTROLLING OFFICE NAME AND ADDRESS Air Force Eastern Test Range (AFSC) R & D Contracts Division (PMRB) Patrick AFB FL 32925		12. REPORT DATE 30 June 1975
		13. NUMBER OF PAGES ~~127~~ 130
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) AFTAC/VSC Patrick AFB FL 32925		15. SECURITY CLASS. (of this report) UNCLASSIFIED
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

APPROVED FOR PUBLIC RELEASE. DISTRIBUTION UNLIMITED.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Seismology
Computer Networks
Data Communication
Packet Switching
Seismic Data
Seismic Network

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This report presents the results of studies performed to support the implementation of a worldwide seismic data collection network using a mix of ARPA Network and leased line communications. Host level protocols for the ARPA Network links are specified. An analysis of the traffic generated by the seismic data subnet identified potential limits imposed by the bandwidth and buffer capacity at two nodes of the network. Operational monitoring and diagnostic procedures are recommended.

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

ABSTRACT CONT.

A backup plan for the system design is developed in case error rates in local communications to the seismic stations are excessive. Design specifications for the Seismic Private Line Interface computer for interfacing existing seismic station controllers to the ARPA Network are included as an appendix.

SUMMARY

The effort under this contract consisted of four tasks intended to verify the design of the seismic data collection network and to support the network implementation efforts.

Task I

This task consisted of maintenance and verification of the Host-to-Host protocols to be used on the ARPA Network paths of the seismic data network. The protocol maintenance function has involved coordinating protocol changes found to be convenient for implementation of the Host programs on the various Host computers.

Both experimental and analytical Techniques were used to verify the protocol designs. As expected, it was found that the total proposed seismic traffic would exceed the reliable capacity of the existing ARPA Network topology. The network topology is routinely reviewed and revised to account for changing loads so this problem will be remedied in the normal course of network operation. The study revealed a more serious bottleneck in the reassembly buffer space available in the IMPs at SDAC and at CCA. At CCA the buffer space is reduced because of a Very Distant Host (VDH) interface. At SDAC the problem is the large number of multi-packet messages that terminate in that IMP. Corrections for these problems are in progress and will be implemented before the problems become a restriction on the seismic traffic.

It was not possible to perform experiments over communication satellite links so the effects of the satellite hops could only be examined analytically.

Task II

The effort under this task was devoted to design of routine monitoring procedures to identify marginal operation of

the communication system. As a Host the CCP does not have access to data concerning the IMP subnet performance. The delay timer statistics for the time-marked input data and the output queue lengths can provide a great deal of diagnostic information about network performance. Procedures for monitoring and use of these indicators have been described.

Task III

In order to use existing seismic stations with a minimum of interference at the station, the system design uses the concept of a mini-Host on the ARPA Network to interface with the seismic station controllers in a manner compatible with the station controller design. Under this task the specifications for the mini-Host, called a Seismic Private Line Interface (SPLI), have been prepared. The specifications include the Host protocol descriptions for communication with the CCP.

Task IV

The study performed under task I indicated that the buffer sizes and Host protocols specified for the seismic data network were adequate for circuits comparable to those that have been used for VDH and RJE connections in previous ARPA Network experiments and applications. However, the quality of the local communication circuits in the vicinity of the seismic stations could conceivably be so poor that network performance would be unsatisfactory. Therefore, under task IV we have examined possible design changes that could be used to compensate for unreliable communication circuits. Additional buffering throughout the system and the implied increased delay times for the fixed delay paths would be the most effective modification but would be quite expensive.

Finally, planned modifications to the ARPA Network that have been influenced to some extent by the studies performed under this contract will be implemented later this year. We recommend that the results of tasks II and III and the protocols be reviewed when these changes are made. Improvement in the seismic network performance and the performance monitoring capability may be achievable as a result of the modifications.

TABLE OF CONTENTS

INTRODUCTION

Background

As part of the effort under the Vela program for improving the capability to detect and identify underground nuclear explosions by seismic means, ARPA is supporting the development of a worldwide network of seismic stations. Some of these stations will communicate on-line with the processing center at SDAC and with a large archival storage system at CCA. Recommendations for the design of the communication, processing, and storage system were prepared under previous studies. Under the present contract we have been performing various engineering studies in support of the implementation of this seismic data collection network, concentrating in particular on the communication subsystem. The communication within the seismic network will use both leased line and ARPA Network paths. We have been most concerned with the latter.

Objectives

Four tasks were assigned in the contract work statement. Task I involved the maintenance and verification of the Host/Host level protocols for the ARPA Network paths in the seismic network. Task II involved design of procedures for operational monitoring of the performance of the seismic subnet in the ARPA Network. Task III required the generation of specifications for the Seismic Private Line Interface (SPLI) which will act as a Host to interface seismic station processors with the ARPA Network. The final task consisted of special studies to be assigned during the contract performance period. As a result of concern about the quality of communication facilities available for connecting to the remote seismic stations, the effort under Task IV consisted essentially of preparing a backup or contingency design which considered such steps as more buffer capability at the stations and more sophisticated Host/Host protocols on these remote links.

In the body of this report, each of these tasks is discussed in more detail. Copies of the protocols, SPLI specifications, and an interim report on the interaction of the seismic traffic and the ARPA Network that provided a basis for the analysis for each of these tasks are enclosed as appendices to this report.

TASK I: PROTOCOL MAINTENANCE AND VERIFICATION

Introduction

The major objective of this task was to verify the adequacy of the planned protocols for use of the ARPA Network to achieve the necessary reliability and throughput. A secondary objective was coordinating and maintaining updated versions of the special Host/Host protocols.

Protocol Maintenance

At the beginning of the contract period it was felt that the protocols were close to being frozen. However, as the implementation of required computer programs proceeded, modifications to facilitate the implementation under the available operating systems and to accommodate changes in the archival file formats were requested. This task then involved achieving agreement on those changes that seemed most valuable and least disruptive to other parts of the system. As a result of these perturbations, the protocol specifications were not finalized until quite late in the performance period. The final versions of the following protocols are included as Appendix A to this document: 1) NORSAR/CCP, 2) DP/CCP, 3) SIP/CCP, 4) 360/44/CCP. The protocols for ILPA/CCP and KSRS/CCP communication are described in the SPLI specifications in Appendix C.

Protocol Verification

The communications protocols for the seismic data subnetwork were verified by a combination of analysis and experiment. The throughput capacity of the ARPA computer network was measured in a series of experiments last autumn with support of the Network Control Center. It was found that the maximum data rate over a standard 50 kilobit/second (Kbs) line is about 37 Kbs. (The remaining 13 Kbs are used by the ARPA

Network communications protocols and routing messages.) As the number of hops in a network path increases, particularly with a satellite hop, the maximum data throughput rate will begin to decrease because the round-trip delay increases and because the network protocols allow only eight messages to be in transit between two Hosts at any one time. This phenomenon begins to occur in paths of about 12-15 conventional hops long. (It is not expected to bother the seismic data transmission since round trip delays are expected to be the order of 3-4 seconds including a satellite hop and the data rate is only one 5 kilobit data message/second from each site.)

A detailed analysis of the data flow paths was performed and transmitted to the sponsor as BBN Report 2995, "Use of the ARPA Network by the Seismic Data Collection Network". It is included in this report as Appendix B.

Basically the report indicated several areas where the ARPA Network would have to be expanded to accommodate the large amount of anticipated seismic data. These areas were the re-assembly buffer space at the CCA and SDAC network connections and the capacity of the lines in the Northeast Corridor between CCA and SDAC.

The report also indicated that the siesmic data would experience a nominal delay of about 3 seconds between the SPLI and the CCP. No experiments could be performed to substantiate this analysis since the proposed satellite link was not available.

TASK II: OPERATION VERIFICATION AND MONITORING

Introduction

Based on the conclusions of BBN Report 2995, monitoring several characteristics of the seismic data subnetwork will provide an adequate measure of system performance. The time delay statistics from the various sites to the CCP will give an indication of the current state of the communications path as well as providing information about changes in path characteristics due to seismic subnetwork or ARPA Network changes. Monitoring the amount of reassembly buffer space available at the SDAC IMP will provide information on the availability of this scarce resource as more and more of the seismic subnetwork comes online. Finally, an indication of the available throughput from the CCP to the mass store is given by monitoring the length of the CCP's output buffer. Of these measures, the CCP has a capability for operator examination of the delays, there is no direct method of determining the reassembly buffer space use in the IMP, and the interrogation of the SIP output queue by the CCP could easily be implemented. Use of these measures for monitoring the seismic network performance is discussed in more detail below.

Potential Problem Areas

Data transmission problems, if encountered, are expected to fall into three categories. First, trouble may occur due to a hardware or software problem in the part of the communication link where SDAC has responsibility (the CCP, the SPLI, the station controller, etc.). In general, these problems will cause a cessation of data being received from the affected sites.

Second, a software or hardware failure in the IMP subnetwork could cause trouble (if the failure influences the route taken by the seismic data). The problem may take the form of increased delays or may result in a cessation of data from one or more sites. In either case, the responsibility lies with the ARPA Network Control Center.

Finally, situations may arise where the ARPA Network is performing nominally, but demand for a limited network resource has increased to the point where data transmission rates begin to suffer. A problem such as this could occur if the amount of non-seismic traffic competing with the seismic data for line bandwidth over a certain path increases to the point where the demand for throughput exceeds the line bandwidth. Another example would be when the seismic network grows to the extent that the scarcity of message reassembly buffers at the CCP IMP severely limits throughput.

Other problems may occur because of the current specifications for a Host-IMP interface. If a Host begins transmitting a multi-packet message and the network cannot allocate reassembly space, the IMP will stop accepting the bitstream from the Host and the interface will become blocked. In this situation, the Host cannot communicate with anyone else or the network until reassembly space is allocated and the IMP starts accepting the bitstream from the Host again. This is an example of the current so-called "blocking" interface. (The possibility of implementing a non-blocking interface is discussed later.)

Monitoring and Diagnostic Procedures

In diagnosing a problem which hampers the operation of the seismic subnetwork, information to help localize the source of the problem and to suggest its possible causes may be gained by observing which sites are affected by the problem. For example, if data from only a single site were being lost or experiencing long delays, one would first suspect the parts of the transmission path for data from that site which did not overlap that of data from any other site. In general, this would be the parts of the transmission path starting with the affected site and working towards the CCP.

Conversely, if several sites begin experiencing difficulties simultaneously, one would suspect problems on the part of the transmission path in common to the two data streams. If all the sites begin experiencing difficulties simultaneously, one would expect the problem to be at or near the CCP-end of the transmission path.

The ARPA Network IMPs can provide snapshots and packet tracing as two debugging tools, but these facilities are not accessible to an ordinary Host. Consequently, the constant monitoring which must be done from the CCP cannot take advantage of these tools. Since data coming into the CCP has een time-stamped at the SPLI, it would be possible to incorporate a program for measuring time statistics (e.g., a short-term maximum delay and a long-term average for each site) directly into the CCP. By examining the pattern of which sites are sending data, which are acknowledging Hello messages, and which ones have increasing delays, a rough diagnosis of any problem is possible without NCC assistance.

For example, figure 1 shows a possible decision flow chart for the situation where the CCP is still getting data from a site, but the delay statistics (short-term, long-term, or both) indicate that a marginal situation has been encountered. Similarly, figure 2 shows a possible decision flow chart for the situation where data from a site (or sites) stops abruptly for an extended period (one minute or greater).

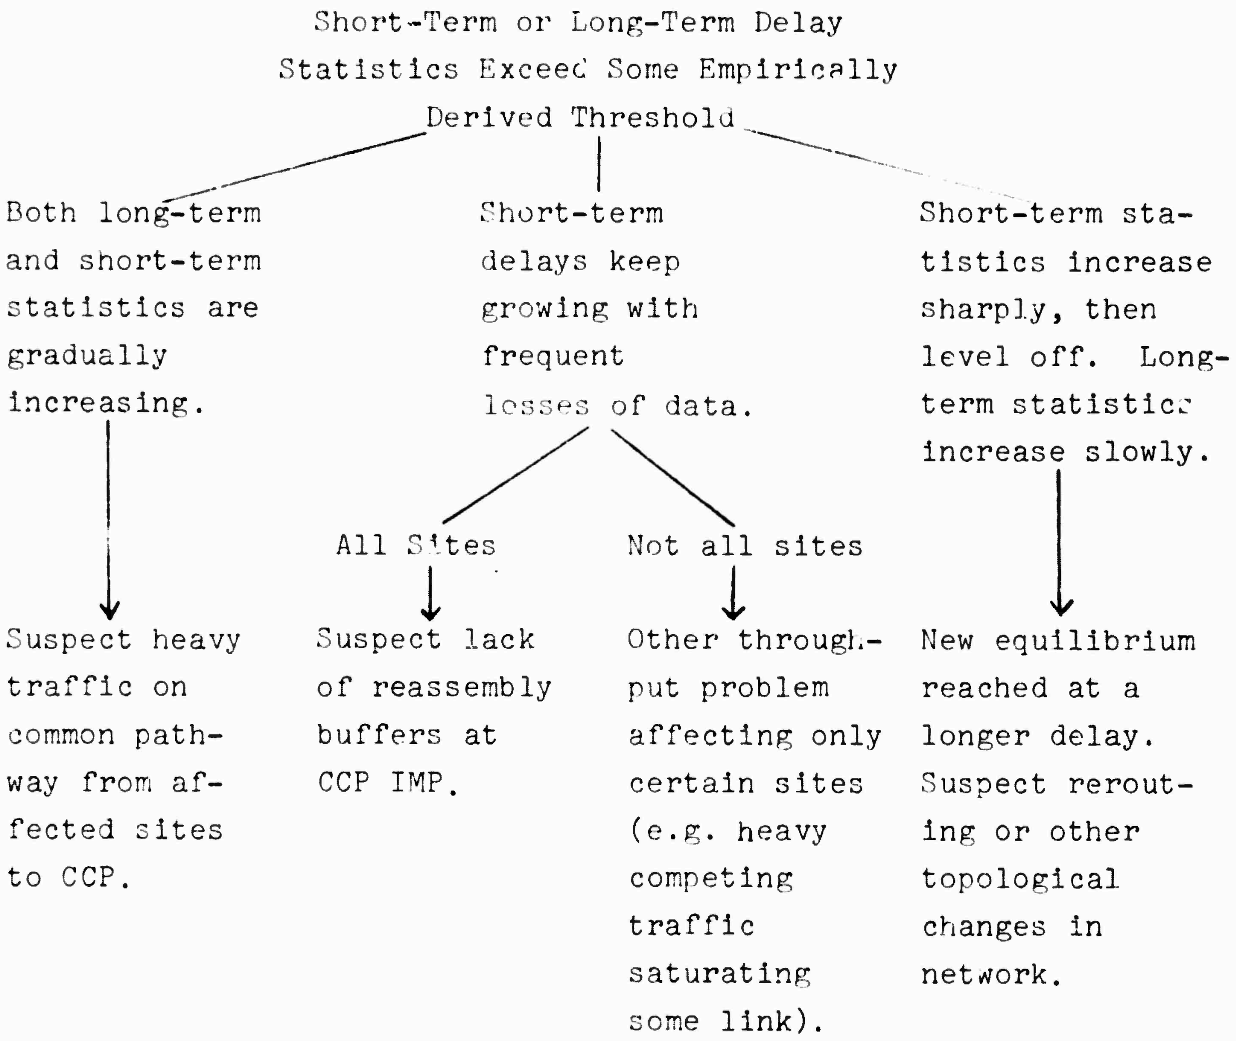

Figure 1. Decision Flow Chart For Delay Statistics Outside Empirically Determined Normal Range

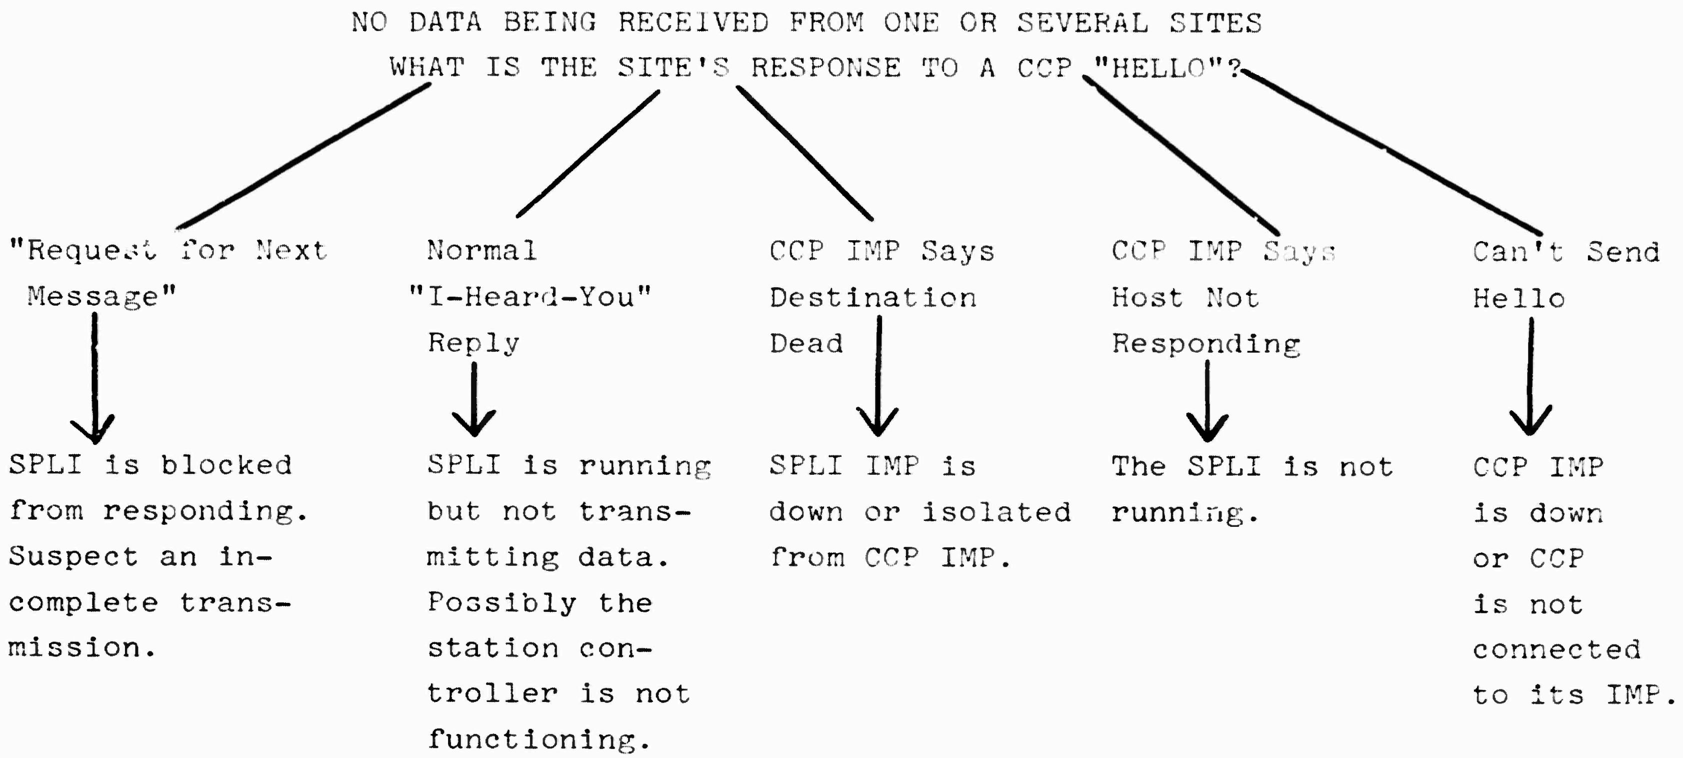

Figure 2. Decision Flow Chart for a Cessation of Data (Blocking Interface)

In detail, figure 1 incorporates the following scenarios. If traffic were heavy on the ARPA Network, it would be competing with the seismic data for network resources, most notably line bandwidth and processor bandwidth. This will cause delays to increase for the seismic data as the amount of competing traffic increases because the seismic data will spend more time waiting in queues for the resources to become available. Consequently, if the CCP operator were observing delay statistics, he would see the short-term and long-term averages fluctuate slowly as a function of other traffic.

If it is observed that the short-term average keeps growing (that is, the data keeps falling farther and farther behind real time) and frequent losses of data occur, then it is likely that the throughput capacity necessary to transmit the seismic data to the CCP is not currently available. This could happen if a line became saturated with traffic, for example, or could not support its usual throughput because of a temporary technical difficulty.

If this kind of behavior is observed for data arriving from all sites, then one should strongly suspect that there is inadequate throughput in the seismic subnetwork bucause of the scarcity of message reassembly buffers at the CCP IMP. Since reassembly space must be reserved before a (multi-packet) data message can be sent, a scarcity of buffer space will cause the data from all sites to become congested ("backed-up").

If there is a failure of one of the components of the ARPA Network, the network will be able to re-route messages around the problem in most cases. If a seismic data stream should have to undergo such a re-routing, it would be reflected in the delay statistics in the following manner. First, the short-term delay would increase rapidly during the short period while the new route is being established. (Actually, there is a brief cessation of data which may or may not last long enough to exceed the SPLI buffer capacity.) Then, the short-term average will level off at new equilibrium. The long-term average will correspondingly grow to a new equilibrium when the new route is established. Again, this phenomenon may be observed in data received from one or several sites depending on how many data streams had to be re-routed because of the failure.

A data stream may also cease abruptly. In this case the CCP begins to send "Hello" messages to the "silent" SPLI. If the SPLI is running, it will reply with an "I-heard-you" (IHY) message which indicates to the CCP that the SPLI is functioning. If the SPLI is not transmitting data because the site station controller is not functioning, the CCP will receive the IHY and the CCP operator can conclude that the problem is upstream of the SPLI. If the SPLI is not running or if the SPLI IMP is isolated from the CCP IMP, the IMP subnetwork will return the appropriate reply in response to the CCP's attempt to send a "Hello" to the SPLI.

If a data message is lost within the ARPA Network, it will take between 30 and 60 seconds for the network protocols to resolve this incomplete transmission. Consequently, it will appear to the CCP that the data stream has ceased suddenly. During this period of time, the SPLI-IMP interface is blocked (the possibility of a non-blocking interface is discussed later). Consequently, the IHY message from the SPLI is not allowed into the network until the incomplete transmission is resolved and the data messages begin to be transmitted again. From the CCP's point of view, it will look as if there has been no response to its "Hello" for some time, until the interface becomes unblocked. However, it is possible for the CCP to know that the "Hello" was received by the SPLI, because it will receive the RFNM (Request for Next Message) from the IMP subnetwork acknowledging the successful delivery of the "Hello" message.

If no message is received by the CCP, the implication is that its "Hello" message got lost by the network. In this unlikely case, subsequent "Hello" messages will be held at the SPLI IMP until the incomplete transmission caused by the missing "Hello" is resolved (30-60 seconds). This is done because the IMP network maintains the order of transmission for messages between a Host pair.

One final possibility is that the CCP IMP is not functioning properly. In this case data from all sites would cease abruptly and the CCP would not be able to send its "Hello" messages.

Note that figure 2 uses the response from the SPLI (I-heard-you) or the IMP subnetwork to determine the situation. Consequently, to fully utilize this decision tree, the CCP program would have to be augmented to communicate these responses to the CCP operator.

Long Term Performance Monitoring

Monitoring the CCP output queue length and keeping statistics about its growth can be done by augmenting the CCP program. This would provide the CCP operator with a direct measure of the ARPA Network's ability to carry the seismic data from the CCP to the SIP at the required throughput rate.

In addition to the statistics taken at the CCP, it is possible for the CCP operator to ask the NCC operator to use the snapshot facility on the source (SPLI) or destination (CCP) IMP if trouble is suspected there. Snapshots of intermediate IMPs can also be taken to determine routing and store-and-forward queue lengths.

Monthly statistics would also be useful in monitoring the growth of the seismic data subnetwork and the effect of changes and growth in the ARPA Network itself. Snapshots of the IMPs along the seismic data transmission path taken once a month would provide information as to changes in the nominal routing and store-and-forward queue lengths. With the cooperation of the Network Control Center, packets may be traced as they are processed by the IMP subnetwork and statistics recorded as to the length of time spent in the various processing tasks in each IMP as well as the send and receive times for each packet.

In addition, excerpts from the NCC monthly statistics could be used to monitor long-term effects of ARPA Network changes on seismic data transmission. These statistics include line usage, number of packets between IMP pairs, and line and IMP down times. The daily and hourly statistics might be helpful in a postmortem evaluation of some situations.

Delay statistics can best be calculated and stored at SDAC since the seismic data will be time-stamped and the delay measured at the CCP. In general, the ARPA Network cannot conveniently measure delay statistics on a regular basis.

For example, a statistic which might be valuable in studying the distribution of competing traffic is the delay as a function of time-of-day. This statistic would indicate if operation of the seismic data collection is nominal, marginal only during "rush" hours, or always marginal due to the amount of non-seismic traffic.

Recommendation for Future Effort

The ARPA Network is continually growing, and new features are continually being added to the message-handling software in the IMPs. Several changes planned for the end of this year*

---

*See ARPA Network Information Center Memo #NIC 32655, June 4, 1975.

have been influenced by this study and could improve seismic data throughput and the capability of the CCP to diagnose marginal situations. The first change is the inclusion of a handling type (currently only regular or priority is used) which would allow Hosts to communicate on several logical channels. If one channel becomes blocked due to an incomplete transmission, the Hosts (SPLI and CCP) could continue communicating on another logical channel until the Network resolves the incomplete transmission (which usually takes 30-60 seconds).

The second change is the implementation of a non-blocking Host-IMP interface. This would be an aid to CCP-generated experiments. For example, if the CCP were receiving data from a site more slowly than usual, it could initiate a Hello message to that site on a different logical channel. If it received a time-stamped I-heard-you reply and the CCP noted that its delay was nominal, one could conclude that lack of message reassembly space at the CCP-IMP was causing the slowdown. This is because the essential difference in the network's handling of single and multi-packet messages is the reservation of reassembly space at the destination IMP for multi-packet messages. Currently, the I-heard-you reply would not be sent out until after the multi-packet message because it would be blocked at the source Host-IMP interface.

Also, the implementation of a non-blocking Host-IMP interface along with the ability of the CCP to send priority Hellos and receive time-stamped priority I-heard-you messages from the SPLI would provide a means of measuring delay due to store-and-forward queue length in the network. This would be done by comparing priority I-heard-you delays with regular I-heard-you delays. (The priority messages are placed at the front of the queues whenever possible.) Delay due to queue length is an indication of the amount of network traffic competing with the seismic data for the network resources.

We therefore recommend a review of the seismic subnet protocols and performance monitoring procedures after these ARPA Network changes have been incorporated.

TASK III: SPLI SPECIFICATIONS

Requirement

In order to minimize the cost of implementing the seismic data collection network, sites which have been or will be implemented to serve other requirements will be integrated into the network. Thus the data collection network function is often a secondary mission for the station. It is, therefore, important that connection to the network should introduce a minimum of redesign and/or operational complexity. Several of the stations were designed to transmit on-line data from the station controllers over leased lines. Rather than modifying the station controllers to implement Host protocols when using the ARPA Network it was decided to introduce a minicomputer Host which would accept data from the station controller in the previously defined format. The minicomputer would then reformat the data and provide Host protocol for transmission over the ARPA Network. To the station controller the SPLI appears as a modem on a leased line and to the ARPA Network the SPLI appears as a Host computer.

In order not to constrain the location of the SPLI, the interface between the SPLI and its ARPA Network IMP uses the Very Distant Host (VDH) design.

Under the initial seismic network configuration, two forms of the SPLI were required, one for ILPA and one for the KSRS class stations.

Under this task we were required to prepare specifications to be used in procuring the SPLIs. Since the SPLI must implement the special Host/Host protocols for communicating with the CCP, the preparation of SPLI specifications was closely related to the effort under task I and the communication protocols are defined in the SPLI specifications. The resulting SPLI specifications are included as Appendix C of this report.

TASK IV: SPECIAL STUDIES

Background and Definition

Because the seismic data collection network will be the secondary mission for some of the seismic stations, the system design attempts to minimize the effect of the network on station operation and maintenance. In this design the problem of adapting to station-to-station peculiarities is handled in the CCP; the CCP is a general-purpose computer located near adequate program support and its use does not put any burden on station personnel.

The design objective was therefore to maintain highly reliable communication with the on-line stations with minimum bandwidth local communications and minimum equipment (for minimum maintenance and operation) at the station. This objective requires a tradeoff between buffer size and program complexity (required by sophisticated Host/Host protocol) vs. minimum equipment and bandwidth in the field. Based on the analysis described in Appendix B and on previous ARPA Network experience a buffer size of 8 to 10 seconds was chosen. It was decided not to attempt Host level retransmission, and the special Host/Host protocol has been kept to a minimum.

This compromise design caused some concern and, under task IV, we were asked essentially to recommend contingency plans in case the resulting reliability was unsatisfactory.

Analysis

For the analysis, the communication with a remote station is considered in four segments with significantly different characteristics. These segments are 1) station controller to SPLI, 2) SPLI to IMP, 3) IMP subnet including the satellite hop, and 4) destination IMP to CCP. These segments are discussed in more detail below.

The communications circuit between the station controller and the SPLI does not have any error detection or correction capability. Consequently, we recommend that the SPLI be co-located with the station controller to avoid reliance on possibly marginal communication circuitry for this unprotected path. Without extensive changes in the station controller there is no way to improve reliability over this path.

The link between the SPLI and the SPLI IMP will be a Very Distant Host (VDH) interface. A VDH interface does provide for checksumming and retransmission of data with detected errors, so that high probability of correct data is possible. One problem which may occur results when the probability of a packet (1024 bits of data) being sent correctly the first time begins to decrease. This causes the number of retransmissions to increase and consequently the effective bandwidth of the communications line to decrease. If this effective bandwidth decreases below the required throughput rate of the seismic data (5066 bits/second for KSRS and 1342 for ILPA), data losses will occur. The solution to this problem is to provide a higher throughput by increasing the circuit bandwidth (use of a higher bandwidth line or several lines in parallel).

Once the data is within the ARPA Network it is reliably transmitted to the destination CCP IMP. The ARPA Network, of course, is designed to permit rapid and reliable digital communications between Host computers. Consequently, responsibility for error checking, retransmission, discarding of duplicates, acknowledgments, and determination of the fastest route through the network rests entirely with the IMP subnetwork. Problems which may arise are expected to fall into one of two categories.

The first category consists of hardware or software failures in the IMP subnetwork. In general, the IMP subnetwork will be able to respond to such failures by rerouting the seismic data (as well as other traffic) around the failure. At the CCP, this will result in longer delays for a few data messages until the new path is established.

If the delay exceeds the buffering capacity, data will be lost. Obviously, increasing the buffer size will decrease the susceptibility of the seismic network to this kind of problem. During the transient, it is possible that a packet can get lost, and this will result in an incomplete transmission for the message containing that packet. It will take the IMP subnetwork between 30 and 60 seconds to resolve this problem in its message accounting and to inform the source host of an incomplete transmission; consequently, to recover from an incomplete transmission it is necessary to have upwards of 60 seconds of buffer space in addition to giving the SPLI the capability to retransmit the message.

Another situation for the CCP operator to be aware of is scheduled outages of key IMPs for preventive maintenance and new releases of IMP software. A backup plan to cover these contingencies might consist of local recording of seismic data to provide a large amount of effective "buffering" during the outage.

Note that larger buffering capacity at the SPLI means that the "catch-up-time" (the time needed to completely empty the buffer while it is being filled with new data) will be proportionately increased. If it is desirable to reduce the catch-up-time, it can be done only by increasing the bandwidth of the lines connecting the SPLI to its IMP.

The second category consists of situations where the IMP subnetwork is performing perfectly, but the seismic data cannot be transmitted quickly enough because of network resource limitations. An example of such a situation is when the throughput demands placed on the network by the seismic and regular traffic exceed the capacity of the network. In this case, possible contingencies include increasing buffer capabilities if the conditions are transient or reducing the amount of seismic data being sent (sacrificing some data so that the rest of it gets through) if the conditions are more long-lived.

For problems in category two, the best long-term approach is to pinpoint the area of congestion to determine precisely what resource (bandwidth on some line, message reassembly buffers, etc.) is scarce. Then, as needs dictate, the ARPA Network can "grow" to handle the expanding demands for throughout from the seismic network.

Finally, it could be argued that a Host-level positive acknowledgment scheme to allow retransmission of a message that does not checksum correctly at the CCP should be implemented to protect data bits from incorrect transmission at the destination Host-IMP interface. (The source interface contains error detection and retransmission capability since it is a VDH interface.) However, such a protocol would duplicate many of the features of the IMP subnetwork. The only new feature it would allow is retransmission of messages which were altered at the IMP-CCP interface, and there is no reason to expect these errors to be frequent enough to compromise the data quality . Such a change to the CCP-SPLI protocol must also allow for several messages "in flight". If it does not, then the maximum throughput possible with the protocol would be smaller than the input data rate. Increased buffering and a retransmission protocol would have to be incorporated in the SPLI so that messages with detected errors could be retransmitted.

Conclusions

As a result of our study of contingencies for the seismi network, we have reached the following conclusions. First, the incorporation of additional memory (buffer space) at the SPLIs and CCP would allow the seismic network to survive all transient conditions likely to be experienced on the ARPA Network. Conditions which could not be survived are insufficient throughput due to a scarcity of network resources for an extended period (several minutes) and isolation of an SPLI from the CCP due to a hardware problem for an extended period.

Secondly, if the amount of memory is increased at an SPLI, it is strongly reccmmended that the bandwidth of the communications circuits from the SPLI to its IMP be increased in order to increase the "catch-up" rate. This would allow a more rapid recovery after a transient slowing or cessation of data and therefore decrease the time it would take for the SPLI to be able to survive another transient without loss of data.

Third, in order to survive an incomplete transmission without loss of data, the SPLI buffering capability would have to be increased to about 60-75 seconds and the protocol changed so that the SPLI can retransmit the incomplete message. The CCP buffering capability would have to be increased accordingly.

Finally, the cost of such changes would be the cost and installation of the additional memory and the cost of the protocol changes in the CCP and SPLIs. Another important point to bear in mind is that the fixed delay in the seismic data seen by the recipients of the data from the CCP would be increased commensurably.

CONCLUSIONS AND RECOMMENDATIONS

As indicated in BBN Report 2995 (included in Appendix B), the design of the seismic communication path using the ARPA Network is basically sound. In that report and in the studies conducted since that report was issued, a few potential problem areas have been discovered and are being corrected.

Some network resources will have to be expanded as the seismic traffic load grows, but this form of expansion is routine. Last January, the ARPA Network had 55 nodes. Currently (July 1975) it has 60 nodes, and the changes necessary to allow more IMPs than the previous limit of 63 are being made. In addition, major topology changes are planned which will increase the number of coast-to-coast lines and East Coast Corridor lines. The result will be that the longest possible network path will be reduced to nine hops. Such changes are proof that the ARPA Network's flexibility allows it to grow to meet users' needs in an organized fashion.

A potential problem is the ability of the seismic data to survive transients in the network. The most difficult transient to predict is a re-routing transient. The length of time needed to change routes may or may not cause data to be lost, depending on the exact circumstances. Our analysis indicates that the amount of buffering currently planned for the SPLI is sufficient for most circumstances; however, if re-routing transients repeatedly cause data losses, additional memory at the SPLIs and the CCP can provide the increased buffering needed to survive these re-routing transients.

Because the ARPA Network is currently undergoing a number of topological and protocol improvements, our final recommendation is that the impact of these changes on the seismic network be explored when they are implemented. In this way the changes, particularly the non-blocking interface and the new handling types, may be studied both analytically and experimentally to see what changes could be made to the CCP and the SPLIs to improve performance, reliability, and on-line monitoring functions.

APPENDIX A

COMMUNICATION PROTOCOLS

From: H. Briscoe
To: Lt. M. Marcus, VSC
Subject: Communication Protocol between the CCP and the SIP.

1. Introduction

1.1 Background

A worldwide seismic data collection network including approximately 6 on-line array stations is being implemented under ARPA sponsorship. The objective of the network is to provide data for research in nuclear test detection and identification. A network event list will be prepared within a couple of days of real-time using on-line event detection and computer aided event analysis at the Seismic Data Analysis Center (SDAC). The data collected by the network will be filed in a large digital Mass Store facility at CCA.

The communication network used to interconnect the primary seismic stations, the SDAC processing facility, and the Mass Store will include a mix of dedicated leased circuits and ARPA Network circuits. The overall design of the seismic data network is described in [1]. The overall design of the ARPA Network is described in [2] and [3]. Protocols for interaction between the communication subnet and the "Host" installations are described in [4].

This document describes the formats and protocols for communication between the CCP and the Seismic Input Processor (SIP). Communications between the CCP and the SIP use the ARPA Network.

1.2 System Configuration and Restraints

The CCP will essentially be interfaced as a Host to both the TIP and the IMP but only one CCP interface will be in use at any time. This machine may, thus, have two possible Host addresses.

A second novel aspect of the seismic data communications use of the ARPA Network is the real-time fixed delay constraint. In particular, it is required that the communications links operate so that the data will be delivered to some destinations with constant delay behind real-time, i.e. the network should appear to be an "ideal (error free) delay line" to that destination.

The communication protocol described in this document builds upon the Host-IMP Protocol [4] and is at the Host-to-Host level although it is not the standard Host-Host Protocol described in [5]. The Host-IMP Protocol [4] is included in this specification by reference.

1.3 Organization of this Document

In section 2 the content and format of the data being exchanged over the communication system are specified.

These data must be embedded in messages that include routing and error control information. The message formats are described in section 3.

Finally, the operating protocols or rules for the exchange of data and control messages are specified in section 4.

2.0 Data Formats

2.1 Data Formats from the CCP to the SIP

Four forms of data will be sent from the CCP to the SIP. One form of data is the file structure parameters which describe the seismic data message field sizes. The second form of data is actual seismic sensor data. The CCP groups the sensor data into frames and then subdivides each frame into messages, each containing data from two source stations, for transmission to the SIP. The third form of data is sensor status change data for sensors whose data is being transmitted. The fourth form of data is operator messages to be typed on the SIP operator console.

The format of the <u>file structure</u> data is as follows:

Field 1: bytes 0 to 1: N = number of sites.

Field 2: bytes 2 to 13: First site parameters coded as follows:
Bytes 2 to 5: Site name plus 1 byte padding.
Bytes 6 and 7: Number of 16 bit words of SP status data.
Bytes 8 and 9: Number of 16 bit words of SP data.
Bytes 10 and 11: Number of 16 bit words of LP status data.
Bytes 12 and 13: Number of 16 bit words of LP data.

Fields 3 to N+1: Nth site parameters coded as for the first site.

The format for the <u>seismic data</u> blocks containing data from two stations for one second is as follows:

Field 1: bytes 0 and 1: number of words in the first site data block including fields 2 to 5 inclusive.

Field 2: bytes 2 to 9: Time-of-day:

16 BCD characters as follows:

char.	use
1	0
2 and 3	two digits of the year
4,5, and 6	three digits of day number
7,8	two digits of hours
9,10	two digits of minutes
11,12	two digits of seconds
13,14	two zeros for hundredths of seconds
15,16	padding-zeros

Field 3: bytes 10 to 12: site identity:
3 ASCII characters of site name

Field 4: byte 13: site status

Field 5: seismic data including
subfield 1: SP status padded to a word boundary
subfield 2: SP data: 10 or 20 samples (depending on sample rate) of each SP channel. Each sample is a 16 bit fixed point number. Samples are ordered in frames of 1/10 or 1/20 second each containing one sample from each seismometer.
subfield 3: LP status padded to a word boundary
subfield 4: LP data: 1 sample from each LP seismometer. Each sample is a 16 bit gain ranged number.

Fields 6 to 10: same as fields 1 to 5 but for the second site data.

The format of the <u>status</u> <u>change</u> data is as follows:

Field 1: bytes 0 to 5: Time-of-day;
Coded as the first 12 characters of time in the seismic data.

Field 2: bytes 6 and 7: Number of channels with changed status.

Fields 3 through N(number of channels): Channel id and status:
14 bytes coded as follows:
3 bytes ASCII characters for site id
If site id =000, complete status of entire network will follow this message.
1 byte ASCII character for channel type.
i = individual seismometer
s = subarray beam
b = array beam

a = adaptive beams
e = all channels from this site.

2 bytes ASCII characters for sample rate in samples per second.
4 bytes ASCII characters for channel id within the site.
1 byte ASCII character for gain code(H or L).
1 Byte of sensor component.
1 byte of sensor status bits.
1 byte offset location of sensor in site data.

The format for the operator message data is as follows:

Field 1: bytes 0 and 1: Character count = N.

Field 2; bytes 2 to N+1 or N+2: Text padded to a full word boundary.

2.2 Data Formats from the SIP to the CCP.

The two forms of data sent from the SIP to the CCP are the data filed list of data that has been passed from the SIP to the Datacomputer and operator messages to be typed on the CCP operator console.

The format of the list of filed data is as follows:

Field 1: bytes 0 to 3: Three bytes of site name and one byte of padding.

Field 2: bytes 4 to 13: File identifier: Ten ASCII character file name.

Field 3: bytes 14 to 19: Starting TOD of the filed data: Coded same as field 1 of the status change data.

Field 4: bytes 20 to 25: End TOD of the filed data: Coding same as field 1 of the status change data.

The format of the operator message data is described in section 2.1.

3. Message formats

3.1 CCP to SIP

The five types of messages that can be sent from the CCP to the SIP are the Structure Check messages (type 7), the Seismic Data message (type 0),the Status Message (type 8), the Acknowledge message (type 1). and the Operator message (type 5).

The message format for Structure Check messages is as follows:

Field 1: bytes 0 to 3: Host-to-IMP leader
(see [4])
bit 1; Priority bit = 0
bits 2-8: zeros
bits 9-10: destination Host number
bits 11-16: destination IMP number
bits 17-28: link number
bits 29-32: zeros

Field 2: byte 4: Source Host id.

Field 3: byte 5: message type = 7

Field 4: bytes 6 and 7: Unique message id.

Field 5: File structure parameter data (see section 2.1)

Field 6: 2 byte checksum for fields 2,3,4,and 5:
checksum defined as a 16 bit number
computed by subtracting the arithmetic sum of
the values of the words in the checked fields
from the number of words in the checked fields
using two's complement binary arithmetic.

The message format for the Seismic Data messages is as follows:

Field 1: bytes 0 to 3; Host-IMP Leader.
(same as Structure Check message)

Field 2: byte 4: Source Host id.

Field 3: byte 5: message type = 0.

Field 4: bytes 6 and 7: Unique message id.

Field 5: Seismic data (see section 2.1).

Field 6: 2 bytes checksum for fields 2,3,4, and 5:
Computed as for Structure Check Message.

The format for the Sensor Status message is as follows:

Field 1: bytes 0 to 3: Host-IMP leader:
(same as for Structure Check Message)

Field 2: byte 4: Source Host id.

Field 3: byte 5: message type = 8.

Field 4: bytes 6 and 7: Unique message id.

Field 5: sensor status (see section 2.1).

Field 6: checksum on fields 2,3,4 and 5:
Computed as for the Structure Check Message.

The format for the Acknowledge message is as follows:

Field 1: bytes 0 to 3: Host-IMP leader:
same coding as for Structure Check Message.

Field 2: byte 4: Source Host id.

Field 3: byte 5: Message type = 1.

Field 4: bytes 6 and 7: Unique id of message being acknowledged.

Field 5: bytes 8 and 9: Checksum on fields 2,3 and 4:
Computed as for the Structure Check Message.

The format of the Operator message is as follows:

Field 1: bytes 0 to 3: Host-to-IMP leader:
Coding same as for Structure Check message.

Field 2: byte 4: Source Host id.

Field 3: byte 5: Message type = 5.

Field 4: bytes 6 and 7: Unique message id.

Field 5: Operator message (see section 2.1).

Field 6: Two byte checksum on fields 2,3,4 and 5:
Computed as for Structure Check message.

3.2 SIP to CCP

The four types of messages sent from the SIP to the CCP include Acknowledge for messages received from the CCP (type 1), Host-Going-Down (type 4), Data Filed message (type 9), and Operator message (type 5).

The format of Acknowledge messages is described in section 3.1.

The format of the <u>Host-Going-Down</u> messages is as follows:

Field 1: bytes 0 to 3: Host-IMP leader:
Same format as for Structure Check messages
(see section 3.1).

Field 2: bytes 4 and 5: message type = 4

The format of the <u>Data Filed</u> message is as follows:

Field 1: bytes 0 to 3: Host-to-IMP leader:
Same coding as for Structure Check message.

Field 2: byte 4: Source Host id.

Field 3: byte 5: Message type = 9.

Field 4: bytes 6 and 7: Unique message id.

Field 5: bytes 8 to 33: Data filed list (see section 2.2).

Field 6: bytes 34 and 35: Checksum on fields 2,3,4 and 5.
Computed same as for Structure Check message.

The format of the <u>Operator</u> message is desribed in section 3.1.

4.0 Operating Procedures

In normal operation, the CCP will send 3 seismic data messages to the SIP each second. Each message contains one second of data from two seismic stations.

Whenever the status of a sensor whose data is being sent to the SIP is changed (manually by the CCP operator or automatically), a sensor Status Message is sent from the CCP to the SIP for that sensor. If the status of all sensors at a given site change at the same time, a special case of the Status Change message will be sent

Approximately once each hour a Structure Check message will be sent from the CCP to the SIP. If the file structure parameters for any site are inconsistent with the current file format, the SIP will stop filing data from that site and will notify the operating personnel.

At midnight G.M.T., at startup after any interruption of communication between the CCP and the SIP, or by operator command, a Structure Check message followed by a special Status Change message (indicating full network status will follow) followed by the full status of the entire network will be sent from the CCP to the SIP.

After recieving and acknowledging a Host-Going-Down message from the SIP, or anytime the CCP determines that communication with the SIP has been interrupted the CCP will send Structure Check messages to the SIP periodically. Receipt of an acknowledge for one of these messages will indicate that the system has returned to normal.

In the other direction the SIP will send a Data Filed message to the CCP for all data successfully passed to the Datacomputer.

Operator messages will be sent in either direction under command of the sending operator.

If the SIP is being taken off-line for maintenance or other predictable reason, a Host-Going-Down message will be sent to the CCP or an equivalent message will be entered by the CCP operator.

All messages except Acknowledge messages exchanged between the CCP and the SIP will be acknowledged at the Host-to-Host level if received with correct checksum.

All data messages will be saved in the source Host until the Acknowledge for that data is received. The source Host will retransmit unacknowledged messages approximately every four seconds until the buffer is deleted. If any Host is buffering ten seconds of unacknowledged messages the oldest buffers may be reused for new data and the operator will be notified.

Since the CCP may be on one of two Host connections to the network, the SIP will send messages to the Host address from which it received the latest Seismic Data message with a good checksum.

5. REFERENCES

[1] BBN Report No. 2632 - Final Report - A Study of the Data Collection, Processing, and Management for a Worldwide Seismic Network, September 1973.

[2] Karp, P.M., Origin, Development and Current Status of the ARPA Network, Digest of Papers, COMPCON 73, Seventh Annual IEEE Computer Society International Conference, 1973.

[3] Heart, F., Kahn, R., Ornstein, S.M., Crowther, W.R., and Walden, D.C., The Interface Message Processor for the ARPA Computer Network, Proc. AFIPS Spring Joint Computer Conference, 1970.

[4] BBN Report No. 1822, Specifications for the Interconnection of a Host and an IMP, April 1973.

[5] McKenzie, A., Host/Host Protocol for the ARPA Network, NIC 8246, January 1972.

From: H. Briscoe
To: Lt. M. Marcus, VSC
Subject: Communication Protocol between NORSAR and SDAC.

1. Introduction

1.1 Background

A worldwide seismic data collection network including approximately 6 on-line array stations is being implemented under ARPA sponsorship. The objective of the network is to provide data for research in nuclear test detection and identification. A network event list will be prepared within a couple of days of real-time using on-line event detection and computer aided event analysis at the Seismic Data Analysis Center (SDAC). The data collected by the network will be filed in a large digital Mass Store facility at CCA.

The communication network used to interconnect the primary seismic stations, the SDAC processing facility, and the Mass Store will include a mix of dedicated leased circuits and ARPA Network circuits. The overall design of the seismic data network is described in [1]. The overall design of the ARPA Network is described in [2] and [3]. Protocols for interaction between the communication subnet and the "Host" installations are described in [4].

This document describes the formats and protocols for communication with one of the primary sites, the Norwegian Seismic Array (NORSAR). Communication with NORSAR uses the ARPA Network. The NORSAR station communication is unique in that a) the station not only sends data but also receives data, and b) the data exchanged includes processed information. This document, therefore, describes both the communication between NORSAR and the CCP and the exchange of processed data between the CCP and the 360/40A Detection Processor (DP) at SDAC.

1.2 System Configuration and Restraints

The NORSAR Detection Processor (DP) will be interfaced to the NORSAR Terminal Interface Processor (TIP) as a "Host". The NORSAR TIP, the SDAC TIP, and the SDAC IMP are each nodes in the ARPA Network. The path between the NORSAR TIP and the SDAC TIP and IMP will include one hop via a communication satellite. The SDAC DP and the CCP will essentially be interfaced as Hosts to both the TIP and the IMP but only one CCP and one DP interface will be in use at any time. These machines may, thus, have two possible Host addresses.

A second novel aspect of the seismic data communications use of the ARPA Network is the real-time fixed delay constraint. In particular, it is required that the communications links operate so that the data will be delivered to some destinations with constant delay behind real-time, i.e. the network should appear to be an "ideal (error free) delay line" to that destination.

The communication protocol described in this document builds upon the Host-IMP Protocol [4] and is at the Host-to-Host level although it is not the standard Host-Host Protocol described in [5]. In spite of the fact that there is considerable variation among the array sites with respect to data format, rate and type of ARPANET access, much of the protocol described below is similar to the protocol for sites other than NORSAR.

1.3 Organization of this Document

In section 2 the content and format of the seismic data being exchanged over the communication system are specified.

These data must be embedded in messages that include routing and error control information. The message formats are described in section 3.

Finally the operating protocols or rules for the exchange of data and control messages are specified in section 4.

2. Data Formats

2.1 CCP to and from NORSAR.

Each second one frame of data will be assembled at the CCP for transmission to NORSAR and one frame of data will be assembled at NORSAR for transmission to the CCP.

2.1.1 CCP to NORSAR

Each frame of data from the CCP to NORSAR will have the following data:

Field 1:	bytes 0 to 3: control characters: character sequence SYN-SYN-DLE-STX.
Field 2:	bytes 4 to 17: LASA Signal Arrival Queue file entry: data from field 1 of the last frame of processed data from the DP (see section 2.2)
Field 3:	bytes 18 to 21: LASA Time-of-Day (TOD): using ISRSPS format
Field 4:	bytes 22 to 28: LASA LP Status and Repeat bits: first 30 bits assigned to 30 components of LP data in the same order as LDC to CCP message. Bits set to 1 when corresponding sensor is down. Remainder of bytes 22 to 27 are zeros. Bit 5 of byte 28 is on if no LASA LP data are present. Byte 28 bit 6 is on if any polycode errors were detected in transmission from LDC to CCP. Byte 28 bit 7 is on if any LP data within the frame are repeated data.

Field 5: bytes 29 to 130: LASA LP data:
LASA long period data values ordered by component and subarray as they are transmitted from LDC to CCP.

Field 6: bytes 131 to 134: ALPA Time-of-day (TOD):
coded same as LASA TOD field 3.

Field 7: bytes 135 to 142: ALPA LP status and repeat indicator:
ALPA LP sensor status arranged by site, 3 bits per site.
For each subarray:

Bit	Sensor	
1	LV	
2	LN	on when sensor down
3	LE	

Byte 142 bit 5 on if no ALPA data present. Byte 142 bit 6 on if any polycode errors were present in transmission of data from ALPA to CCP. Byte 142 bit 7 on if any LP data in the frame are repeat data. Unused bits are zero.

Field 8: bytes 143 to 256: ALPA LP data:
Data values ordered by site and component for a total of 57 values in 16 bit gain ranged format.

Field 9: bytes 257 to 282: zeros

Field 10a: bytes 283 to 294: Text from CCP operator:
twelve bytes of operator message from the CCP operator to NDPC. Text is broken into messages up to 128 bytes long and messages formatted as follows.

Byte	Description
0	Start of message characters - must equal X´FF´
1	Count of NDPC to SAAC messages required to complete the message
2-5	Time of day, in ISRSPS format
6-10	Message identification (S70CA) or (N700*)
11	Message control field
	Bits
	0 — Set to one to indicate last message or single message
	1 — Set to zero to not ring alarm bell
	2-7 — Spare - encoded as zeros.
12-N	Message text

Field 10b: bytes 283 to 291: SP request data:
This is a special format of operator message that may replace field 10a.
Format for this field is as follows:

Byte	Description
0	Data Request Message Identifier set equal to X´FE´
1	Number of short period channels requested (0-3)

2	Number of subarray beam/array beam values requested (0-3)
3-8	Channel/Beam numbers (2 bytes each)

Field 11: bytes 295 to 298: Control characters:
character sequence DLE,ETB,0,0

Field 12: byte 299: spare-coded as zeros:

2.1.2 NORSAR to CCP.

Each frame from NORSAR to the CCP will have the following data:

Field 1: bytes 0 to 3: Control Characters:
character sequence SYN, SYN, DLE, STX.

Field 2: bytes 4 to 9: Time of day:
12 BCD characters as follows:

character	use
1	0
2&3	two digits of the year
4,5,&6	three digit of day number
7,8	two digits of hours
9,10	two digits of minutes
11,12	two digits of seconds

Field 3: bytes 10 to 27: NORSAR detection log reduction groups: processed data to be transmitted to the SDAC DP in the next frame (see section 2.2)

Field 4: bytes 28 to 37: LP status and repeat indicators:
This field contains NORSAR LP sensor status arranged continuously by subarray 3 bits per subarray. For each subarray:

Bit	Sensor	
1	LV	
2	LN	set when sensor is down
3	LE	

Byte 36, bits 2 through 7, and byte 37, bits 0 through 5, are spare, encoded zero.
Byte 37, bits 6 and 7 are repeat indicators which are set whenever any SP or LP data within the record is repeated data.

**

When set (1) the bits indicate the following:
bit 6 - SP data is repeated (at least .5 second of data is repeated)
bit 7 - LP data is repeated (at least 1 value is repeated).

Field 5: bytes 38 to 169: NORSAR LP data:
This field contains long period data values ordered by component and subarray for a total of 66 long period values. Each value is in 16 bit gain ranged form. The components are ordered LV, LN, and LE. The 22 long period values (one for each subarray) for component LV are recorded in the field first, followed by the 22 for LN, followed by the 22 for LE.

Field 6: bytes 170 to 175: NORSAR SP channel identification: Three two-byte entries containing the type and channel or beam number of the short period data in field 7. The data may be short period channel, subarray beam or array beam data, as selected by the SDAC operator. Each entry has the following format:

bits	Content
0-3	Type of data identifier
4-15	Channel or beam number of data

If no short period data has been requested by SDAC, this field will be all zeros.

Field 7: bytes 176 to 235: NORSAR SP data: This field contains one second´s worth of short period data from each of the channels and/or beams identified in Field 6; the data for each channel and beam is recorded at a 10 Hz rate. The field is divided into ten subfields, each containing a decisecond´s worth of data. Each subfield is six bytes in length and contains one two-byte data value for each channel and/or beam identified in Field 6. If only one or two channels and/or beams are being transmitted, the latter two-byte entries of each subfield will contain zeros (the latter two if one channel, the latter one if two channels). If no channels and/or beams are being transmitted, this entire field will contain all one bits.

Field 8: bytes 236 to 283: NORSAR off-line results: processed data to be transmitted to SDAC DP in the next frame (see section 2.2).

Field 9: bytes 284 to 291: Operator text messages: Operator messages from NORSAR to the CCP operator. Maximum message length of 128 bytes. A twelve byte header is added and the resulting message is broken into 8 byte groups for transmission in successive frames.

Message format is as follows:

Byte	Description	
0	Start of message characters - must equal X´FF´	
1	Count of NDPC to SDAC messages required to complete the message	
2-5	Time-of-day, in ISRSPS format	
6-10	Message identification (N700A) or (N737A)	
11	Message control field	
	Bits	
	0	Set to one to indicate last message or single message
	1	Set to zero to not ring alarm bell
	2	Set to one to indicata NDPC "A" system
	3-7	Spare - encoded as zeros
12-N	Message text	

Field 10: Bytes 292 to 295: control characters:
character sequence DLE, ETB, 0,0
Field 11: bytes 296 to 299: spares coded as zeros.

2.2 CCP to and from DP

Whenever processed data is available in the DP, the data will be assembled in the DP for transmission to the CCP to be forwarded to NORSAR. Each second the processed data received from NORSAR during the previous second will be assembled in the CCP and passed to the DP.

2.2.1 DP to CCP

Each frame of data from the DP to the CCP will have the following data:

Field 1: bytes 0 to 5: Time-of-day:
12 BCD characters as follows:

char.	use
1	0
2&3	two digits of the year
4,5,&6	three digit of day number
7,8	two digits of hours
9,10	two digits of minutes
11,12	two digits of seconds

Field 2: bytes 6 to 19: LASA Signal Arrival Queue File Entry: data for field 2 of a CCP to NORSAR message (see section 2.1)

2.2.2 CCP to DP

Each frame of processed data from the CCP to the SDAC DP will have the following data:

Field 1: bytes 0 to 5: time-of-the-day: as in field 1 above
Field 2: bytes 6 to 23: NORSAR detection log reduction groups:
data from field 3 of the last NORSAR to CCP message.
Field 3: bytes 24 to 71: NORSAR off-line results:
data from field 8 of the last NORSAR to CCP message.

3. Message Formats

3.1 CCP/NORSAR

The only messages exchanged between the CCP and NORSAR are data messages as follows:

Field 1: bytes 0 to 3: Host-to-IMP leader (see [4])
 bit 1: Priority bit = 0
 bits 2-8: zeros
 bits 9-10: destination host number
 bits 11-16: destination IMP number
 bits 17-28: link number
 bits 29-32: zeros
Field 2: bytes 4 and 5 message type: 0 for data from NORSAR to CCP. 11 for data from CCP to NORSAR
Field 3: bytes 6 to 305: Data: As described in section 2.1
Field 4: bytes 306-307: checksum for fields 2 and 3 = number of words in checksum minus the arithmetic sum of the word values using binary two's complement arithmetic.

3.2 CCP/DP

Two classes of message are used between the CCP and the DP for the exchange of processed data. A unique data message class is used to separate these messages from the sensor data messages. Messages are acknowledged using the standard acknowledge message.

The data messages have the following format:

Field 1: bytes 0 to 3. Host-to-IMP leader (see [4]) coded as in section 3.1
Field 2: byte 4: source Host id:
Field 3: byte 5: message type - 10
Field 4: bytes 6 and 7: Unique message id
Field 5: data field described in section 2.2
Field 6: 2 bytes; checksum on fields 2,3,4 and 5 (see section 3.1 for coding)

The format of the standard acknowledgment message is as follows:

Field 1: bytes 0-3: Host-to-IMP leader coded as in section 3.1
Field 2: byte 4: source Host id:
Field 3: byte 5: message type =1
Field 4: bytes 6 to 7: Unique id of message being acknowledged
Field 5: bytes 8 and 9: checksum on fields 2,3,and 4 (see section 3.1 for coding)

4. Operating Procedures

4.1 CCP/NORSAR

In normal operation of the link between the CCP and the NORSAR, both Hosts transmit one data message each second. No attempt will be made to retransmit lost or garbled messges. If data is lost it can be replaced off-line.

The only complication in the operating procedures is the result of the fact that the CCP will have two network addresses on different IMPs and will use only one at a time. The NORSAR DP must, therefore, monitor the source address on incoming data messages in order to determine the correct address for the outgoing data.

If the NORSAR sends a data message to the wrong CCP address and the addressed Host interface is looped for testing, the IMP subnet will cause the message to be sent back to NORSAR as a message from the erroneous CCP address. The NORSAR program must be prepared to ignore any type 0 messages it receives.

4.2 CCP/DP

The exchange of processed data between the CCP and the DP at SDAC will use a positive acknowledgment procedure in both directions.

Whenever the CCP has processed data from NORSAR available it will attempt to send a processed data message to the DP. All messages transmitted to the DP will be saved in the CCP until an acknowledge message for that data message is received or the message is at least ten seconds old. Unacknowledged messages will be retransmitted approximately every four seconds until the buffer is overwritten. When a message is over ten seconds old the Host may delete the space and reuse the buffer space. The operator will be notified whenever unacknowledged messages are deleted.

Each time the CCP receives a processed data message from the DP with a correct checksum an acknowledge message for that data will be sent to the DP.

Whenever processed data is available, the DP will attempt to transmit a processed data message to the CCP. All messages transmitted to the CCP will be saved in the DP until an acknowledge message for that data message is received. If no acknowledge is received after 5 seconds the data message will be retransmitted.

Each time the DP receives a processed data message from the CCP with a correct checksum an acknowledge message for that data will be sent to the CCP.

Control of the link between the CCP and the DP is described in the specification of the protocol for communication between the CCP and the DP.

5. References

[1] BBN Report No. 2632 - Final Report - A Study of the Data Collection, Processing, and Management for a Worldwide Seismic Network, September 1973.

[2] Karp, P.M., Origin, Development and Current Status of the ARPA Network, Digest of Papers, COMPCON 73, Seventh Annual IEEE Computer Society International Conference, 1973.

[3] Heart, F., Kahn, R., Ornstein, S.M., Crowther, W.R., and Walden, D.C., The Interface Message Processor for the ARPA Computer Network, Proceedings. AFIPS Spring Joint Computer Conference, 1970.

[4] BBN Report No. 1822, Specifications for the Interconnection of a Host and an IMP, April 1973.

[5] McKenzie, A., Host/Host Protocol for the ARPA Network, NIC 8246, January 1972.

From: H. Briscoe
To: Lt. M. Marcus, VSC
Subject: Communication Protocol between the CCP and the DP.

1. Introduction

1.1 Background

A worldwide seismic data collection network including approximately 6 on-line array stations is being implemented under ARPA sponsorship. The objective of the network is to provide data for research in nuclear test detection and identification. A network event list will be prepared within a couple of days of real-time using on-line event detection and computer aided event analysis at the Seismic Data Analysis Center (SDAC). The data collected by the network will be filed in a large digital Mass Store facility at CCA.

The communication network used to interconnect the primary seismic stations, the SDAC processing facility, and the Mass Store will include a mix of dedicated leased circuits and ARPA Network circuits. The overall design of the seismic data network is described in [1]. The overall design of the ARPA Network is described in [2] and [3]. Protocols for interaction between the communication subnet and the "Host" installations are described in [4].

This document describes the formats and protocols for communication of sensor data between the CCP and the 360/40 acting as the Detection Processor (DP). Communication between the CCP and the DP use the ARPA Network. Communication of the processed data fields for the interaction with NORSAR also involves a CCP/DP interaction that is described in the document on communication protocol beteen NORSAR and SDAC.

1.2 System Configuration and Restraints

The SDAC DP and the CCP will essentially be interfaced as Hosts to both the TIP and the IMP but only one CCP and one DP interface will be in use at any time. These machines may, thus, have two possible Host addresses.

A second novel aspect of the seismic data communications use of the ARPA Network is the real-time fixed delay constraint. In particular, it is required that the communications links operate so that the data will be delivered to some destinations with constant delay behind real-time, i.e. the network should appear to be an "ideal (error free) delay line" to that destination.

The communication protocol described in this document builds upon the Host-IMP Protocol [4] and is at the Host-to-Host level although it is not the standard Host-Host Protocol described in [5]. In spite of the fact that there is considerable variation among the seismic communication paths with respect to data format rate and type of ARPANET access, much of the protocol described below is similar to the protocol for communication with the stations and for communication with the SIP.

1.3 Organization of this Document

In section 2 the content and format of the data being exchanged over the communication system are specified.

These data must be embedded in messages that include routing and error control information. The message formats are described in section 3.

Finally the operating protocols or rules for the exchange of data and control messages are specified in section 4.

2. Data formats

2.1 Data Formats from the CCP to the DP

Four forms of data will be sent from the CCP to the DP. One form of data is the processed data from NORSAR. Processed data formats are described in the Communication Protocol between NORSAR and SDAC. A second form of data is the file structure parameters which describe the seismic data message field sizes. The third form of data is actual seismic sensor data. The CCP groups the sensor data into one second frames and then subdivides each frame into messages each containing data from two source stations for transmission to the DP. The fourth form of data is sensor status change data for sensors whose data is being transmitted.

The format of the file structure data is as follows:

Field 1: bytes and 1: N=number of sites.

Field 2: bytes 2 to 13: First site parameters coded as follows:
    Bytes 2 to 5:   Site name.
    Bytes 6 and 7: Number of 16 bit words of SP status data.
    Bytes 8 and 9: Number of 16 bit words of SP data.
    Bytes 10 and 11: Number of 16 bit words of LP status data.
    Bytes 12 and 13; Number of 16 bit words of LP data.

Fields 3 to N+1: Nth site parameters coded as for the first site.

The format for the seismic data blocks containing data from two stations for one second is as follows:

Field 1: bytes 0 and 1: number of words in the first site data block including fields 2 to 5 inclusive.

Field 2: bytes 2 to 9: Time-of-day:

16 BCD characters as follows:

char.	use
1	0
2&3	two digits of the year
4,5,&6	three digits of day number
7,8	two digits of hours
9,10	two digits of minutes
11,12	to digits of seconds
13,14	two zeros for hundredths of seconds
15,16	padding-zeros

Field 3: bytes 10 to 12: site identity:
3 ASCII characters of site name

Field 4: byte 13: site status

Field 5: seismic data including
subfield 1: SP status padded to a word boundary

subfield 2: SP data: 10 or 20 samples (depending on sample rate) of each SP channel. Each sample is a 16 bit fixed point number. Samples are ordered in frames of 1/10 or 1/20 second each containing one sample from each seismometer.

subfield 3: LP status padded to a word boundary

subfield 4: LP data: 1 sample from each LP seismometer. Each sample is a 16 bit gain ranged number.

Fields 6 to 10: same as fields 1 to 5 but for the second site data.

The format of the status change data is as follows:
Field 1: bytes 0 to 5: Time-of-day;
Coded as first 6 bytes of time in the seismic data message.

Field 2: bytes 6 and 7: Number of channels with changed status.

Fields 3 through N(number of channels): Channel id and status:
14 bytes coded as follows:
- 3 bytes ASCII characters for site id
  If site id = 000, complete staatus of the entire network will follow this message.
- 1 byte ASCII character for channel type
  - i = individual seismometer
  - s = subarray beam
  - b = array beam
  - a = adaptive beam
  - e = all channels from this site
- 2 bytes ASCII characters for sample rate in samples per second.
- 4 bytes ASCII characters for channel id within the site.
- 1 byte ASCII character for gain code(H or L).
- 1 byte of sensor component.
- 1 byte of sensor status bits.
- 1 byte offset location of sensor in site data.

2.2 Data Formats from the DP to the CCP.

The only data sent from the DP to the CCP is processed data for NORSAR described in the Communication Protocol between NORSAR and SDAC.

3. Message formats

3.1 CCP to DP

The five types of messages that can be sent from the CCP to the DP are the NORSAR Processed Data messages (type 10), the Structure Check messages (type 7), the Seismic Data message (type 0), the Status Message (type 8),and the Acknowledge message (type 1).

Type 10 messages are described in the Communication Protocol between NORSAR and SDAC.

The message format for Structure Check messages is as follows:

Field 1: bytes 0 to 3: Host-to-IMP leader
(see [4])
bit 1; Priority bit = 0
bits 2-8: zeros
bits 9-10: destination Host number
bits 11-16: destination IMP number
bits 17-28: link number
bits 29-32: zeros

Field 2. byte 4: Source Host id.

Field 3: byte 5: message type = 7

Field : bytes 6 and 7: Unique message id.

Field 5: File structure parameter data (see section 2.1)

Field 6: 2 byte checksum for fields 2,3,4,and 5: checksum defined as a 16 bit number computed by subtracting the arithmetic sum of the values of the 16 bit words in the checked fields from the number of 16 bit words in the checked fields using two's complement binary arithmetic.

The message format for the Seismic Data messages is as follows:

Field 1: bytes 0 to 3; Host-IMP Leader.
(same as Structure Check message)

Field 2: byte 4: Source Host id.

Field 3: byte 5: message type = 0.

Field 4: bytes 6 and 7: Unique message id.

Field 5: Seismic data (see section 2.1).

Field 6: 2 bytes checksum for fields 2,3,4, and 5:
Computed as for Structure Check Message.

The format for the Sensor Status message is as follows:

Field 1: bytes 0 to 3: Host-IMP leader:
(same as for Structure Check Message)

Field 2: byte 4: Source Host id.

Field 3: byte 5: message type = 8.

Field 4: bytes 6 and 7: Unique message id.

Field 5: sensor status (see section 2.1).

Field 6: checksum on fields 2,3,4 and 5:
Computed as for the Structure Check Message.

The format for the Acknowledge message is as follows:

Field 1: bytes 0 to 3: Host-IMP leader:
same coding as for Structure Check Message.

Field 2: byte 4: Source Host id.

Field 3: byte 5: Message type = 1.

field 4: bytes 6 and 7: Unique id of message being acknowledged.

Field 5: bytes 8 and 9: Checksum on fields 2,3 and 4:
Computed as for the Structure Check Message.

3.2 DP to CCP

The three types of messages sent from the DP to the CCP include processed data for NORSAR (type 10), acknowledge for messages received from the CCP (type 1), and Host-Going-Down (type 4).

Processed Data messages are described in the Communication Protocol beteen NORSAR and SDAC.

The format of Acknowledge messages is described in section 3.1.

The format of the host-Going-Down messages is as follows:

Field 1: bytes 0 to 3: Host-IMP leader:
Same format as for Structure Check messages (see section 3.1).

Field 2: byte 4· may contain source Host id.
Field 3: byte 5: message type = 4

4. Operating Procedures

In order to clarify the range of possible operating procedures for communication from the CCP to the DP, three classes of information available from the CCP are defined. These include NORSAR processed data, seismic data, and auxiliary data consisting of Status messages and Structure Check messages (see Communicatin Protocol Between the CCP and the SIP). Seismic data is subdivided into long period and short period data from each reporting station. In normal operation the CCP will transmit one NORSAR processed data message to the DP each second if processed data is being recieved from NORSAR. In normal opreration the CCP will send one frame of seismic data to the DP each second. A frame of seismic data may reequire from one to three messages each containing data from up to two stations. The seismic data transmitted is under control of the CCP operator. He may designate data type (long period, short period, or both) to be transmitted for each station. Finally the CCP operator may command that auxiliary data will be transmitted from the CCP to the DP whenever auxiliary data messages are generated.

If the CCP determines that the mass store is not accepting data or if the CCP operator initiates backup operation the CCP will replace the previously defined DP seismic and auxiliary data set with the data being transmitted to the mass store. The mass store data stream is the same as long and short period data from all stations and auxiliary data. Whenever possible, data will be transmitted in chronologic order.

The DP will send an acknowledge message to the CCP for each data message in any of the above classes recieved with a correct checksum.

In the other direction the DP will send processed data messages to the CCP for transmission to NORSAR whenever there are processed data to be sent. The CCP will acknowledge each processed data message from the DP that is received with the correct checksum.

All of the data messages described above will be saved in the source Host until either an acknowledge for the message is received from the receiving Host or the message is ten seconds or more old. The source Host will retransmit unacknowledged messages approximately every four seconds until the buffers are deleted. When a buffered message is more than ten seconds old, the source Host may delete the message and reuse the buffer space. The operator at the source Host will be notified whenever unacknowledged messages are deleted.

Since both the DP and the CCP may be on one of two Host connections to the network, the correct Host address for sending messages will be determined as follows:

a) the DP will send messages to the Host address from which it received the latest Seismic Data message.

b) the CCP will send messages to the Host address entered by the CCP operator.

If the DP is being taken off-line for maintenance or other predictable reason, a Host-Going-Down message will be sent to the CCP or an equivalent message will be entered by the CCP operator.

5. References

[1] BBN Report No. 2632 - Final Report - A Study of the Data Collection, Processing, and Management for a Worldwide Seismic Network, September 1973.

[2] Karp, P.M., Origin, Development and Current Status of the ARPA Network, Digest of Papers, COMPCON 73, Seventh Annual IEEE Computer Society International Conference, 1973.

[3] Heart, F., Kahn, R., Ornstein, S.M., Crowther, W.R., and Walden, D.C., The Interface Message Processor for the ARPA Computer Network, Proceedings. AFIPS Spring Joint Computer Conference, 1970.

[4] BBN Report No. 1822, Specifications for the Interconnection of a Host and an IMP, April 1973.

[5] McKenzie, A., Host/Host Protocol for the ARPA Network, NIC 8246, January 1972.

From: H. Briscoe
To: Lt. M. Marcus, VSC
Subject: Communication Protocol between the CCP and the 360/44.

1. Introduction

1.1 Background

A worldwide seismic data collection network including approximately 6 on-line array stations is being implemented under ARPA sponsorship. The objective of the network is to provide data for research in nuclear test detection and identification. A network event list will be prepared within a couple of days of real-time using on-line event detection and computer aided event analysis at the Seismic Data Analysis Center (SDAC). The data collected by the network will be filed in a large digital Mass Store facility at CCA.

The communication network used to interconnect the primary seismic stations, the SDAC processing facility, and the Mass Store will include a mix of dedicated leased circuits and ARPA Network circuits. The overall design of the seismic data network is described in [1]. The overall design of the ARPA Network is described in [2] and [3]. Protocols for interaction between the communication subnet and the "Host" installations are described in [4].

In order to allow the entry of data from sources not in the on-line network the CCP will accept data in VELA block data form from the 360/44 at SDAC. Data entered in this fashion will be placed in the block data section of the messages on the VELA-3 circuit. This document describes the formats and protocols for the communication of block data from the 360/44 to the CCP.

1.2 System Configuration and Restraints

The CCP will essentially be interfaced as a Host to both the TIP and the IMP but only one CCP interface will be in use at any time. This machine may, thus, have two possible Host addresses. The 360/44 will be a Host on the network.

A second novel aspect of the seismic data communications use of the ARPA Network is the real-time fixed delay constraint. In particular, it is required that the communications links operate so that the data will be delivered to some destinations with constant delay behind real-time, i.e. the network should appear to be an "ideal (error free) delay line" to that destination.

The communication protocol described in this document builds upon the Host-IMP Protocol [4] and is at the Host-to-Host level although it is not the standard Host-Host Protocol described in [5]. In spite of the fact that there is considerable variation among the seismic communication paths with respect to data format rate and type of ARPANET access, much of the protocol described below is similar to the protocol for communication with the stations and for communication with the SIP and DP.

1.3 Organization of this Document

In section 2 the content and format of the data being exchanged over the communication system are specified.
These data must be embedded in messages that include routing and error control information. The message formats are described in section 3.
Finally the operating protocols or rules for the exchange of data and control messages are specified in section 4.

2. Data formats from the 360/44 to the CCP.

The only data sent from the 360/44 to the CCP is block data for the VELA-3 circuit. The format of the block data is as follows:

Field 1: bytes 0 and 1: Number of bytes in field 3

Field 2: bytes 2 and 3: Id and status word for the VELA-3 header

Field 3: VELA-3 block data for one second.

3. Message Formats

3.1 360/44 to CCP

The message format for the VELA-3 block data message is as follows:

Field 1: bytes 0 to 3: Host-IMP leader
(see [4])

bit 1; priority bit = 0
bits 2-8; zeros
bits 9-10; destination Host number
bits 11-16; destination IMP number
bits 17-28; link number
bits 29-32; zeros

Field 2: byte 4; Source Host id

Field 3: byte 5: Message type = 12

Field 4: bytes 6 and 7: Unique message id

Field 5: VELA-3 block data (see section 2.0)

Field 6: 2 byte checksum for fields 2,3 4 and 5: Checksum defined as a 16 bit number computed by subtracting the arithmetic sum of the values of the 16 bit words in the checked fields from the number of 16 bit words in the checked fields using two's complement binary arithmetic.

3.2 CCP to 360/44

The only messages sent from the CCP to the 360/44 are acknowledge messages with the following format:

Field 1: bytes 0 to 3: Host-IMP leader: (same format as for VELA-3 block data message in section 3.1)

Field 2: byte 4: Source Host id .

Field 3: byte 5: message type =1

Field 4: bytes 6 and 7: Unique id of message being acknowledged

Field 5: bytes 8 and 9: Checksum on fields 2,3,and 4 (computed as for VELA-3 block data message described in section 3.1)

4. Operating Procedures

Whenever there is a new sequence of block data to be transmitted , the 360/44 begins the transmission by sending the first VELA-3 block data message to the CCP. Thereafter, the 360/44 will send the next message in the sequence whenever an acknowledge is received for the previous message. If no acknowledge is received after approximately 2 seconds the previous message is repeated.

Whenever the CCP receives a VELA-3 block data message with the correct checksum and a new message id, the block length field will be compared with the available space for block data in the VELA-3 format. If the block will fit, the message is placed in a queue. When the queue length is less than 4 messages, the CCP returns an acknowledge message to the 360/44. Thus the acknowledgment provides flow control.

If the block data will not fit in the available space on the VELA-3 line, the operator will be notified and the message will not be acknowledged.

Since the CCP may be on either of two Host ports, the 360/44 must determine the correct CCP address. This may be done by operator input or the 360/44 may send the first message of a block data sequence to both possible CCP addresses and send the rest of the data to the address from which it receives an acknowledgment.

5. References

[1] BBN Report No. 2632 - Final Report - A Study of the Data Collection, Processing, and Management for a Worldwide Seismic Network, September 1973.

[2] Karp, P.M., Origin, Development and Current Status of the ARPA Network, Digest of Papers, COMPCON 73, Seventh Annual IEEE Computer Society International Conference, 1973.

[3] Heart, F., Kahn, R., Ornstein, S.M., Crowther, W.R., and Walden, D.C., The Interface Message Processor for the ARPA Computer Network, Proceedings. AFIPS Spring Joint Computer Conference, 1970.

[4] BBN Report No. 1822, Specifications for the Interconnection of a Host and an IMP, April 1973.

[5] McKenzie, A., Host/Host Protocol for the ARPA Network, NIC 8246, January 1972.

APPENDIX B

USE OF THE ARPA NETWORK
by the
SEISMIC DATA COLLECTION NETWORK

BOLT BERANEK AND NEWMAN INC

CONSULTING · DEVELOPMENT · RESEARCH

Report No. 2995

USE OF THE ARPA NETWORK BY THE SEISMIC DATA COLLECTION NETWORK

30 January 1975

Submitted to:

VELA Seismological Center
312 Montgomery Street
Alexandria Virginia 22314

Sponsored by the Advanced Research Projects Agency
ARPA Order No. 2551

55<

USE OF THE ARPA NETWORK BY THE SEISMIC DATA COLLECTION NETWORK

30 January 1975

PRINCIPAL INVESTIGATOR: Howard W. Briscoe
Telephone (617) 491-1850 Ext. 415

This research was supported by the Advanced Research Projects Agency of the Department of Defense and was monitored by AFTAC/VSC, Patrick AFB FL 32925, under Contract No. F08606-75-C-0027

AFTAC Project Authorization No. VELA T/5701/BETR
Program Code No. 5F10
Effective Date: 1 August 1974
Total Amount of Contract: $70,000.00
Expiration Date: 30 June 1975

Prepared by:

Computer Systems Division
Bolt Beranek and Newman Inc.
Cambridge, Massachusetts 02138

TABLE OF CONTENTS

SUMMARY

Summarized below are the results of a preliminary analysis, presented at length in this interim report, of the impact of the seismic network on the ARPA Network. During the course of this study, no insurmountable problems were uncovered which would discourage the use of the ARPA Network by the seismic network. The ARPA Network provides a flexible, reliable data communication medium which the seismic network can use to advantage immediately, and which is readily capable of being adapted to support future seismic network growth.

The impact of the seismic data on the ARPA Network will be significant. Rough estimates indicate that with the completion of the seismic network, the average number of packets/day will double. Consequently, some growth of the ARPA Network will be necessary to support both the seismic application and the other users. In particular, adding the seismic traffic will overload the set of current ARPA Network lines in the region between SDAC and CCA, thus requiring some increase in the capacity of those lines or modification of topology. Also, as presently structured, the ARPA Network IMPs at SDAC and CCA cannot support the throughput necessary for the seismic traffic because of buffer space limitations; some modification to the affected IMPs will therefore be needed to achieve the throughput desired.

This study also shows that the delay encountered by the on-line seismic data from the station controllers in the field to the CCP will be well within the ten-second requirement.

Summary of Expected Delays

Link (refer to fig. 1)	Delay (sec)
a: to SPLI	1 sec
b: SPLI-IMP	0.14 *
c: Satellite	0.27-0.32 **
d+e: ARPA Network	0.3-0.5 ***
f: IMP-CCP	0.71-0.85 ****
CCP Processing	0.1
Total	2.67-3.06 sec

*Assuming a 7200-bit/sec line. An additional 0.4 second delay will occur if a packet must be retransmitted.

**Assuming a reservation protocol. An additional 0.56 to 0.71 second delay will occur if a packet must be retransmitted.

***Marginal lines and heavy traffic can cause additional delays of the order of tenths of seconds, and in extremely rare cases as high as 5 seconds.

****Includes delay of waiting for remaining packets in a multi-packet message.

Summary of Problem Areas

Problems:

(a) Limited amount of message reassembly buffer space the CCA-TIP and at SDAC.

(b) Congestion of lines from SDAC to CCA.

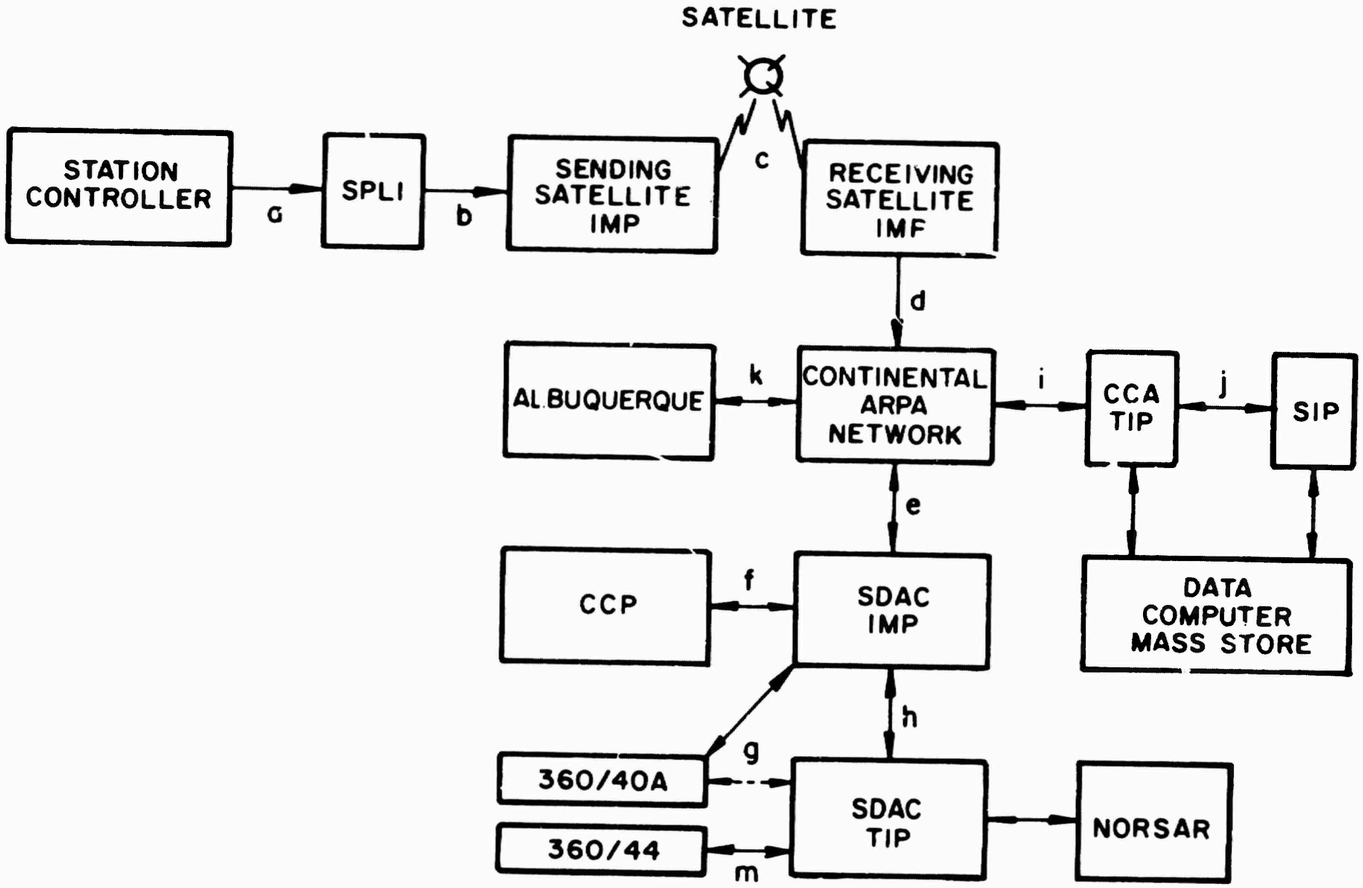

FIGURE 1 SEISMIC NETWORK LOGICAL MAP

Areas of Concern:

(c) The effect of the hostile environment on the VDH interface error rate (between the SPLI and Satellite IMP).

(d) The ability of the network routing algorithm to adapt quickly enough after a failure (line or IMP outage) to prevent excessive delays from being encountered by the on-line seismic data.

(e) Throughout this report, it has been assumed that satellite communication channels of the appropriate bandwidth will exist for use by the seismic network. Obviously, if such channels do not come into existence, this will create a major problem for the seismic network.

Summary of Possible Solutions to Problems

(a) Limited Buffer Space--This problem will begin to be noticeable at SDAC when the CCP and the 360/40A are on the same IMP and the 360/44 and NORSAR are both putting data into the system, possibly as early as the fall of 1975. It will become critical when the first SPLI comes on-line, possibly as early as late 1975 but more likely in mid-1976. The problem will appear at CCA when non-seismic inputs to the Datacomputer become significant. We see three possible solutions to this problem:

1. Replace the affected IMPs by Pluribus IMPs. This allows additional memory and processing capability to be added modularly when needed.
2. Install a dedicated IMP for the SIP, a brute-force but straightforward way to add the necessary additional memory.
3. Add memory plus the necessary software changes to allow existing IMPs to use the additional memory.

Our recommendations are as follows: For SDAC, solution 1. is the cleanest. Solution 3. is the next alternative, but a second IMP must be provided to allow the CCP an alternate means of connecting to the ARPA Network. The SDAC-TIP is not able to provide adequate reassembly buffer space. For CCA, either solution 1. or solution 2. will suffice.

(b) Limited Channel Capacity--This problem will appear gradually as the seismic load increases, probably late 1975 to mid-1976 depending on when the SIP comes on-line. Possible solutions:

1. Replace affected lines by 230.4 Kbit/sec lines. (20%)*
2. Install a direct 50 Kbit/sec line from SDAC to CCA via terrestrial or domestic satellite circuits. (60%)
3. Establish a nearly direct line by installing a Harvard-CCA line. (80%).
4. Implement bandwidth routing** with no topology changes. (70% on SDAC-Belvoir line)
5. Implement bandwidth routing with two independent paths. (50-60% on each path)

We recommend a combination of the direct line and bandwidth routing approaches to provide reliability and adaptability. For example, bandwidth routing could be implemented as the first parts of the seismic network come on-line (early 1976), with a nearly direct line added later when needed. Finally, as the full seismic network is completed, the nearly direct line could be replaced by a direct satellite or terrestrial circuit.

---

* Figures in parentheses give percentage of available data bandwidth used at tightest bottleneck. SRO-LPE data is not included.

** Bandwidth routing allows traffic to be routed simultaneously on multiple paths to achieve higher throughput than any line can support singly.

1. INTRODUCTION

1.1 Scope and Purpose

This interim report contains a preliminary analysis of the expected data flow paths in the proposed seismic network. The purpose of this report is to outline potential problem areas in the use of the ARPA Network to convey seismic data from the field to a centralized mass store, and to suggest possible ways in which the ARPA Network should evolve in order to satisfy the needs of the seismic application.

1.2 Organization of the Report

This report is divided into three sections. The first, besides introducing the body of the report, gives a very brief description of the ARPA Network which is intended to augment the reader's familiarity with that network and to provide a framework for the preliminary analysis.

In the second section, the preliminary analysis of the data flow paths is presented. Section 2.1 contains an overview of the various data flow paths. Sections 2.2 through 2.6 describe the various parts of the data path from the station processor to the seismic Communications and Control Processor (CCP). Delay is the primary concern for this data path, since the data must reach the CCP and be processed by it within a time window of ten seconds.

Sections 2.7, 2.8, and 2.9 deal with other data flows for which throughput is the vital parameter, especially in view of the convergence of data at SDAC (where the CCP will be located) and at CCA (where the mass store is located).

Conclusions and a discussion of potential problems are presented in section 3.

1.3 A Brief Description of the ARPA Network

The ARPA Network is a packet-switching communication network which allows dissimilar and geographically separate Host computers to communicate reliably and rapidly. The network consists of a number of Interface Message Processors (IMPs) interconnected by 50 Kbit/sec capacity lines; the Host computers are connected to the IMPs (up to four on an IMP). In addition, there are Terminal IMPs (TIPs) which allow computer terminals to be connected to the network, and through the network to a Host, without having to be connected directly to a Host computer.

The IMPs themselves are actually small mini-computers which perform the functions of forming messages into a succession of smaller packets, storing and forwarding the packets through the network to the destination IMP, retransmitting them as necessary to ensure correct receipt, checking the status of neighboring lines and IMPs, exchanging routing information to find the most efficient path through the network from source to destination, and reassembling the packets into messages at the destination IMP for transmission to the destination Host computer. A basic description of the issues of network design may be found in reference 1.

The ARPA Network has grown from 29 nodes and an average of 807,164 internode packets/day in Jan. 1972, to 36 nodes and 1,455,325 internode packets/day in Jan. 1973, to 45 nodes and 2,965,220 internode packets/day in Jan. 1974. During 1974 the traffic growth rate has slowed considerably, until currently the average is between 3.8 and 4.0 million packets/day; the number of nodes has grown to 55. Figure 2 shows the current network topology.

ARPA NETWORK, LOGICAL MAP, JANUARY 1975

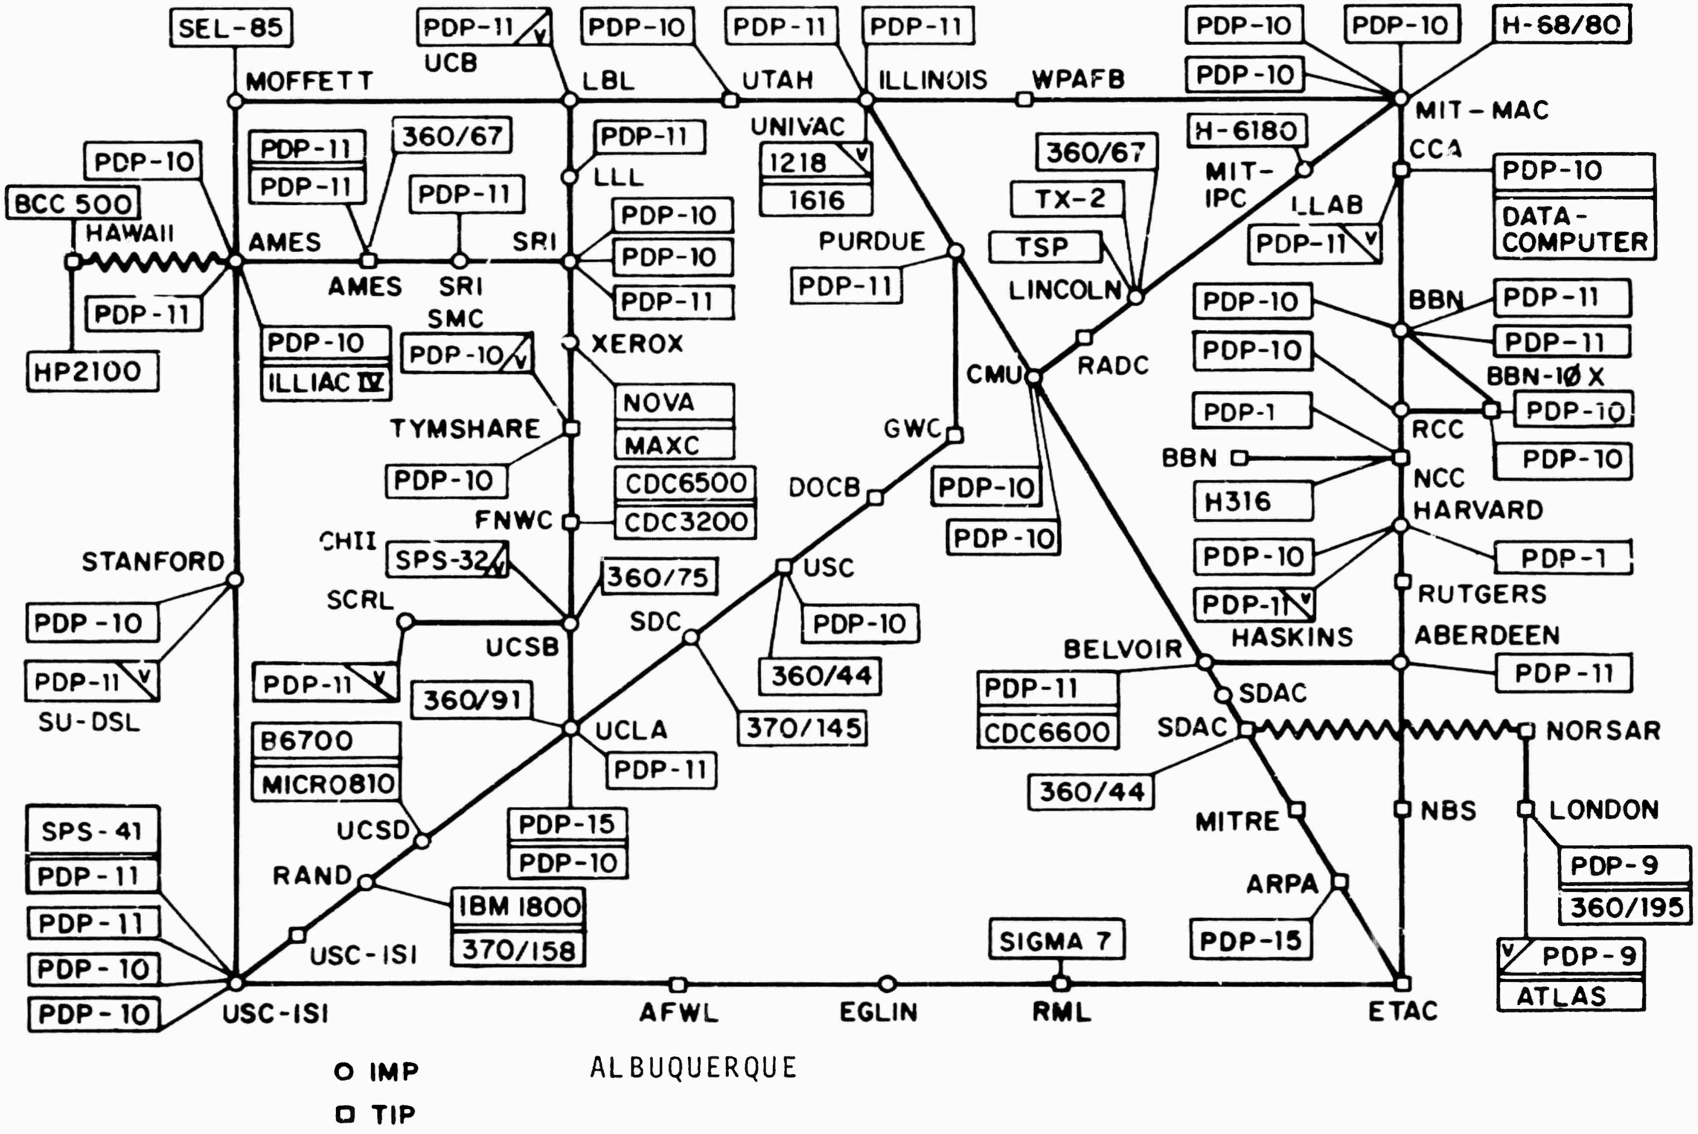

Figure 2

2. PRELIMINARY ANALYSIS OF DATA FLOW PATHS

2.1 An Overview

A logical map of the seismic network was shown as figure 1. Seismic data will be collected at the overseas sites by a station processor at each site (KSRS, ILPA, SITE II) and passed to a Seismic Private Line Interface (SPLI) which looks like a modem to the station processor and like a Host computer to the ARPA Network (link a in figure 1).

The SPLIs will be connected to overseas Satellite IMPs using a Very Distant Host (VDH) interface (see reference 2). The data will be formed into data messages and transmitted from an SPLI to its IMP (link b). The IMP will break the data messages into packets and transmit them across a satellite channel to a receiving Satellite IMP (link c). We shall assume that the overseas network is a natural extension of the continental ARPA Network, rather than a separate network connected by a "gateway", and that Satellite IMPs will be installed on the west coast and east coast. Conventional satellite (or underwater) circuits could also be used to transmit data to the continental United States. From the receiving Satellite IMP the packets will be passed from IMP to IMP (link d) through the continental ARPA Network to the IMP recently installed at SDAC (link e). Assuming the present network topol gy, this amounts to about 9 hops for data from the Pacific and 1 hop for data from the Atlantic.

All the on-line seismic data will converge at the SDAC-IMP, which will have as one of its Host computers the Communications and Control Processor (CCP) for the seismic network. The SDAC-IMP will reassemble the transmitted packets into data messages and pass them to the CCP (link f).

After being processed by the CCP, the data will be output onto the VELA link as well as formed into messages and sent via the ARPA Network to the Seismic Input Processor (SIP) at the CCA-TIP in Cambridge, Mass. (Again, the SDAC-IMP will break the data messages into packets which will be transmitted and reassembled at the CCA-TIP. This is a path of about 8 hops.) In addition, data will be sent from the CCP to the SDAC 360/40A, from the CCP to NORSAR, and long period data will be sent from the mass store to the SDAC 360/40B for analysis. Periodically, long period data recorded on magnetic tape will be sent via the ARPA Network from Albuquerque, N.M., destined for the SIP.

It is important to bear in mind the delay and throughput constraints which must be met in order that the seismic network be effective. First, the delay from the station processor through the CCP must be kept below ten seconds. This maximum delay was chosen because it provides a good compromise between expected transmission delays and buffering needed to average out those delays. In this way the network can be used in a real-time fixed-delay application.

There is no explicit delay requirement for data being transmitted from the CCP to the SIP, but a bound is implied because of the finite buffering ability of the CCP. It is more important to consider throughput for this data, since adequate steady throughput implies that the CCP will never become "backed up" with data destined for the SIP.

Similarly it is important to consider the required throughput for the other data flow paths in the seismic network.

We will begin by examining the delay constraints along the various parts of the path from a station processor through the CCP, and then examine the throughput requirements for the data flow paths.

2.2 Station Processor to SPLI (link a)

The SPLI will be designed to look like a modem from the station processor's point of view. The rate of data transfer will be real-time; that is, it will take one second to transfer a second's worth of seismic data from the station processor to the SPLI.

If the SPLI is located at the station processor site, the propagation, or speed-of-light, delay will be negligible. If the SPLI is located some distance away, the delay is 1 msec/186 miles, or about 5.4 msec/1000 miles.

It takes one second for the station processor to assemble a second's worth of data into a message. As the message begins to be transmitted to the SPLI, it is stamped with the time of day. From the time of this stamp through the processing by the CCP, a maximum of ten seconds can elapse. One second is used to transmit the data message from the station processor to the SPLI (excluding propagation delays). Therefore, nine seconds remain to cover the variable delays encountered in the remainder of this data flow path.

2.3 SPLI - Satellite IMP (link b)

There are two possible locations for the SPLI. The first is at the site of the station processor. This location is preferable from the point of view of reliability, since the line traversing the most hostile environment will contain error-detection and receipt acknowledgment as well as data. The second possible location for the SPLI is at the site of the Satellite IMP. In this case the data line will be traversing the most hostile environment and the data may be erroneously altered without detection.

We will assume, for the sake of argument, that the SPLI will be located at the station processor site. The SPLI is interfaced to the Satellite IMP using the Very Distant Host (VDH) protocol. Briefly, this protocol calls for a Reliable Transmission Package (RTP) for the data including "Hello" and "I heard you" messages to verify the integrity of the communication channel. The sending side of the RTP will break up the data messages into packets of 1008 bits or less and transmit them. The receiving side of the RTP will transmit an acknowledgment for every packet correctly received. If an acknowledgment is not received at the sending side after some time (called the retransmit time), the packet is retransmitted.

Currently this retransmit time is on the order of 250 msec, but this may need to be adjusted. The tradeoff is this: if the retransmit time is too short, the sending side will overburden the receiving side by sending packets more often than necessary, thus cutting the effective bandwidth of the SPLI-Satellite IMP link. If the retransmit time is too long, inordinate delays for lost or altered packets will be introduced which will increase the effective delay.

A good compromise is to set the retransmit time equal to the maximum expected round-trip delay time (the time from packet transmission to receipt of acknowledgment). This can be calculated as follows. First calculate the propagation delay L based on the distance between the SPLI and the Satellite IMP. Then calculate the one-way transmission delay T for a full packet by dividing the packet length by channel capacity. It will take L+T seconds for the packet to reach the receiving side. Acknowledgments are "piggybacked" on return packets (when possible).

If there is one packet in the output queue ahead of the packet on which the acknowledgment will ride, it will take T seconds to output the first packet, T seconds to output the second, and L seconds before the second packet (and the acknowledgment) arrives at the sending side. Thus the maximum expected round-trip delay is 2L plus 3T. Table 2.1 shows this maximum delay value for various channel capacities.

TABLE 2.1 186 miles

(times in milliseconds)

		L	T	2L plus 3T
bits/sec	2400	1	417	1253
	4800	1	208	626
	7200	1	139	419
	9600	1	104	314
	50,000	1	20	62

The additional increase in delay for distances up to 1000 miles is less than 10 msec.

The round trip delay will, in general, be less than this computed maximum since the output queue may be empty and the acknowledgment may be returned on its own short packet instead of piggybacked on a long one.

In order to keep throughput high, the VDH protocol allows two packets to be "in flight", but the sending side must receive an acknowledgment for the first packet before the third packet can be sent. However, it should be noted that it is relatively simple to alter the VDH protocol to allow four packets in flight, if that is necessary for higher throughput.

It is also important to consider the "catch up" capability of the VDH link. Since this link is over a hostile environment and one of the lowest capabilities in the proposed system, it is suspected *a priori* that it may be the tightest bottleneck. Therefore we would like to examine how fast the VDH link allows the SPLI to catch up to real time if it has fallen behind.

For KSRS and Site II, the data rate is 5066 bits/sec including overhead. Consequently, a 4800 bits/sec line is not adequate for these sites. If a 7200, 9600 bits/sec, or 50 Kbit/sec line is used, the time needed to catch up is given in Table 2.2.

TABLE 2.2 Catch-up Times

	1 data frame (5066 bits)		8 data frames
bits/sec	7200	2.374 sec	18.992 sec
	9600	1.117 sec	8.939 sec
	50,000	0.113 sec	0.902 sec

For ILPA, the data rate is 1342 bits/sec including overhead. Consequently the catch-up rates for the ILPA SPLI using a 2400 or 4800 bits/sec line are shown in Table 2.3.

TABLE 2.3 Catch-up Rates for ILPA

	1 data frame (1342 bits)		8 data frames
bits/sec	2400	1.268 sec	10.147 sec
	4800	0.388 sec	3.105

Obviously, the higher the channel capacity, the faster the catch-up rate. However, high channel capacity is expensive or perhaps not available, and at some point (between 18 Kbit/sec and 37 Kbit/sec) the satellite link or the ARPA Network itself will become the tightest bottleneck.

Perhaps another important measure to keep in mind is the amount of time needed to catch up after a delay introduced by a lost or altered packet. If we take the maximum round-trip delay (Table 2.1) as the retransmit time, and use the retransmit time as the estimate of the delay introduced by a lost or altered packet (generally the delay will be less, since the VDH protocol allows two packets in flight), then the required catch-up time is given for KSRS and SITE II in Table 2.4 and for ILPA in Table 2.5.

TABLE 2.4 Catch-up times to offset delay due to a lost packet

	KSRS & Site II (5066 bits/sec)		Maximum tolerable ave. error rate
	7200	0.995 sec	17%
bits/sec	9600	0.357 sec	36%
	50,000	0.0075 sec	96%

TABLE 2.5 Catch-up times to offset delay due to a lost packet

	ILPA (1342 bits/sec)		Maximum tolerable ave. error rate
bits/sec	2400	1.585 sec	33%
	4800	0.243 sec	76%

These times are directly related to the average error rates which the VDH link can sustain without getting behind, also shown in Tables 2.4 and 2.5. For example, the first entry in Table 2.4 indicates that it would take roughly a second to catch up after losing one packet (on a 7200 bits/sec line with a 5066 bits/sec data rate). So this line could sustain an average error rate of about one bad packet in six, or 17%. Again, the higher the channel capacity, the higher the average error rate the channel will be capable of sustaining while maintaining the necessary 5066 bits/sec throughput rate. Keep in mind that this is an average error rate. It is certainly possible to overflow the SPLI's buffering capacity if a number of packets in succession are lost and must be retransmitted.

In the event that channels with a high capacity are difficult or impossible to acquire, it should be possible to use lower capacity lines in parallel. Since an IMP has provision for up to four Hosts, lines from the SPLI running in parallel could be attached to different Host ports at the IMP.

It would be reassuring perhaps to simulate this link with an "error-prone" line to estimate error rates and the effect that altering the retransmit time in the VDH protocol may play at various error rates. On the other hand, it may be difficult to accurately simulate the actual conditions which will be encountered in the field and get meaningful results.

2.4 Satellite hop (link c)

Various methods have been proposed for allowing packet-switching communications over satellite channels. Reference 3 describes the Slotted ALOHA protocol (among others), which is well suited for bursty, interactive communication. It is not very well suited to high-throughput continuous data transfers because the maximum effective bandwidth is only 36% of the circuit bandwidth.

For the seismic network, one would obviously want to use the circuit bandwidth more efficiently, and so some reservation protocol should be implemented. A reservation protocol would avoid the problem of packets from different sources "colliding" at the satellite and having to be retransmitted, since time slots would be reserved in advance for the seismic data packets. Reference 3 also describes such protocols.

Let us assume that 14.4 Kbits/sec is reserved for the on-line seismic data (7200 bits/sec each for KSRS and SITE II for example) out of a 50 Kbit/sec satellite circuit. Then the expected delay is 20 msec for transmission plus 250 msec for propagation. Additional delays of as much as 100 msec may be incurred by some packets while waiting for a reserved slot.

The transmission time will increase for slower circuits up to a maximum of 70 msec for a 14.4 Kbit/sec channel. (Circuitry with channel capacities below 14.4 Kbits/sec would not provide adequate bandwidth.) So we may expect delays of 270 to 320 msec (depending on channel capacity) for the satellite hop using a reservation protocol.

If a retransmission is necessary, this would add a delay of 560 to 710 msec using the 2L+3T formula.

It is possible to use conventional IMPs and communication circuits (which happen to use satellite circuits or underwater cable circuits) instead of Satellite IMPs and a reservation protocol. This is done now to allow NORSAR, London, and Hawaii to connect to the ARPA Network. These existing satellite links have too low a throughput rate (insufficient for the proposed seismic data flow) because of circuit limitations for the Atlantic hop and because of buffering limitations for the Pacific hop.

If conventional leased satellite circuits are used, adequate bandwidth must be provided for the proposed data rates plus additional bandwidth to allow "catching up". This would mean leasing circuits of at least 7200 bits/s[illegible]h for KSRS and SITE II, 4800 bits/sec for NORSAR, and 240[illegible]s/sec for ILPA. The expected delay for a satellite hop configured in this way would be about 250 msec due to propagation, and about 140-420 msec due to transmission (depending on channel capacity), for a total delay of around 400-700 msec. Retransmissions would introduce delays of the order of 1-2 seconds.

2.5 The Continental ARPA Network (links d and e)

The ARPA Network in the continental United States is the strongest link in the path from the SPLI to the CCP. A variety of delay and throughput statistics reports exist for the network which indicate that delay will be no problem, except in extraordinary circumstances.

Theoretically, the minimum round trip delay for a coast-to-coast path of 10 hops is about 400 msec for a 5-packet data message, with the one-way transmit time for the message accounting for about 325 msec. Data from the Network Measurement Center (reference 4) indicates that the average round trip delay times for a coast-to-coast path are about 150 msec. (This shorter time is due to using shorter messages, one or two packets long.)

Another interesting, and relevant, measurement (reference 5) was made on messages sent from Norway to Hawaii using the ARPA Network. This included two satellite hops as well as the coast-to-coast hop. Short (one packet) messages had an average round trip delay of 1.62 seconds, and long (eight packet) messages had an average round trip delay of 4.76 seconds.

Longer delays can be introduced by heavy traffic, marginal lines, or equipment outages. In the case of heavy traffic, additional delay will be caused by a packet having to wait in the output queue of each IMP before being transmitted to the next IMP. Each IMP can have up to eight packets in its output queue, and it takes about 20 msec to transmit each packet in the queue. In the nine hops from the west coast Satellite IMP to the SDAC-IMP, this would add at most (7 x 9 x 20 =) 1.26 seconds of delay. However, if all of the IMP's output buffers were constantly full, this may imply a throughput bottleneck. So in the case of heavy traffic, throughput is the more critical variable.

Throughput summaries are compiled monthly for the ARPA Network. Recent summaries (reference 6) indicate that the most heavily traveled link which the seismic network would use has an average traffic density of 11% of the usable data channel capacity.* A rule of thumb for estimating the peak traffic density is that it is between three and four times the average traffic density. (This rule comes from observation of daily and hourly throughput summaries. A plausibility argument can be made by averaging over a 40-50 hour work week rather than a full 168 hour week.) Thus, at peak times the usual traffic will use up to 44% of the usable capacity.

---

*Of course, this figure may grow before the seismic network becomes operational

Data from KSRS and Site II will be using this link on the coast-to-coast hop, and will need about 25% of the usable capacity. This leaves about 30% (at peak times) for other traffic, or for growth of existing traffic.

Delays may also be caused by marginal lines, leading to altered or lost packets. Between IMPs, the retransmit time is 200 msec. A sending IMP will retransmit a packet every 200 msec until it receives an acknowledgement from the receiving IMP. The average error rate between IMPs is about one bad packet in 10^4, so rarely will any appreciable delay result from IMP-to-IMP retransmissions. (Even if every line in the path from Ames to SDAC were marginal so that each packet had to be transmitted three times before it was received correctly, the resulting additional delay would only be 3.6 seconds.)

Equipment outages (lines or IMPs) pose the most interesting delay problems. If a certain path is no longer working (because a line is declared dead or an IMP has crashed), then the routing algorithm must reroute the packets around the dead path. Order of magnitude estimates indicate that the routing algorithm may take from 1-10 seconds before settling down to a new equilibrium. Therefore it is possible that an equipment failure may, without causing the failure of the network, cause delays of the order of 10 seconds for a few packets.

Again, it would be reassuring to make appropriate delay measurements on the ARPA Network while selectively disabling an IMP or a line. However, this may not be possible to do because of the interference it would cause with other users. In the absence of these measurements, the following comments about the reliability of the ARPA Network are reassuring.

Recent statistics (reference 6) indicate that the average IMP outage is around 1% (one-third of that for preventive maintenance, retrofitting, and other scheduled downs). The statistics for the month of August '74 reveal that the mean time between failures (MTBF) for a particular line is 477.5 hours, or about 20 days. The MTBF for an IMP is 243 hours. (If only hardware/software failures are included, and not scheduled down time, the MTBF is 473 hours.)

If we examine an ARPA Network path with 10 IMPS and 9 lines, and consider the failures between the lines and IMPs to be independent, the MTBF for the path is 16.7 hours (25.0 hours if only the hardware/software failures are included for IMPs). However, this does not necessarily mean that data will be lost or delayed beyond ten seconds whenever a line or an IMP fails. Depending on circumstances, the routing algorithm will usually be able to re-route data packets around the problem and to the destination within the ten-second window.

2.6 SDAC-IMP to CCP (link f)

The Communications and Control Processor (CCP) will act as a network control center for the seismic network. It will be a Host on the SDAC-IMP, with the possibility of being a Host on the SDAC-TIP (should the IMP fail). Since all the seismic data is destined for the CCP, the SDAC-IMP must receive all incoming data packets, reassemble them into messages, and transmit them to the CCP.

After the first packet of a message has arrived, the SDAC-IMP must wait for the remaining packets to arrive in order to reassemble the message. This delay will be about 0.56-0.7 seconds for the Pacific sites, and somewhat less for ILPA and NORSAR.

It has been estimated that the amount of data being received by the CCP over the Network will be about 15 Kbits/sec. Thus,

78

it would take at most 150 msec to transmit the reassembled messages from the SDAC-IMP to the CCP, assuming a Host-IMP interface capacity of at least 100 Kbits/sec. In addition, the CCP will output about 45 Kbits/sec which must pass through the SDAC-IMP. About 22 Kbits/sec will be destined for the 360/40A, about 21 Kbits/sec will be destined for the Seismic Input Processor (SIP), and about 2.4 Kbits/sec will be destined for NORSAR.

It may be desirable to have the CCP and the 360/40A as Hosts on the same IMP. Then the 22 Kbits/sec data from the CCP to the 360/40A would not traverse any ARPA Network lines and would not add to the throughput burden of those lines. Alternatively, if the CCP is a Host on the SDAC-IMP and the 360/40A is a Host on the SDAC-TIP, then approximately 25-30 Kbits/sec of data will be on the link from the SDAC-IMP to the SDAC-TIP. Obviously, in this case it is advisable to use a high capacity modem by-pass for that link.

The final step in this path where delay is the primary concern is processing by the CCP and outputting data on the VELA link. The processing time should consume at most 100 msec, and so does not add significantly to the delay.

2.7 CCP-SIP (links f,e,i, and j)

After the CCP has received the seismic data (within the ten-second delay window), the data is reformatted and transmitted to the Seismic Input Processor (SIP) at CCA. The expected data rate from the CCP to the SIP is about 21 Kbits/sec, and the path through the ARPA Network is eight hops long. (Actually, at present there are three possible eight-hop paths with some lines in common.)

The 21 Kbits/sec data rate is about 55% of the available data bandwidth. Currently the average traffic on the paths between SDAC and CCA is about 8% with estimated peaks of about 30%. Therefore, there is barely enough bandwidth available for the transmission of this seismic data, and there is very little room for growth or transmission of other seismic data during prime usage hours.

In addition, there are some characteristics of the CCA-TIP which will have to be altered in order to fully utilize the available bandwidth. First, since the CCA-TIP is a Terminal IMP, its processor must spend time servicing the user terminals. Preliminary experiments indicate that this will not hinder the throughput of the seismic data. Sufficient processor bandwidth is present for doing both tasks adequately.

Second, one of the Hosts (the Lincoln Laboratory Cambridge Field Station PDP-11) is connected to the CCA-TIP via the VDH protocol. This protocol "borrows" buffer space from the IMP program which the IMP program could use for reassembling the packets into messages. If buffer space is not reserved for reassembly at the destination IMP, the source IMP will not transmit the message. Consequently, if the VDH protocol "borrows" too many buffers, the throughput rate will suffer.

Preliminary analysis indicates that removal of the VDH interface would provide enough buffer space for the seismic network alone. However, as the CCA Datacomputer is used more and more for file storage and retrieval by others, a shortage of message reassembly buffers will eventually develop even with the removal of the VDH.

Currently, throughput experiments are being planned to determine to what extent any throughput limitations are line dependent, buffer dependent, VDH dependent, or processor dependent.

2.8 SRO and LPE Data (links k,i, and j)

The Seismic Research Observatories (SROs) and Long Period Experiment (LPE) sensors will have their data recorded on magnetic tape. It is estimated that data from all the SROs and LPEs will amount to about 1600 bits/sec. Every two weeks, the recorded data (about 1.9 x 10^9 bits) is sent to the Albuquerque Seismological Center for examination. Then the data is to be sent to the mass store via the ARPA Network.

Obviously, such a tremendous amount of data cannot be dumped onto the present network without causing network congestion. For example, the lines from Albuquerque to the east coast are the most heavily utilized in the Network, and the on-line seismic data from the Pacific sites will use these lines as well. Dumping would saturate these lines causing "back-ups" and long delays for the on-line seismic data. In this case the network would be forced to "meter" the incoming data at a non-detrimental rate, or "grow" (by installing higher capacity lines) to absorb it.

A simple way to avoid this problem is to transmit the SRO-LPE data at low rates over an extended period of time. For example, if the data transfer is spread out over two weeks and is done only at night (to avoid competition with normal network traffic), the bandwidth needed would be around 3200 bits/sec.

Along this same line, it might be prudent to consider the use of other means of transmitting the SRO-LPE data to the mass store (e.g. mailing the tapes to CCA).

2.9 Long Period Data Examination

The long period data from all sites (ILPA, ALPA, KSRS, Site II, NORSAR, LASA, SROs, and LPEs) stored at CCA will be examined by seismic analysts at SDAC. Thus a user on the SDAC 360/40B will be retrieving files through the SIP at various times during the working day.

The long period data is accumulated at the rate of about 5 Kbits/sec or about 4.3×10^8 bits/day. If the analyst at SDAC wishes to see 50% of this data during an 8-hour work day, that amounts to an average throughput of 7.5 Kbits/sec during the day, with peaks much higher. This represents a significant throughput demand (about a 20% average demand) on the ARPA Network. In addition, the retrieval and examination of short-period SRO data could add another 1 Kbit/sec to this data flow. Fortunately, since full duplex lines are used, this demand will not interfere with traffic traveling in the other direction (from SDAC to the SIP). However, it will add to the competition for buffer space at SDAC.

For example, at the CCA-TIP we will have 21 Kbits/sec coming from the CCP, in addition to the SRO and LPE data, and an average of 7.5 Kbits/sec going to the 360/40B at SDAC. In addition the TIP must support the usual ARPA Network traffic (average about 2.4 Kbits/sec with peaks estimated at 9.6 Kbits/sec) plus any demands for data from the CCA Datacomputer by other users. At SDAC, about 15 Kbits/sec is coming into the CCP and about 45 Kbits/sec is going out, 21 Kbits/sec destined for the SIP, 22 Kbits/sec destined for the 360/40A and 2.4 Kbits/sec for NORSAR. The SDAC 360/40B will receive an average 7.5 Kbits/sec (with higher peaks). With such throughput, the demands on buffer space at CCA and SDAC are severe.

3. CONCLUSIONS AND DISCUSSIONS OF PROBLEMS

The preliminary study reported on here considered the effect on the ARPA Network of the proposed volume of seismic data to be carried by the network. The study consisted of examining the proposed seismic data flow, theoretical delay and throughput results, as well as available delay and throughput measurements.

The study indicates that seismic data from the field stations will be able to reach the CCP within the allotted ten-second interval.

There are two serious expected problems which will affect the network's throughput capability for seismic data. (See reference 7 for a discussion of network throughput.) The problems arise from a limited amount of buffer space for message reassembly in the destination IMPs and the channel capacity of certain lines which the seismic data is expected to employ.

In the day-to-day use of the network, these problems are not encountered since traffic is more distributed and "bursty". In contrast, the seismic data will be a continuous stream of large volume, and more importantly, it will be converging at the CCP and the mass store.

Specifically, the present CCA-TIP will not be able to handle the expected amount of seismic data being sent to the mass store because of insufficient messages reassembly buffer space for the volume of seismic data converging at CCA. In addition, a Very Distant Host (VDH) interface used by one of the CCA-TIP's Hosts contributes to the problem, it would be sensible to move the VDH interface to a different IMP (or have another IMP support the mass store). Throughput experiments are being planned to more precisely determine what measures will be sufficient to alleviate this deficiency.

In addition, a deficiency of buffer space will also exist at the SDAC-IMP. This is due to the amount of data arriving at the CCP from several sources with round-trip delays of several seconds. Before a data message is sent, reassembly buffers are reserved for it. If the time to complete transmission of the message is long (2-3 seconds), no other messages can use these reserved reassembly buffers during that time. With several sources transmitting simultaneously to the CCP, it is easy to exhaust the buffering capability of the SDAC-IMP.

In view of the expected seismic data flows and current network usage, certain network lines presently do not have adequate bandwidth available for carrying the seismic data as well as other traffic. These are the lines in the northeast corridor from SDAC to CCA.

Basically, the problems encountered are a scarcity of certain network resources (buffer space and channel capacity). Thus, the possible remedies are acquisition of more of the scarce resource, or rationing of its use.

Throughput experiments are being designed to determine to what extent bottlenecks are due to buffer scarcity of lack of channel capacity. Also, an experiment is being planned to observe the time variation of the available bandwidth on the network. This variation in throughput capability could prohibit successful transmission of seismic data even though there is adequate average throughput.

3.1 Message Reassembly Buffer Space

One way of providing more reassembly buffer space would be to replace the affected IMP (or IMPs) by a high-speed modular (Pluribus) IMP (reference 8). In this way, additional memory, I/O, or processing capability could be added easily and only if needed, and taken away to be used elsewhere if not needed. This degree of flexibility may be desirable in view of the evolutionary character of the ARPA Network.

Along this same line, installing a separate (dedicated) IMP at CCA would alleviate the buffer space problem there. This is a brute-force but straightforward way to add the necessary memory.

Adding more memory to affected IMPs is also a solution, but it would entail extensive software modification to make use of the additional memory.

Of course, Host-IMP interfaces should have as high a bandwidth as is practical to facilitate rapid emptying of buffers.

Specifically, at the CCA-TIP, removal of the VDH interface is an adequate short-term solution. This solution would suffice until such time as usage of the CCA Datacomputer begins to put significant demands for reassembly buffer space on the CCA-TIP. Currently, we have no estimate of projected Datacomputer usage by those outside the seismic community. After such time, a Pluribus IMP or a dedicated IMP (with the SIP as its only Host) would provide adequate buffer capacity.

At SDAC, the problem is more severe, since a dedicated IMP (with the CCP as its only Host) would not provide enough reassembly buffers when the entire seismic network is on-line.

Here the cleanest solution is the Pluribus IMP. The alternative is to add more memory to the SDAC-IMP and make the necessary modifications to the IMP software to allow the additional memory to be used.

It should be pointed out that the CCP must have the ability to be connected to the network in at least two ways, so that the CCP can remain on-line when the SDAC-IMP is off-line. The SDAC-TIP, which was to provide the second connection, will not have adequate buffering capability in its present form.

IMPs and TIPs can each support a maximum of 32K words of memory. Currently, each IMP has at most 16K, the SDAC-TIP has 28K and the CCA-TIP has 28K. However, it is expected that TIPs will need the last 4K of memory for terminal-handling at some future time. This will preclude the possible use of this 4K of memory for reassembly buffers in TIPs. Consequently, while it is possible to add memory for reassembly buffers to the SDAC-IMP, it is not possible to use additional memory on the SDAC-TIP.

The implication of this conclusion is far-reaching. The CCP will not be able to be a Host on the SDAC-TIP when the full seismic network is in operation. A second IMP with the memory and software modifications recommended for the SDAC-IMP would be needed to provide the CCP with another means of access to the ARPA Network. (Equivalently, the terminal-handling capability of the SDAC-TIP could be moved to some other node. This would make the SDAC-TIP into an IMP).

The problem at SDAC will not manifest itself until an overseas site (other than NORSAR) begins transmitting seismic data to the CCP. Therefore, a solution does not have to be implemented when the CCP comes on-line, but one will have to be implemented as overseas sites come on-line.

3.2 Channel Capacity

Obviously, upgrading the affected lines from 50 Kbits/sec to 230.4 Kbits/sec would solve the channel capacity problem. We believe that such a drastic solution is not warranted at this time, and therefore we suggest two other approaches.

One approach is to connect a 50 Kbit/sec line directly from SDAC to CCA to bypass the northeast corridor traffic. This could be done by conventional terrestrial circuits or by a domestic satellite circuit and would provide the additional bandwidth necessary to accommodate the CCP-SIP data flow. (The CCP-SIP data would use about 60% of the available data bandwidth of this line.) One possible drawback is that a failure of this new circuit at an inopportune time (such as during a day with heavy network traffic) would force the seismic data onto other lines which then would become nearly overloaded (roughly 92% of the available data bandwidth would be used).

A nearly direct path (5 hops) could be established by installing a line from Harvard to CCA (a distance of a few miles). This is less expensive than a direct line, and would serve to improve throughput by reducing the round-trip delay (fewer hops) and by bypassing the busy BBN nodes. The heaviest traffic would be between SDAC and Harvard, and would be about 80% of the available data bandwidth. However, this would allow little expansion of either seismic or non-seismic traffic along this route.

The second approach places more emphasis on software changes. Since there are two paths between SDAC and CCA (with only the SDAC-Belvoir line in common) the IMP routing algorithm could be altered to allow the CCP-SIP data to be split between the two paths. At peak times, each path would be loaded to about

70% of its data capacity. Two completely independent paths could be created by reconnecting the Belvoir-Aberdeen line between SDAC and Aberdeen, for example.

Such load splitting by using bandwidth routing could be useful for the SRO-LPE data as well. A second independent path from Albuquerque to CCA could be established by installing a line from Albuquerque to DOCB.

The development of this type of routing would also be beneficial to other network users who need to transmit large quantities of data.

It is recommended that a combination of the two approaches outlined above be considered. A direct line from SDAC to CCA, as well as two secondary paths, and a high bandwidth routing algorithm, would provide a great degree of reliability. If the direct line should fail, the routing algorithm would be able to split the CCP-SIP load over the two secondary paths in such a manner that neither would become overloaded.

Such a solution could be reached progressively, as the seismic network takes shape. For example, bandwidth routing could be implemented as the first parts of the seismic network come on-line (early 1976). As the amount of seismic data sent to the SIP increases, the nearly direct path could be established inexpensively. Finally, as the seismic network nears completion, the direct line via satellite or terrestrial circuits could be established.

REFERENCES

1. F.E. Heart, et al, "The Interface Message Processor for the ARPA Computer Network," Spring Joint Computer Conference, 1970, AFIPS Conference Proceedings, vol. 36, pp. 551-567.

2. BBN Report No. 1822, "Specifications for the Interconnection of a Host and an IMP", Appendix F.

3. S. Butterfield, R. Rettberg, and D. Walden, "The Satellite IMP for the ARPA Network," Computer Nets Supplement to the Proceedings of the Seventh Hawaii International Conference on System Sciences, Jan. 1974, pp. 70-73.

4. W. Naylor, H. Opderbeck, "Mean Round-Trip Times in the ARPANET", Network Measurement Group Note #19, NIC 21791, March 7, 1974.

5. William E. Naylor, "On Message Delay Over Worldwide Channels" Network Measurement Group Note #14, NIC #19013, July 26, 1973.

6. Data collected and compiled monthly by the Network Control Center, BBN, Cambridge, Mass.

7. John M. McQuillan, "Throughput in the ARPA Network--Analysis and Measurement," BBN Report No. 2491, January 1973.

8. F.E. Heart, et al, "A New Minicomputer/Multiprocessor for the ARPA Network," National Computer Conference, 1973, AFIPS Conference Proceedings, vol. 42, pp. 529-537.

APPENDIX C

PRIME ITEM DEVELOPMENT SPECIFICATIONS
for the
SEISMIC PRIVATE LINE INTERFACE

PRIME ITEM DEVELOPMENT SPECIFICATIONS
FOR THE
SEISMIC PRIVATE LINE INTERFACE

CONTENTS

1. SCOPE

This specification establishes the performance, design, development, and test requirements for the Seismic Private Line Interface (SPLI) prime item.

2. APPLICABLE DOCUMENTS

The following documents of the exact issue shown form a part of this specification to the extent specified herein. In event of conflict between the documents referenced herein and the contents of this specificaion, the contents of this specification shall be considered a superseding requirement.

MIL-STD-188C, 24 November 1969

"Interface Message Processor - Specifications for the Interconnection of a Host and an IMP", Bolt, Beranek, and Newman, Report No. 1822, March 1974.

"Final Report - Seismic Network Systems Study", Bolt, Beranek, and Newman, Report No. 2865, 9 August 1974.

3. REQUIREMENTS

3.2 GENERAL INTRODUCTION

The SPLI (Seismic Private Line Interface) will be implemented as part of a worldwide seismic data collection and processing network. Its specific task will be to interface a seismic array station controller to an ARPA Network communications link with the Communications and Control Processor (CCP) in the seismic network. The SPLI will be joined to the ARPA Network link by connection to the nearest Satellite Interface Message Processor (Satellite IMP) or, in an expanded overseas Network, possibly to either a TIP or an IMP.

The design specifications given here for the SPLI are based on the data formats and communication protocol defined in BBN Report No. 2865 Appendix A to be used for communication between the CCP and the overseas seismic sites. SPLIs are required for three sites in the currently planned seismic network: one for KSRS, one for SITE II, and one for ILPA. One SPLI design is given for both the KSRS and SITE II sites because the station controllers at those sites are identical. The ILPA station controller differs in both the format and rate of data output, hence a specific ILPA SPLI design is also presented.

In this specification, sections describing SPLI characteristics that are based on design criteria peculiar to an individual site will be divided into two parts: one describing the KSRS/SITE II SPLI and another describing the ILPA SPLI. All other sections establish design specifications that are identical for both kinds of SPLI.

3.2 SYSTEM OVERVIEW

The station controller at a seismic array site generates a real-time data message once each second that consists primarily of one second's worth of seismic data from that site (KSRS/SITE II = 4800 bits, ILPA = 1200 bits). The main task of the SPLI is to transmit these real-time data messages via the ARPA Network to the CCP. To accomplish this task, a communications protocol and message format suitable for ARPA Network transmission must be used by both the SPLI and the CCP. A first draft of such a CCP-SPLI protocol has already been defined and is described in BBN Report No. 2865, Appendix A. The SPLI specifications presented herein employ a somewhat modified version of that CCP-SPLI protocol and should be considered as superseding requirements.

In order to accomodate any short duration circuit outages that may occur in communications to the CCP, the

SPLI will provide limited backup buffering for the real-time data messages that it sends. The CCP must receive a real-time message within 10 seconds after it is generated by the station controller and thus, assuming an overall transmission delay of slightly more than 1 second to the CCP, the SPLI need only provide buffering for up to 8 one second real-time messages.

The SPLI will also respond to a special HELLO message from the CCP by sending an I-HEARD-YOU message to the CCP, in order to allow the CCP to determine whether or not the data link is properly functioning. This status assessment scheme will be used whenever the CCP is not receiving data messages.

Command messages to the KSRS/SITE II station controller and operator messages between the CCP and the ILPA SPLI operators will be communicated by means of a Host level acknowledgement scheme. That is, although the ARPA Network insures reliable transmission of these messages through the IMP subnetwork, a "higher" level message acknowledgement scheme is needed to control the flow of these messages between the CCP and the SPLI. This Host level flow control is necessary in order to minimize the input buffering required at either Host. A Host will acknowledge the receipt of such a message only after it is ready to receive the next message.

KSRS/SITE II SPLI:

Two options are given here for the configuration of the KSRS/SITE II SPLI. The first option assumes that the SPLI will be located in the same facility that houses the station controller. The connection to the Satellite IMP will be as a Very Distant Host* (VDH) by means of a leased line to the satellite ground station (or TIP/IMP). This situation is depicted in Figure 1. The second option assumes that the SPLI will be located at the satellite ground station (or TIP/IMP) where it will be connected to the Satellite IMP as an ordinary Host. The station controller will communicate with the SPLI by means of a leased line. This configuration is depicted in Figure 2.

In both configurations, a leased line is needed between the facility housing the station controller and the satellite ground station (or TIP/IMP) which will contain the Satellite IMP.

ILPA SPLI:

---

* see Appendix F of BBN Report No. 1822.

The ILPA SPLI will be located in the same facility as the station controller and will be connected to the Satellite IMP by means of a leased line to the satellite ground station (or TIP/IMP), employing the error-protected VDH protocol. Operator messages will be communicated with the CCP operator using the console connected to the SPLI. This configuration is shown in Figure 3.

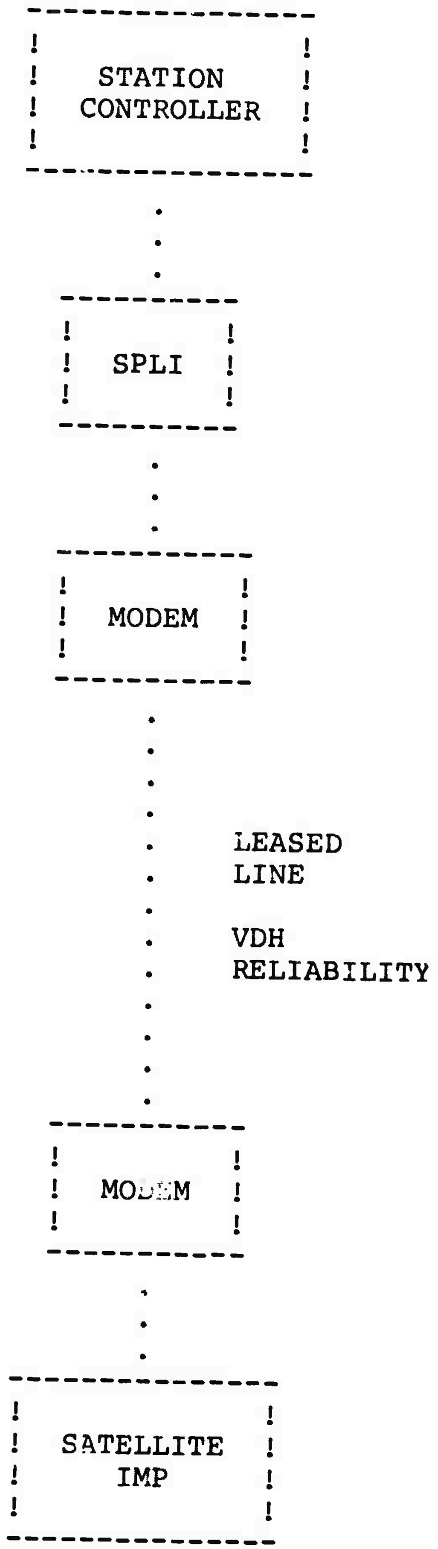

Figure 1. KSRS/SITE II SPLI Configuration - Option 1

```
 ----------------
 !              !
 !   STATION    !
 !  CONTROLLER  !
 !              !
 ----------------
         .
         .
         .
   -----------
   !         !
   !  MODEM  !
   !         !
   -----------
         .
         .
         .
         .
         .
         .
         .      LEASED
         .      LINE
         .
         .
         .
         .
         .
   -----------
   !         !
   !  MODEM  !
   !         !
   -----------
         .
         .
         .
   ----------
   !        !
   !  SPLI  !
   !        !
   ----------
         .
         .
         .
 ---------------
 !             !
 !  SATELLITE  !
 !     IMP     !
 !             !
 ---------------
```

Figure 2. KSRS/SITE II SPLI Configuration - Option 2

```
                       ----------------
                       !              !
                       !   STATION    !
                       !  CONTROLLER  !
                       !              !
                       ----------------
                              .
                              .
                              .
                         ----------     ---------
                         !        !     !       !
                         !  SPLI  !-----!  TTY  !
                         !        !     !       !
                         ----------     ---------
                              .
                              .
                              .
                         -----------
                         !         !
                         !  MODEM  !
                         !         !
                         -----------
                              .
                              .
                              .
                              .
                              .    LEASED
                              .    LINE
                              .
                              .    VDH
                              .    RELIABILITY
                              .
                              .
                              .
                              .
                         -----------
                         !         !
                         !  MODEM  !
                         !         !
                         -----------
                              .
                              .
                              .
                       ---------------
                       !             !
                       !  SATELLITE  !
                       !     IMP     !
                       !             !
                       ---------------
```

Figure 3. ILPA SPLI Configuration

3.5 SPLI/CCP MESSAGE FORMATS

The message format that is to be used in accordance with the communication protocol between the SPLI and the CCP is shown in Figure 4A. It consists of a 2-word Host-IMP leader, a 1-word message class identifier, and a message body.

The specific format of the 32-bit Host-IMP leader that is to be used with messages transmitted by the SPLI is given in Figure 4b. The Destination Address used will be that of the CCP either connected to the SDAC Tip or the SDAC IMP, determined by the SPLI as described in a following section. The message ID is used for Host/IMP message identification.

A message received from the CCP by the SPLI will be considered as valid only if the Host-IMP leader conforms to figure 4c. The Source Address field (bits 9-16) except a HELLO message must contain the CCP address that is currently being recognized by the SPLI (as specified below). The SPLI shall discard all ARPA Network Satellite IMP messages it receives that contain an invalid Host-IMP leader.

It is important that the SPLI maintains a responsive connection with the ARPA Network. That is, it must receive incoming messages held by the Satellite IMP with highest priority in order not to load down that node of the network. If the SPLI fails to take a message within a reasonable time, the Satellite IMP will declare it dead.

The 16-bit message class identifier is used to differentiate between the several types of messages that are communicated

```
----------------------------------------------------------------
!                      !                        !                !
!  HOST-IMP LEADER     !  MESSAGE CLASS ID      !  MESSAGE BODY  !
!                      !                        !                !
----------------------------------------------------------------
       2 words                 1 word            see figs. 5 and 6
```

(a) Format of CCP-SPLI Messages

```
-----------------------------------------
!          !                  !MESSAGE ID !
!   ZERO   !   DESTINATION    !  FOR      !
!          !     ADDRESS      !HOST/IMP   !
!          !                  !  USE      !
-----------------------------------------
   8 bits        8 bits          16 bits
```

(b) Format of Host-IMP Leader from SPLI to CCP

```
--------------------------------
!          !           !         !
!   ZERO   !  SOURCE   !   NOT   !
!          !  ADDRESS  !  USED   !
!          !           !         !
--------------------------------
   8 bits     8 bits     16 bits
```

(c) Format of Host-IMP Leader from CCP to SPLI

Figure 4. Message Format used between the CCP and the SPLI

between the SPLI and the CCP. The classes of messages that are used, along with the definitions of the corresponding message bodies, are shown in Figure 5 for the KSRS/SITE II and in Figure 6 for the ILPA SPLI.

Class 0 messages are sent to the CCP by the SPLI once each second and contain the real-time data messages generated by the station controller.

Unusual delays and circuit outages will cause a backlog of Class 0 messages in the SPLI output queue. As explained above, the SPLI will provide buffering for only eight of such backlogged Class 0 messages. Assuming an overall transmission delay of slightly more than 1 second, this will insure that the oldest saved message will be processed by the CCP within the required 10 second limit. An outage longer than 8 seconds will cause the SPLI to overwrite the oldest of the 9 one second real-time message buffers.

Class 1 (Acknowledgement) messages are sent whenever a host has received a Class 5 or 6 message and is ready (i.e. it has the necessary buffer space) to receive another. The body of an Acknowledgement message contains the 1 word unique identifier of the particular message that is being acknowledged. When the connection is first declared alive, it shall be assumed that the other host is prepared to receive one Class 5 or 6 message.

Message Class ID	Definition of Fields in Message Body	Field Size (words)	Message Interpretation
2		0	HELLO
6	Message identifier	1	Commands to Station Controller
	Time Code	3	
	Character count (≤ 600)	1	
	Command to KSRS or SITE II	≤300	
	Checksum*	1	

(a) Messages from the CCP to the KSRS/SITE II SPLI

Message Class ID	Definition of Fields in Message Body	Field Size (words)	Message Interpretation
0	Data from station controller	300	Real-time Data
	Checksum* (Generated by SPLI)	1	
1	Identifier of message being acknowledged	1	Acknowledgment of Command Message
	Checksum*	1	
3		0	I-HEARD-YOU
4		0	Host Going Down

* checksum = length - (ε data)

(b) Messages from the KSRS/SITE II SPLI to the CCP

Figure 5. Message Classes between CCP and KSRS/SITE II SPLI

Message Class ID	Definition of Fields in Message Body	Field Size (words)	Message Interpretation
1	Unique Identifier of message being acknowledged	1	Acknowledgement of Operator Message
	Checksum*	1	
2		0	HELLO
5	Message identifier	1	CCP to SPLI Operator Message
	Time Code	3	
	Character count	1	
	ASCII Text	≤300	

(a) Messages from the CCP to the ILPA SPLI

Message Class ID	Definition of Fields in Message Body	Field Size (words)	Message Interpretation
0	Data from station controller	51	Real-time Data Message
	Error Count	1	
	Polycode*	1	
1	Identifier of message being acknowledged	1	Acknowledgment of Operator Message
	Checksum*	1	
3		0	I-HEARD-YOU
4		0	Host Going Down
5	Unique Identifier	1	SPLI to CCP Operator Message
	Time Code	3	
	Character count (≤ 600)	1	
	ASCII Text	≤ 300	
	Checksum*	1	

* checksum = length - (ε data)

polycode = cyclic redundancy code assuming the error generator polynomial 11000000000000101

(b) Messages from the ILPA SPLI to the CCP

Figure 6. Message Classes between CCP and ILPA SPLI

Class 2 (HELLO) and Class 3 (I-HEARD-YOU) messages have no body and are exchanged when no data is being received by the CCP from the particular site. Whenever the SPLI receives a HELLO message, it must respond with highest priority by sending a Class 3 (I-HEARD-YOU) message. The I-HEARD-YOU is a Host-to-Host acknowledgement of the corresponding HELLO. The SPLI sets the recognized CCP address to the source address of the HELLO message.

A Class 4, Host going down message will be sent from the SPLI to the CCP whenever the SPLI or station controller is deliberately taken down for any reason such as preventative maintenance.

Class 5 or 6 messages include either command messages to the KSRS/SITE II station controller or operator messages between the CCP and the ILPA SPLI operators. The body of a Class 5 message consists of a 1 word unique message identifier, a 3 word Time-of-day, a 1 word character count, up to 300 words of either station controller commands or ASCII text, and a checksum. Longer commands or text will be sent as a sequence of Class 5 messages. If an end does not receive a corresponding Acknowledgement within a 2 minute time-out period after having sent a Class 5 message then it will retransmit the message.

Class 6 messages are similar to Class 5. The data format is defined as the station controller format.

The time code used with Class 5 messages will be a 48-bit time code identifier specifying when the message was originally sent. The structure of this time code is depected in Figure 7 and consists of eleven BCD 4-bit subfields (year-tens, year-units, day hundreds, ..., second-units).

The ILPA SPLI will generate a unique identifier for each Class 5 message that it sends to the CCP by using a sequential 16-bit counter. When a Class 5 message is formed, the value of the counter is used as the 16-bit unique identifier field and then the counter is incremented. At initialization, the counter is set to the low order 16 bits of the time code.

When an end receives a Class 5 message it must save a copy of the unique identifier as well as returning a corresponding Acknowledgement. If the Acknowledgement is delayed too long or becomes lost then a duplicate Class 5 message will be received. Such duplicates will be detected by noting that their unique identifiers are not greater than the one that has last been saved. At initialization, the saved identifier register is set to zero. Acknowledgements will be returned for the duplicate messages, but the saved unique identifier will not be changed and the duplicates will be discarded.

Messages communicated between the SPLI and the CCP will be sent with the following priorities, from highest to lowest:

Class 3 (I-HEARD-YOU)
Class 2 (HELLO)
Class 0 (real-time data)
Class 1 (Acknowledgement)
Class 5 or 6 (Commands or Operator Message)
Class 4 (Host going down)

```
     ------------------------48 bits--------------------------
--------------------------------------------------------------------------
!               !      !   !       !      !   !      !   !      !   !      !   !
!               !  Y   ! Y !   D   !  D   ! D !  H   ! H !  M   ! M !  S   ! S !
!       0       !  10  ! 1 !  100  !  10  ! 1 !  10  ! 1 !  10  ! 1 !  10  ! 1 !
!               !      !   !       !      !   !      !   !      !   !      !   !
--------------------------------------------------------------------------
                         (4 bits each)
```

Figure 7. Detailed Structure of the Time Code Field

3.6 STATION CONTROLLER/SPLI INTERFACE

KSRS/SITE II SPLI:

Seismic data is output from the KSRS/SITE II station controller at 4800 bits/second (synchronously) in accordance with the low level, serial, digital data specification of MIL-STD-188C. The logical format of a one second 300 word data frame is shown in Figure 8. The SPLI will transmit a data frame to the CCP in a Class 0 message. The SPLI, will compute a single checksum on the message and add it at the end of the message.

The KSRS/SITE II SPLI will receive Class 6 messages from the CCP containing commands to the station controller. This command data will be output to the station controller at 75 bits/second (asynchronously) in accordance with MIL-188C. The logical format of such a command is shown in Figure 9. Option 1 --with the SPLI located in the same facility as the station controller (see Figure 1)--permits direct connection between the SPLI and the station controller. The SPLI will require both a 4800 bits/second synchronous modem simulator interface and a 75 bits/second asynchronous device interface.

Option 2--with the SPLI located at the satellite ground station (see Figure 2) requires a leased modem line connection between the SPLI and the station controller. The SPLI will require both a 4800 bits/second synchronous line interface and a 75 bits/second asynchronous device interface.

ILPA SPLI:

Seismic data is output from the ILPA station controller at 1200 bits/second (asynchronously) in accordance with EIA-RS-232C. Data will be transmitted in 11-bit bytes, each byte consisting of:

start bit
8 bits of data
odd parity bit
stop bit

Only the 8-bit portion of the byte will be included in the real-time data frame sent to the CCP. The logical format of the one-second 51 word data frame received by the SPLI from the station controller is shown in Figure 10.

The SPLI will check the odd parity on each incoming byte and keep a count of the number of parity errors detected during each one second frame of data input. This count will be kept in a 16-bit word and will be appended (as word 52) to the one second data frame that is sent to the CCP each second in a Class 0 message. The SPLI will compute a polycode redundancy check on the one second message to be transmitted and will add it to the end of the message.

The ILPA SPLI will be connected directly to the station controller and thus requires a 1200 bits/second asynchronous device interface.

```
----------------------------------------------------------------
!           !              !                  !         !         !
!  HEADER   !  REAL-TIME   !  WAVEFORM DATA   !  ERROR  !  IDLE   !
!           !    DATA      !  OR ASCII TEXT   !  CODE   !  CODE   !
!           !              !                  !         !         !
----------------------------------------------------------------
   96 bits      -----4672 bits---------         16 bits   14-18
                                                           bits
```

(a) Real-Time Message Format from KSRS and SITE II

```
---------------------------------------------------------
!         !      !           !             !              !
!  SYNC   !  ID  !  STATUS   !  UNDEFINED  !  TIME CODE   !
!         !      !           !             !              !
---------------------------------------------------------
  16 bits 8 bits   8 bits       20 bits        44 bits
```

(b) Format of Header in Real-Time Message from KSRS and SITE II

Figure 8. Real-Time Message Format from KSRS and SITE II

≤ 688 bits*

START CODE	BODY OF COMMAND	TERMINATION CODE
7 bytes	≤ 72 bytes*	7 bytes

* except for "MESS" command which can contain text of up to 150 lines of 72 characters each

BODY OF COMMAND

COMMAND CODE	REQUEST ID	FUNCTION CODE	TYPE CODE	OPTION CODE	PARAMETERS

(ASCII CHARACTERS)

START CODE	SOH	SOH	CR	CR	LF	NUL	NUL
COMMAND CODE	/	/					
REQUEST ID	H	Q	Q	Q	Q	Q	,
TERMINATION CODE	CR	CR	LF	NUL	NUL	EOT	EOT

Figure 9. Command Message Format to KSRS and SITE II

Length (words)	Contents	Description
1	F09A (IN HEX)	SYNC
1	'IL' (IN EBCDIC)	STATION CODE
4	One status byte* each for seven LP sites and one SP site	DATA STATUS
3	0YYDDDHHMMSS	TIME CODE
21	1 FRAME 21 CHANNELS	LP DATA
20	20 FRAMES 1 CHANNEL	SP DATA
1	C8C8 (IN HEX)	END MESSAGE

* Each data status byte will have the format:

Bit On	Description
0	Sync error (remote site to CRS)
1	Faulty or missing LP data
2	Calibration in progress
3	Deleted from beamforming by operator
4	Faulty or missing SP data
5	Extraneous data

Figure 10. Real-Time Message Format from ILPA station controller.

3.7 ILPA SPLI TTY INTERFACE

The ILPA SPLI operator, must be able to exchange text messages with the CCP operator. This implies that the SPLI be interfaced with a low-speed, hardcopy, keyboard-equipped terminal.

A Class 5 message received by the SPLI will contain up to 4800 bits of an ASCII text message from the CCP operator. The time code identifier of the message should be printed on the SPLI console preceding the ASCII text.

The SPLI will store text entered by the SPLI operator in a 80 byte buffer. The SPLI will transmit the contents of a buffer in a Class 5 message to the CCP when either the buffer is filled, a carriage return is typed, or an end-of-transmission character is typed.

To indicate a buffer overflow when the buffer is full, the SPLI should ring the bell and should type a "/" character for each character that is typed wherever the text buffer is full.

3.8 SPLI/SIMP INTERFACE

KSRS/SITE II SPLI:

Option 1--with the SPLI located in the same facility as the station controller (see Figure 1)--requires a dedicated full duplex synchronous modem line between the SPLI and the SIMP. The SPLI will be connected using the Very Distant Host (VDH) interface and will employ the error-controlling VDH line protocol described in BBN Report No. 1822, Appendix F.

In addition to the bandwidth required for the 4800 bits/second of real-time data, the leased line from the SPLI to the SIMP must provide sufficient excess bandwidth for:

1. other classes of CCP-SPLI messages
2. ARPANET and VDH overhead (see Table I)
3. VDH packet retransmission
4. catch-up transmission of backlogged real-time data

An estimate of 1. and 2. is computed in Table II.

The necessary bandwidth then, assuming ideal circuit bahavior, is approximately 5060 bits/second.

To allow for 3. and 4., a minimum leased line bandwidth of 7200 bits/second is recommended.

Option2--with the SPLI located at the satellite ground station (see Figure 2) permits direct connection of the SPLI to the SIMP as an ordinary Host. The specific Host interface required for the SPLI as described in detail in BBN Report No. 1822.

TABLE I

ARPANET and VDH Overhead

(where M = length of message in words without Host-IMP leader)

component	bits
Host-IMP	32
VDH packet headers	$\left(1 + \left\lceil \frac{M}{63} \right\rceil\right) \times 16$
VDH DLE repetition	$\approx \left\lceil \frac{2M}{256} \right\rceil \times 8$ on the average
SYN, SYN DLE, STX, DLE, ETX.	48
VDH checksum	24

TABLE II

KSRS/SITE II SPLI to SIMP Excess Bandwidth Requirements

message class	M (words)	overhead (bits)	messages/ second	total of 1. and 2. (bits)
0	302	224	1	5056
1	4	144	$\frac{1}{16}$	4
		TOTAL		5060

ILPA SPLI:

The ILPA SPLI will require a dedicated full duplex synchronous modem line to the SIMP to which it will be connected as a Very Distant Host (VDH). The SPLI will be equipped with a VDH interface that uses the error-controlling VDH line protocol described in BBN Report No. 1822, Appendix F.

In addition to the bandwidth required for the 832 bits/second of real-time data, the leased line from the SPLI to the SIMP must provide sufficient excess bandwidth for:

1. other classes of CCP-SPLI messages
2. ARPANET and VDH overhead (see table I)
3. VDH packet retransmission
4. catch-up transmission of backlogged real-time data

The expected total of 1. and 2. is computed in table III.

The necessary bandwidth then, assuming ideal circuit behavior, is 340 bits/second.

To allow for 3. and 4., a minimum leased line bandwidth of 2400 bits/second is recommended.

TABLE III

ILPA SPLI to SIMP Excess BAndwidth Requirements

message class	M (words)	overhead (bits)	messages/ second	total of 1. and 2. (bits)
0	54	144	1	1008
1	4	144	$\frac{1}{16}$	13
5	305	224	$\frac{1}{16}$	319
			TOTAL	1340

3.9 ENVIRONMENT

The SPLIs will operate over a temperature range of 0° to 45°C (32° to 110°F) and humidity from 0 to 90% with no condensation. The equipment will be able to be stored or shipped over a range of -20° to 50°C (0° to 120°F).

ABBREVIATIONS

SPLI	Seismic Private Line Interface
IMP	Interface Message Processor
CCP	Communication and Control Processor
KSRS	Korean Seismic Research Station
ILPA	Iranian Long Period Array
SDAC	Seismic Data Analysis Center
SIMP	Satellite Interface Message Processor