

AD-787 693

SEISMIC NETWORK SYSTEMS STUDY

Howard W. Briscoe

Bolt Beranek and Newman, Incorporated

Prepared for:

Air Force Technical Application Center
Advanced Research Projects Agency

9 August 1974

DISTRIBUTED BY:

NTIS

**National Technical Information Service
U. S. DEPARTMENT OF COMMERCE
5285 Port Royal Road, Springfield Va. 22151**

DOCUMENT CONTROL DATA - R & D

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

1. ORIGINATING ACTIVITY (Corporate author) Bolt Beranek and Newman Inc. 50 Moulton Street Cambridge, Mass. 02138		2a. REPORT SECURITY CLASSIFICATION	
		2b. GROUP	
3. REPORT TITLE FINAL REPORT: SEISMIC NETWORK SYSTEMS STUDY			
4. DESCRIPTIVE NOTES (Type of report and, inclusive dates) Final Technical Report			
5. AUTHOR(S) (First name, middle initial, last name) Bolt, Beranek and Newman Inc.			
6. REPORT DATE August 1974		7a. TOTAL NO. OF PAGES 78 84	7b. NO. OF REFS 5
8a. CONTRACT OR GRANT NO F08606-74-C-0028		9. ORIGINATOR'S REPORT NUMBER(S) BBN Report No. 2865	
b. PROJECT NO VELA T/4701/B/ETR		10. OTHER REPORT NO(S) (Any other numbers that may be assigned this report)	
c.			
d.			
10. DISTRIBUTION STATEMENT APPROVED FOR PUBLIC RELEASE, DISTRIBUTION UNLIMITED			
11. SUPPLEMENTARY NOTES		12. SPONSORING MILITARY ACTIVITY Advanced Research Projects Agency Arlington, Virginia 22209	
13. ABSTRACT <p>This report presents design recommendations by Bolt Beranek and Newman for the data collection function of a worldwide seismic data network. Particular attention has been paid to integrating both ARPA Network packet switching technology and conventional leased channels. It appears not only feasible but attractive to employ such a mixed approach. The resulting communication system should be reliable, flexible, and cost-effective. Two detailed protocols included in this report specify the communications between four remote array sites and a central Communications and Control Processor (CCP) and the communications between the CCP and the seismic network mass storage facility.</p>			

Reproduced by
NATIONAL TECHNICAL
INFORMATION SERVICE
U.S. Department of Commerce
Springfield VA 22151

84

14

KEY WORDS

LINK A

LINK B

LINK C

ROLE

WT

ROLE

WT

ROLE

WT

Seismology
 Computer Systems
 Computer Networks
 Data Communication
 Packet Switching
 Seismic Detection

ia

Report No. 2365

Bolt Beranek and Newman Inc.

FINAL REPORT
SEISMIC NETWORK SYSTEMS STUDY

9 August 1974

Sponsored by:

Advanced Research Projects Agency
Arlington, Virginia 22209

Prepared by:

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

APPROVED FOR PUBLIC RELEASE,
DISTRIBUTION UNLIMITED

ib

TABLE OF CONTENTS

	Page
SUMMARY	1
1. INTRODUCTION	3
2. TECHNICAL DISCUSSION	5
2.1 Purpose of the Investigation	5
2.2 Technical Background - ARPA Network Communications	5
2.3 Description of Work	7
2.3.1 Data Collection System Overview	7
2.3.2 Results	9
2.3.2.1 Phase I Results	9
2.3.2.2 Phase II Results	10
2.4 Discussion of Results	10
2.4.1 Existing Constraints	10
2.4.2 Special Purpose Protocols	11
2.4.3 Protocol Commonality	12
2.4.4 Protocol Flexibility	12
3. CONCLUSIONS	14
4. RECOMMENDATIONS FOR FUTURE WORK	15
REFERENCES	17
GLOSSARY	18
Appendix A - CCP - Site Protocol	A-1
Appendix B - CCP - SIP Protocol	B-1

LIST OF ILLUSTRATIONS

	Page
Figure 1. Data Collection System Overview	8

SUMMARY

In order to provide a large data base for research in detection and identification of contained nuclear explosions and to demonstrate the feasibility of an on-line system for monitoring a nuclear test ban, ARPA plans to expand the existing digital seismic network to provide more extensive worldwide coverage. This network will utilize state-of-the-art technology, particularly in the area of data collection and storage. Bolt Beranek and Newman Inc. (BBN) has been asked, under the present contract, to work on the design of the on-line data communication system integrating both ARPA Network packet switching technology and conventional leased channels.

The approach taken has been to start with a broad system definition and to make successive refinements of the design. Under a previous contract [1], BBN described the system overview. Under the current contract BBN has considered specific alternatives to the previous design and has described the interfacing of the packet switching links to the rest of the system in detail. Experience provided by previous work at interfacing computers on the ARPANET has been combined with the constraints existing in the seismic data network.

The main results of this work are the two communication protocols included as appendices to this report. Additional results summarized in section 2.3.2.1 concern possible options for specific components of the seismic data network. The protocols specify in detail the manner in which array sites use the ARPANET to communicate with the Communications and Control Processor (CCP) and how the CCP communicates with the Seismic Input Processor (SIP) located at the mass storage site. By

specifying what data is exchanged between the various pieces of the on-line data collection system the protocols effectively define that system.

One general implication of the current work is the feasibility of using packet switched and conventional technology in a mixed communication network, even where various real-time constraints may exist. The overall seismic data network design demonstrates the potential of the ARPA Network for providing access to shared hardware and software resources.

The work done under the current contract suggests three specific areas for further research and development. First, the protocols complete the input/output specifications of the Seismic Private Line Interfaces (SPLIs). The design of these devices can, therefore, begin. Second, a study should be initiated to determine how the seismic data network will be affected by various changes in structure or loading of the ARPANET. The design of a routine monitoring procedure to be used by the CCP would be desirable. Finally, the interface between the individual research seismologist and the data base has not yet been addressed. The design and implementation of a seismically oriented man-machine interface would capitalize on the large data base which will exist on the mass storage facility.

1. INTRODUCTION

In a continuing program to provide high quality digital seismic data for research in detection and identification of contained nuclear explosions, and to demonstrate the feasibility of an on-line system for monitoring nuclear test bans, ARPA plans to expand the existing digital seismic network to obtain more extensive worldwide coverage, particularly for long period data. Three additional small on-line digital arrays and 15 to 20 single point stations recording on digital magnetic tape are to be operational in 1975. The resulting seismic data collection network will produce a formidable library of valuable digital seismic data which must be organized and made available to the seismic research community.

Under a previous contract, BBN prepared a recommendation [1] for an overall system design for the world-wide seismic data handling system. The design included use of state-of-the-art techniques in the rapidly developing areas of communication, processing and data storage.

The objectives of the present contract have been to update the recommended system design to account for changes in the seismic network configuration or developments in the relevant technology and to provide more detailed design and perhaps breadboard implementation of parts of the system as directed by the project officer.

For system analysis purposes it has been convenient to think of the seismic network as serving or performing three interrelated functions: 1) data capture, 2) event processing, and 3) seismic research. Three phases or levels of design have

also been defined. Phase I consists of examining optional configurations to arrive at comparative estimates of costs, schedules, and capabilities. Phase II consists of defining recommended formats, protocols, and flow charts. Phase III consists of designing and implementing experiments to test or develop techniques required by the recommended design.

As directed by the project officer, the work carried out by BBN under the present contract has focused primarily on the Phase I and Phase II design of the data capture function.

2. TECHNICAL DISCUSSION

2.1 Purpose of the Investigation

The work carried out and reported on in the following sections was part of a larger effort to design and implement the seismic data collection and analysis network mentioned in section 1.

In addition to providing data to support seismic research, this network should demonstrate the feasibility of an on-line system for monitoring nuclear explosions which will be essential for implementation of any underground nuclear test ban.

The primary goal of BBN's work has been to design the communication system for the on-line data collection system integrating both ARPA Network packet switching technology and conventional channels leased from the common carriers. The design objective was to develop a cost-effective communication system flexible enough to adapt to changes in the deployment of sensors. BBN has also assisted with the application of other state-of-the-art technology in the seismic data network.

2.2 Technical Background - ARPA Network Communications

The ARPA Network [2,3] provides a flexible, reliable, and economically attractive means for dissimilar, geographically distributed computers (Hosts) to communicate via common-carrier circuits. Each Host connects into the network through a small local computer called an Interface Message Processor (IMP); each IMP is connected to several other IMPs via wideband (typically 50 kilobit) communication lines. The IMPs, all of which are virtually identical, are programmed to store and

forward messages to their neighbors based on address information contained in each received message. The route that a message will take through the network is determined dynamically and depends on network loading as well as IMP and line failures.

The store-and-forward logic and the error control procedures implemented in the IMP subnetwork (which attempt to insure correct delivery of a message by retransmitting a section of a message when an error is detected) are optimum for asynchronous and very bursty communication between various nodes in the network. Many of the major resources connected to the ARPANET as Hosts also assume interaction over a bursty communication system with other Hosts which are tolerant of large data rate variations and variable transmission delays. Real-time data sources and real-time recording systems such as the seismic observatories and the VELA link, however, cannot tolerate variable-speed bursty system components. An important function of the seismic communication system design, therefore, is to interface these two classes of resources in the seismic data network.

The general approach to this problem is to provide large buffers at both ends of each communication link which use the ARPA Network to implement an apparently non-bursty channel. These buffers are used to smooth out the variations in message delivery time. Messages are delivered a fixed delay behind real-time, the delay being specified by the amount of buffering provided. When transients in the bursty part of the system exceed the delay time, of course the on-line subsystems will see a condition consistent with a transient phone line outage which they are designed to tolerate. The delay is specified

so as to reduce the probability of this occurring while keeping the required buffer space within reasonable bounds.

2.3 Description of Work

2.3.1 Data Collection System Overview

The structure of the on-line seismic data collection network is shown in Figure 1. There are six array sites which are the sources of the on-line seismic data. The destinations for this data are (1) the Seismic Input Processor (SIP) at The Computer Corporation of America (CCA), (2) the 360/40A at the Seismic Data Analysis Center (SDAC), and (3) the VELA Links at SDAC. Rather than have each source be concerned with each of the destinations, each source simply sends its data to a central site, Communications and Control Processor (CCP), in a convenient form. The CCP takes care of any required reformatting and transmission to the individual destination sites. The CCP also provides a central site for overall seismic network control, performance monitoring, and maintenance.

Two of the array sites (ALPA and LASA) are connected to the CCP over conventional leased channels. The remaining source sites are connected utilizing ARPA Network connections. The NORSAR 360/40A is interfaced as a local Host on the NORSAR TIP while the station controllers at IRAN, KSRS, and Site II are interfaced to Satellite IMPs (SIMPs) by means of Seismic Private Line Interfaces (SPLI). These devices act as Hosts on the ARPANET and provide a convenient approach to interface computers (i.e., the station processors) which are designed to communicate over standard modem interfaces. The 360/40A at SDAC and the SIP at CCA are connected as local Hosts on the

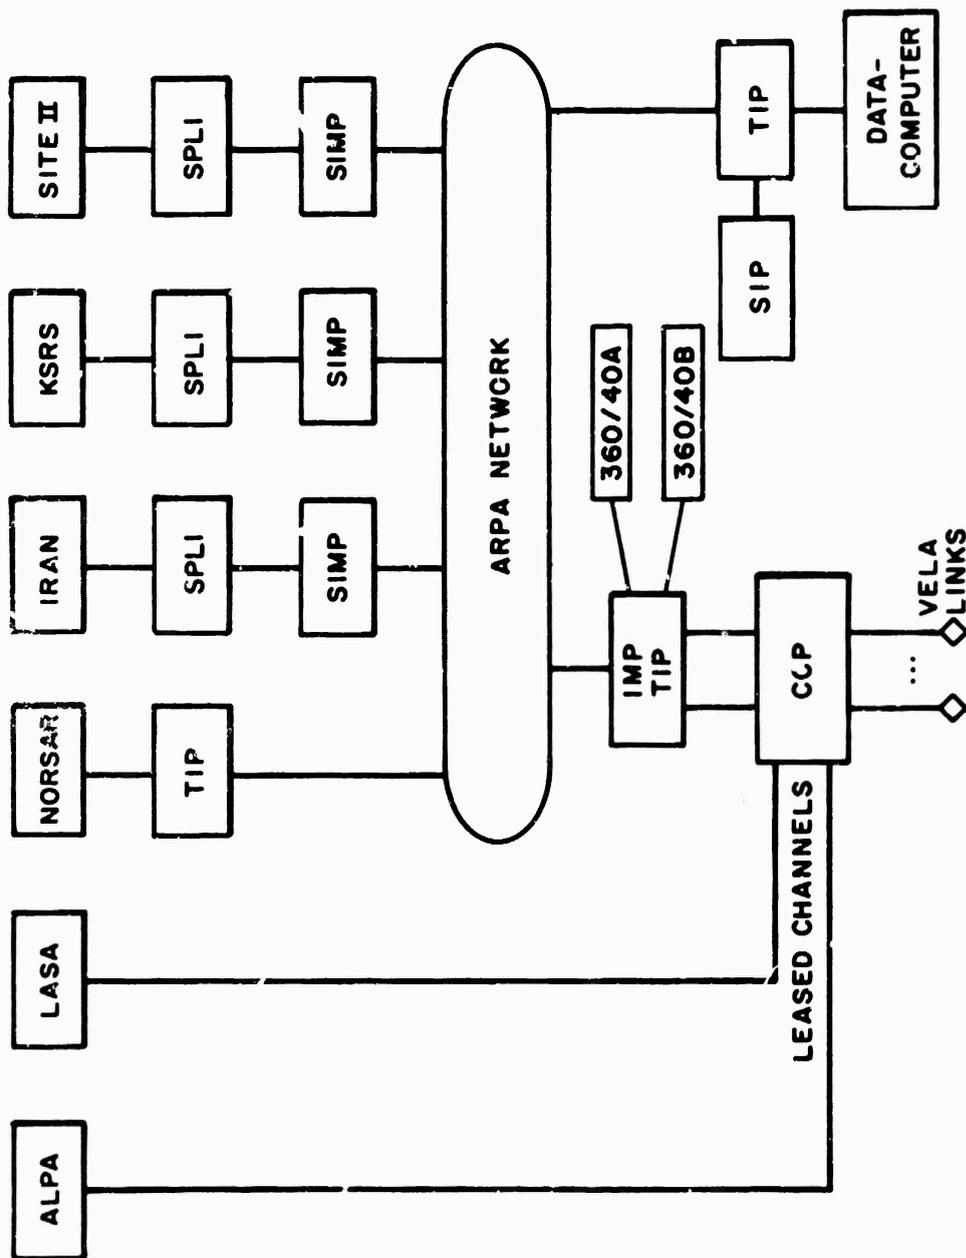


Figure 1 Data Collection System Overview

ARPA Network. The VELA links consist of a number of 4800 baud leased channels.

2.3.2 Results

2.3.2.1 Phase I Results

Several trade-offs not considered in the previous network design [1] were considered under the current contract. Options involving the areas of data communication and man-machine interaction were presented to the project officer and are summarized below.

LASA: In light of the planned reconfiguration of the LASA seismic array, a study was carried out to determine the most appropriate method of providing LASA ARPANET access. Alternatives involving either use of the LDC 360/44 as a Very Distant Host (VDH) or the acquisition of a new minicomputer to replace the 360/44 and provide a VDH interface were presented to the project officer. It was subsequently decided that connecting LASA to the CCP via a standard leased line would be more cost-effective than either of the alternatives mentioned above.

NORSAR: BBN was asked to explore the possibility of modifying the current on-line seismic data network by replacing the existing Trans-Atlantic Link (TAL) by an ARPANET connection. A detailed analysis of the resource availability (processor, core, disk, and channel) in the Detection Processing (DP) systems at SDAC and NORSAR indicated that this would be feasible. Estimates of the time and costs required to make the necessary changes were presented to the project officer.

NEPT: Several options concerning device types and interfacing schemes were considered for the Network (Analyst) Event Processing Terminal (NEPT) at SDAC. After considering storage tube devices, conventional refreshed displays, and TV scan displays it was recommended that a conventional refreshed display be employed. Such a graphic display could be interfaced as 1) a terminal on the SDAC TIP, 2) a mini-host on the SDAC TIP or IMP, or 3) a peripheral device on the SDAC 360/40B. In light of the expected use of the NEPT, either option (2) or (3) was recommended. Cost estimates for the recommended options were also presented to the project officer.

2.3.2.2 Phase II Results

The majority of the work done under the current contract involved Phase II level design of communication formats and protocols. The protocols specify the data exchange between four of the array sites (NORSAR, IRAN, SRS, and Site II) and the CCP and between the CCO and the SIP. The protocols, in their current form, are included in Appendices A and B. They should be considered as working documents and may be subject to further revision. Various issues related to these protocols are discussed in section 2.4.

2.4 Discussion of Results

2.4.1 Existing Constraints

A basic design decision for the on-line seismic data collection network has been that the data communication system should be adapted to the existing site computer systems wherever possible. Since these systems were originally designed to communicate over conventional leased communication channels,

the resulting interfaces are not necessarily the most straightforward. The data formats from the source computers (station processors) also reflect the original leased channel assumption. The communication system and protocols might be considerably different if ARPANET communications had been assumed from the start.

The other existing constraint has been the continuous transmission fixed-delay constraint (see section 2.2) implied by the VELA links.

2.4.2 Special Purpose Protocols

In order to effect communications between ARPANET Host computers, several levels of communication protocols are required. All Hosts must implement the Host/IMP protocol described in BBN Report No. 1822 [4]. In addition, any pair of Hosts that expect to communicate with each other over the network must implement some mutually acceptable Host/Host protocol. An elaborate Host/Host protocol for interconnecting processes in large general-purpose computer facilities is described in [5]. Finally, programs running in Hosts communicating with each other must implement the protocol or language of the other user program and/or operating system.

Since the seismic data collection function does not involve communication with either large general-purpose computing centers on the network or Hosts outside of the seismic network (data collection and archiving require that the CCP communicate with the SIP, not the Datacomputer system itself), special purpose Host/Host protocols meeting the particular real-time requirements of the seismic network have been proposed. The

CCP resources that must be devoted to Host/Host protocol implementation using these protocols are significantly less than required for a generalized Host/Host protocol.

2.4.3 Protocol Commonality

The initial versions of the protocols for the 3 different array sources (NORSAR, IRAN, and the pair KSRS, Site II) were developed and described independently. It soon became apparent, however, that a common seismic data input protocol was possible. The merging of the three original protocols resulted in the protocol included as Appendix A. This common input protocol should simplify implementation of data communication software both at the remote sites and at the CCP.

A similar set of statements can be made concerning the protocol in Appendix B which applies to CCP communication with both the SIP and the 360/40A at SDAC.

2.4.4 Protocol Flexibility

Flexibility has been an important consideration in the design of the current protocols especially in light of the research nature of the seismic data network. The protocols should not constrain the deployment of channels in terms of either the number of array sites or the number of long and short period channels returned from the sites. The protocols should also allow any subset of the received data to be stored at the seismic network mass storage facility.

Data is returned in terms of long and short period channels, therefore, predetection processing can be done either centrally, at the remote site, or at both. Raw and processed

waveforms are treated symmetrically in the CCP-Site protocol and may both be present in the real-time data messages from the sites.

The CCP to SIP protocol is extremely flexible. Periodic control messages identify the specific channels which will be included in subsequent data frames, and the destination file of each channel. Data flow to the SIP can be modified under command from the CCP operator.

Adaptation to changes in channel deployment at the remote sites is straightforward, however it is not automatic and source-initiated as in the case of the CCP-SIP protocol. If a new site were to come on-line, the same protocol used for NORSTAR, IRAN, KSRS, and Site II would be employed. The format particular to this new site would be defined and the CCP input format definition tables and buffer allocation would be modified to reflect the addition. A similar change would be required if additional channels were returned from an existing site. The elimination of an entire site or a reduction in the number of channels sent to the CCP can be handled by considering the changes as a temporary loss of the associated channels or by modification of the CCP format definition tables and buffer allocation as mentioned above.

3. CONCLUSIONS

It appears to be both feasible and cost-effective to implement the on-line seismic data collection and storage network using a combination of ARPA Network links and conventional leased channels. Initial steps toward the design of this network can be found in section 2.3.2 and the Appendices.

Appendices A and B should be considered as working documents. Although these specifications incorporate the latest design decisions related to the seismic data network, they may require minor modification as hardware and seismic data requirements become firm.

4. RECOMMENDATIONS FOR FUTURE WORK

There are several projects which would be logical follow-ups to the work reported on here. These areas for future work are listed below:

SPLI Design - Based on the attached protocols it appears appropriate to begin final design of the SPLI devices for the 3 array sites. This would involve a specification of the required hardware capabilities and a detailed design of the software.

System Throughput Analysis - The seismic data network will utilize the ARPANET to transmit real-time data in a way previously not tried. Assuming that there is sufficient bandwidth available on the ARPANET links, there is every reason to believe that the protocols mentioned previously will work. It will be important, however, for the seismic data network to anticipate any changes in the ARPANET structure or load which will significantly affect the seismic data subnetwork. A study should be made to determine what measurements and monitoring of the ARPANET is appropriate and how it can be accomplished.

Seismologist Interface System - Given that a large and interesting seismic data base will be accumulated by the on-line system it is important to insure that this data can be accessed easily by individual research seismologists. The research seismologist should not have to learn Datalanguage to access data for his local computer. Instead, we believe there should exist an interface system which translates between a language oriented toward the seismologist and Datalanguage. The design of this system must begin as soon as possible if it

is to be available when a significant volume of data accumulates in the Datacomputer. The design of the data access procedures could also provide valuable insight into the design of file formats that facilitate the necessary interaction between data and status files and between raw data and processed data files.

REFERENCES

- [1] Bolt Beranek and Newman Inc., Final Report - A Study of the Data Collection, Processing, and Management for a Worldwide Seismic Network, BBN Report No. 2632, September 1973.
- [2] Roberts, L. G., and Wessler, B. D., "Computer network development to achieve resource sharing," AFIPS 1970 Spring Joint Computer Conference Proceedings, Volume 36, pp. 543-549.
- [3] McQuillan, J. M., Crowther, W. R., Cosell, B. P., Walden, D. C., and Heart, F. E., "Improvements in the design and performance of the ARPA Network," AFIPS 1972 Fall Joint Computer Conference Proceedings, Volume 41, pp. 741-754.
- [4] Bolt Beranek and Newman Inc., Specifications for the Interconnection of a Host and an IMP, BBN Report No. 1822, March 1974.
- [5] McKenzie, A., Host/Host Protocol for the ARPA Network, NIC #8246, January 1972.

GLOSSARY
of Acronyms and Abbreviations

ALPA	Alaska Long Period Array
ARPA	Advanced Research Projects Agency
ARPANET	ARPA Network
BBN	Bolt Beranek and Newman Inc.
CCA	Computer Corporation of America
CCP	Communications and Control Processor
DP	Detection Processing
IMP	Interface Message Processor
IRAN	Iranian Seismic Array
LASA	Large Aperture Seismic Array
KSRS	Korean Seismic Research Station
LDC	LASA Data Center
NEPT	Network Event Processing Terminal
NORSAR	Norwegian Seismic Array
SDAC	Seismic Data Analysis Center
SIMP	Satellite Interface Message Processor
SIP	Seismic Input Processor
SPLI	Seismic Private Line Interface
TAL	Trans-Atlantic Link
VDH	Very Distant Host

Table of Contents

1. Introduction and Background.....	3
2. Real-Time Data Transmission Formats.....	9
2.1 Real-Time Data from KSRS and SITE II to the CCP.....	11
2.2 Real-Time Data from IRAN to the CCP.....	13
2.3 Real-Time Data Messages between NORSAR and the CCP.....	15
3. Format for Control Messages.....	18
3.1 Commands to KSRS and SITE II.....	24
3.2 Operator Messages between the IRAN SPLI and the CCP.....	25
4. Detailed Operation of the Communications Links.....	26
4.1 Transmission of Real-Time Data.....	26
4.2 Flow of Commands and Operator Messages.....	28
4.3 The Question of End-to-End Acknowledgement of Real-Time Data.....	29
4.4 Error Recovery and Initialization.....	30
5. References.....	32

List of Figures

1. Site-CCP Data Link for Sites Connected to a SPLI.....	7
2. Site-CCP Data Link for NORSAR.....	8
3. Format of Site-CCP Real-Time Messages.....	10
4. Real-Time Message Format from KSRS and SITE II.....	12
5. Real-Time Message Format from IRAN.....	14
6. Real-Time Message Format from NORSAR.....	16
7. Real-Time Message Format to NORSAR.....	17
8. Format of CCP-SPLI Control Messages.....	19
9. Classes of Control Messages from the CCP to the SPLI.....	20
10. Classes of Control Messages from the SPLI to the CCP.....	21
11. Command Message Format to KSRS and SITE II.....	23

1. Introduction and Background

The overall design of a worldwide seismic data network utilizing ARPANET communications has been described in [1]. In that design six array sites enter raw and processed data by transmitting it to a central Communications and Control Processor (CCP). This document specifies in detail the communications formats and protocols between the CCP and four of these Sites--the Korean and Site II Seismic Arrays (KSRS and SITE II), the Iranian Long Period Array (IRAN), and the Norwegian Seismic Array (NORSAR).

The communication protocols described below build upon the Host-IMP Protocol [2] and are at the Host-to-Host level although they are not the standard Host-Host Protocol described in [3]. 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, the overall communications requirement for each site is in essence the same. This document presents communications formats and protocols that could readily be used with new array sites in the network, each having its own individual peculiarities.

These protocols are designed to provide two unusual features in the ARPANET links. First, since the CCP is central to the operation of the on-line data collection system it must be designed to operate continuously. To achieve a reliable seismic communications subnet it is desirable to connect the CCP to at least two IMPs (more precisely, an IMP and a TIP in the case of SDAC). With this arrangement, the CCP can be switched from one IMP to the other if

the one on which it is operating is having troubles or is scheduled for preventive maintenance. Although such a duplex arrangement has not yet been tried, such a scheme should be straightforward in the current seismic data network application.

Second, because of the need to continuously retransmit and plot waveforms from several different sources with a known shift between the individual waveforms, it is necessary to operate the ARPANET links in the seismic network so that the data will be delivered to its destination some constant period behind real-time, i.e. the network should appear to be an "ideal (error free) delay line" between the data source and destination. Although this mode of operation is somewhat unnatural on the ARPANET, it can be effected by providing sufficient buffering for the data at both ends of the communication link. There are tradeoffs, however, relating to the length of outages (delays) that can be handled and the size of buffers and bandwidth required to allow catchup.

The amount of data buffered by a Site will be approximately 10 seconds. This period was selected in order to (1) reduce the loss of real-time data due to noise burst errors, retransmissions, and network delays, (2) to be frugal in the amount of space allocated to buffering, and (3) to provide the option of a uniform delay across all six of the array sites. For array sites with a 4800 baud data rate, 10 seconds of data equals 3K 16-bit words of buffer at each end. Longer delays leading to larger buffer sizes were judged inappropriate since in many cases these will be core buffers.

It is possible, of course, that the network will be unavailable for periods sufficient to overflow this 10 second buffer. Assuming all of the data is recorded by the Site on tape, however, the data which overflows the buffer is not actually lost. A copy of this tape could be sent by mail to SDAC or selected portions of the data could be transmitted over the ARPANET.

Two slightly different physical configurations will be used for the ARPANET communications links between the overseas array sites and the CCP. Figure 1 depicts the configuration that will be employed by KSRS, SITE II, and IRAN. The station controller at each of these array sites is interfaced to the ARPANET by means of a SPLI (Seismic Private Line Interface). The SPLI is connected as a Host using the VDH (Very Distant Host) Host-IMP Protocol [2] to the nearest SIMP (Satellite IMP) by means of a leased line. This SIMP will be located at a satellite communications ground station (the Kum San ground station for KSRS and the Asadabad ground station for IRAN). Communications over the Atlantic will be via the SIMP located at the Etam, West Virginia ground station. KSRS and SITE II will be linked to the continental U.S. ARPANET by means of the SIMP at the ground station in Jamesburg, California.

The physical configuration that will be used for the ARPANET communications between NORSAR and the CCP is shown in figure 2. The NORSAR Data Processing Center (NDPC) is interfaced as a Host to the NORSAR TIP (Terminal IMP). The TIP is connected indirectly to a SIMP located at the Goonhilly Downs ground station. Communications

over the Atlantic are again via the SIMP in Etam, West Virginia.

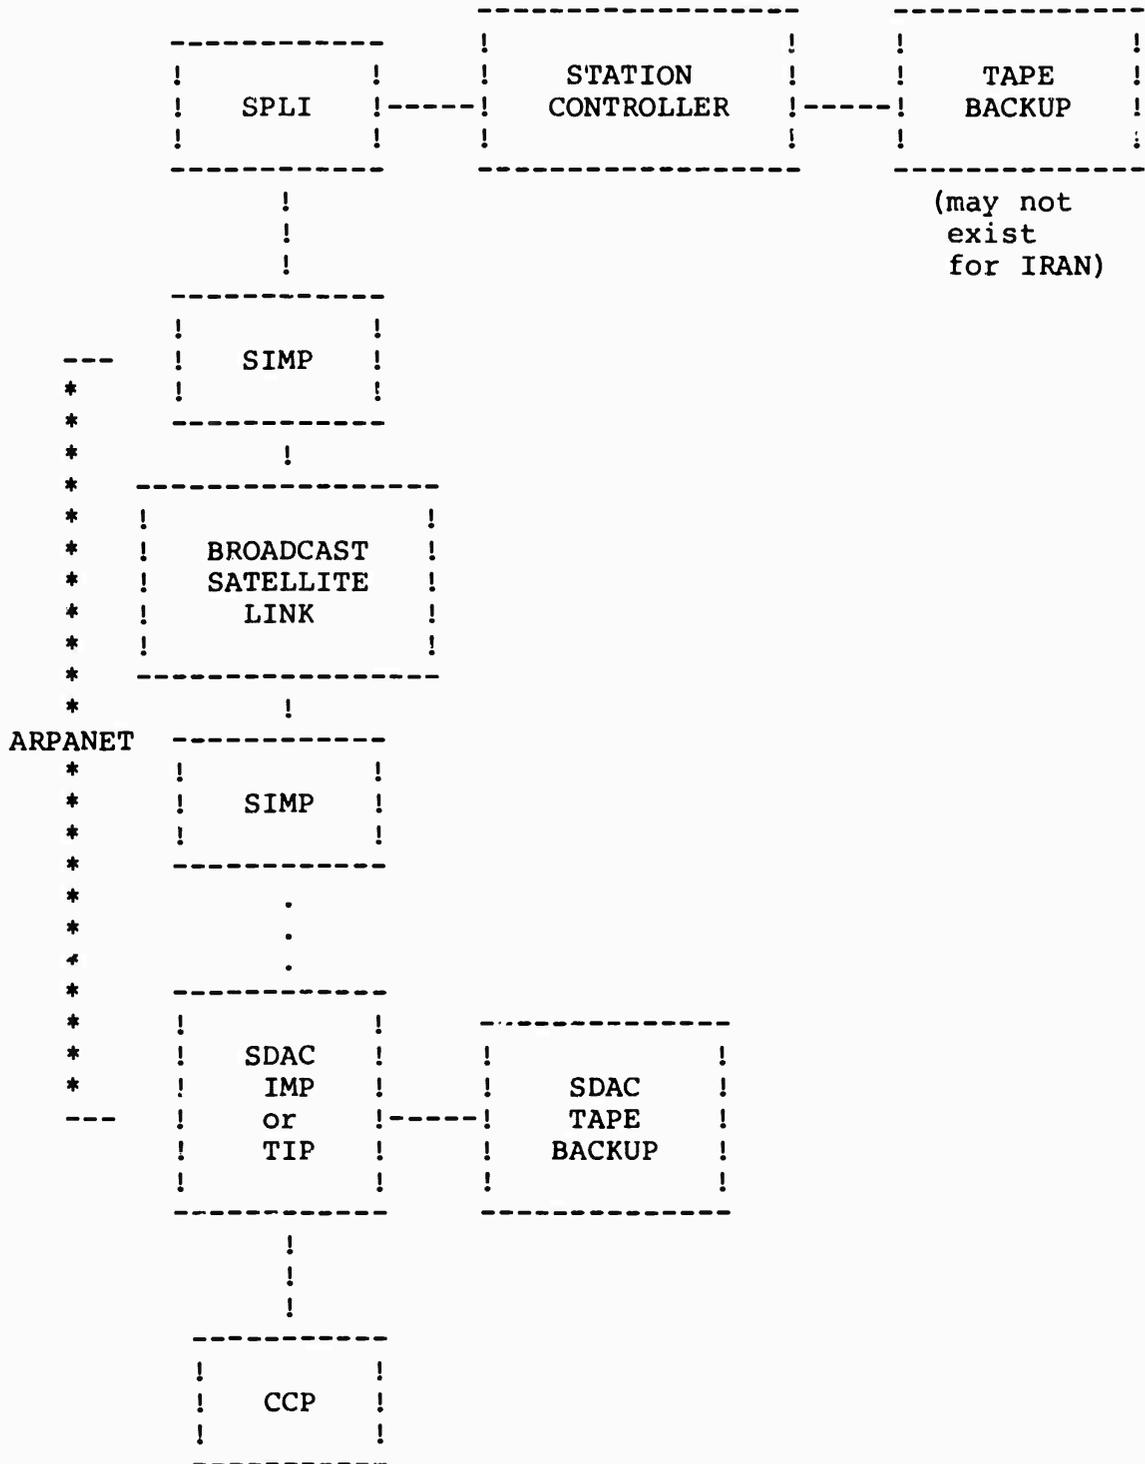


Figure 1. Site-CCP Data Link for Sites Connected to a SPLJ

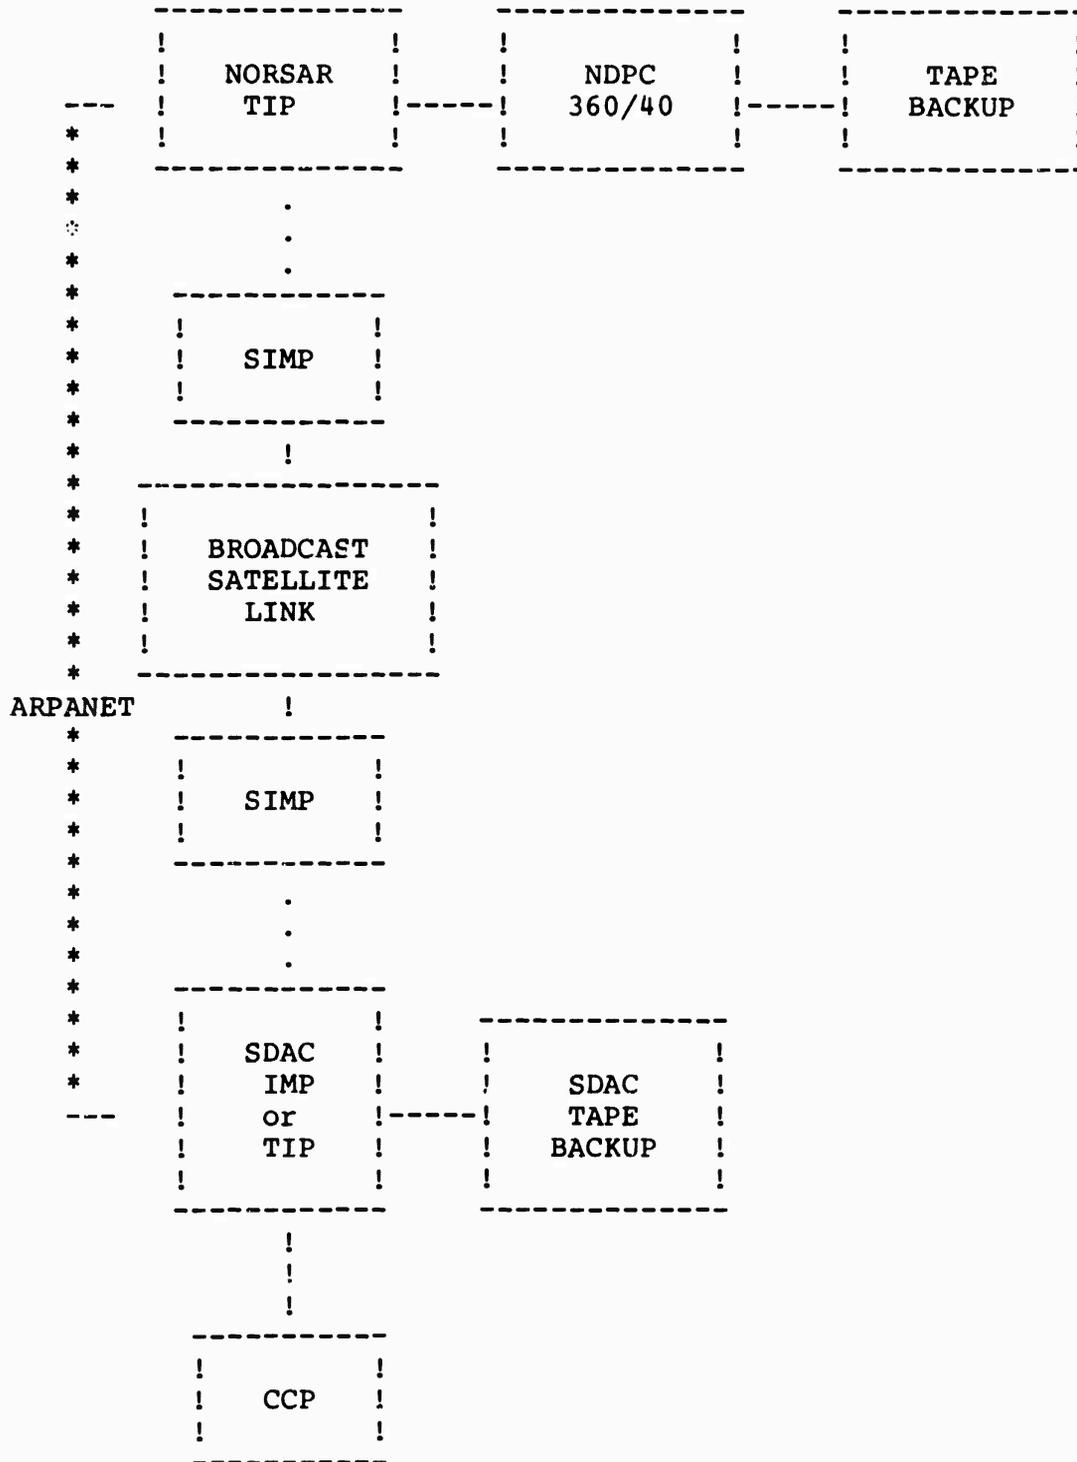
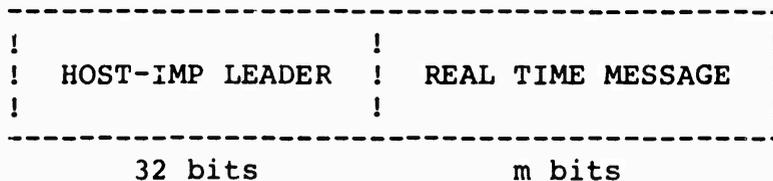


Figure 2. Site-CCP Data Link for NOR SAR

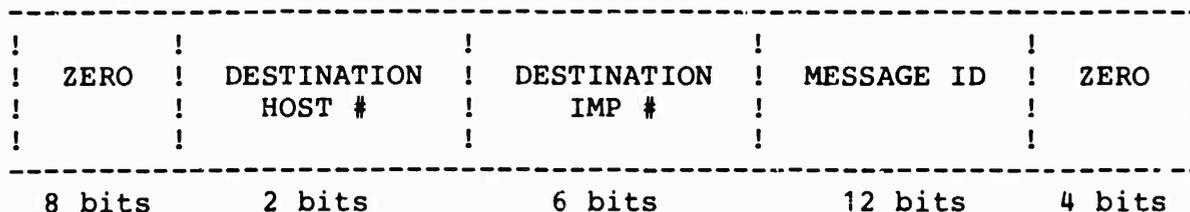
2. Real-Time Data Transmission Formats

The format for ARPANET transmission of real-time data between the overseas seismic sites and the CCP is shown in figure 3. It consists of a 32-bit Host-IMP leader [2] followed by the particular real-time message. The ARPANET employs RFNM (Ready For Next Message) and Incomplete Transmission messages for reliable, efficient communication. The message ID field in the leader is constrained to be nonzero (see section 3) and is used to allow up to four messages to be transmitted concurrently along an individual data path. This "pipelining" procedure decreases the time required to empty an output buffer that has been filling due to a data link outage (see section 4.1). The formats of the particular real-time messages are exactly those specified to be used by the array sites and are described in the following sections.



(a) Format of Site-CCP Real-Time Messages

(where: m=4800 from KSRS and SITE II
 m=816 from IRAN
 m=2400 for NORSAR)



(b) Host-IMP Leader for Real-Time Messages

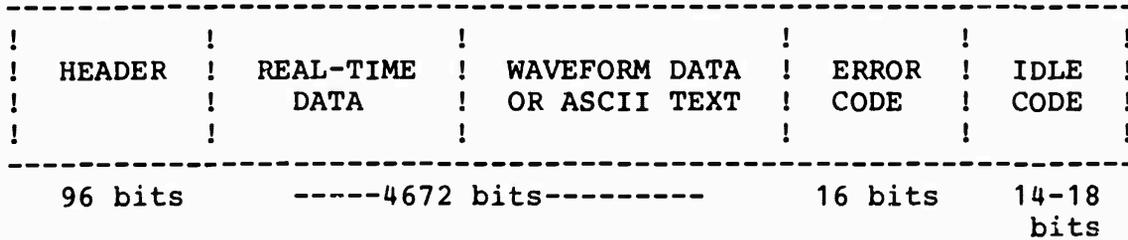
Figure 3. Format of Site-CCP Real-Time Messages

2.1 Real-Time Data from KSRS and SITE II to the CCP

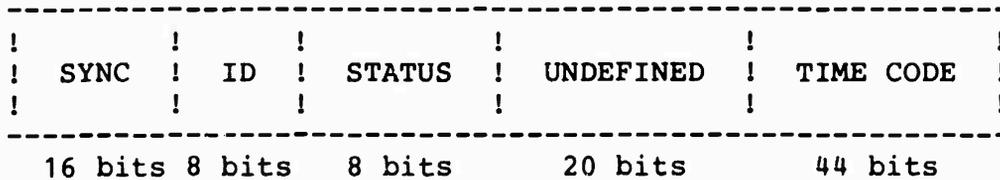
The real-time messages from the station controllers at KSRS and SITE II have the same format and are described in detail in [4]. Data frames are 4800 bits plus or minus 2 bits in length and will be transmitted by the station controller continuously at a rate of 1 frame per second. The mode of transmission will be synchronous.

As shown in figure 4, each data frame consists of a 96-bit header, a real-time data portion, a block data portion, and a 32-bit plus or minus 2 bits portion which is hardware generated. The header consists of 5 fields. A 16-bit message SYNC field identifies the start of each 1 second data frame. The 8 bits of the ID field constitute a single ASCII character used to identify the station controller. The 8 bits of the status field individually convey information, primarily relating to requested data in the block data portion of the data frame. Twenty bits in the middle of the header are currently undefined. The last field of the header is the 44-bit time code field which contains the BCD time in eleven 4-bit subfields (year-tens, year-units, day-hundreds, ..., second-units).

The real-time data portion has two parts. The first part is the short period data field. This field is made up of 20 frames (corresponding to the 20 Hz sampling rate) each of which has NSP 12-bit data samples (NSP = number of short period channels). The long period data field consists simply of NLP 16-bit data samples (NLP = number of long period channels). The data format constrains



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



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

Figure 4. Real-Time Message Format from KSRS and SITE II

the real-time data portion to be less than 4560 bits. In addition, NSP must not exceed 19 and NLP must not exceed 30.

The block data portion of the one second data frame provides a field in which specific requested data can be transmitted. Requested waveform data can be returned as a sequence of either 12-bit or 16-bit samples. ASCII data can also be returned as a sequence of 8-bit characters. Typically, the data associated with a particular request will be multiplexed over a sequence of one second frames because of the 4672-bit limit for both the real-time data and the block data together in an individual message. The block data is described in more detail in [4].

The error and idle codes are hardware generated by the 4800 bps modem interface on the station controller. The error code provides a polynomial-type error check mechanism. The idle code consists of a sequence of 14 to 18 zeros. The flexible length permits variations between the clock in the 4800 bps modem and the time standard of the station controller.

2.2 Real-Time Data from IRAN to the CCP

The format of the once-per-second real-time messages from IRAN is specified in figure 5. This format is derived from Addendum No. 1 to the Iranian Long Period Array Design Review [5].

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	0YYDDDDHHMMSS	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 5. Real-Time Message Format from IRAN

2.3 Real-Time Data Messages between NORSAR and the CCP

Real-time data is transmitted in both directions of the NORSAR-CCP data link. Figure 6 specifies the format of the real-time messages from NORSAR and figure 7 gives the format of the real-time messages sent to NORSAR by the CCP. These formats are essentially the same as those given in [6] except that the time specifications have been changed to 36-bit abbreviated station controller time codes [4], each consisting of 9 BCD 4-bit subfields (day-hundreds, day-tens, ..., second-units).

Bits	Contents	Description
0000-0031	Field 1	Control Characters
0032-0067	Field 2	Time (36-bit abbreviated format: DDDHHMMSS)
0068-0211	Field 3	NORSAR Detection Log Reduction Groups
0212-0291	Field 4	NORSAR LP Status and Repeat Indicators
0292-1347	Field 5a	NORSAR LP Data
1348-1395	Field 5b	NORSAR SP Channel Identification
1396-1875	Field 5c	NORSAR SP Data
1876-2259	Field 6	NORSAR Offline Results
2260-2323	Field 7	Program Coordination Data
2324-2355	Field 8	Control Characters
2356-2399	Field 9	Spare (encoded as zeros)

Figure 6. Real-Time Message Format from NORSAR

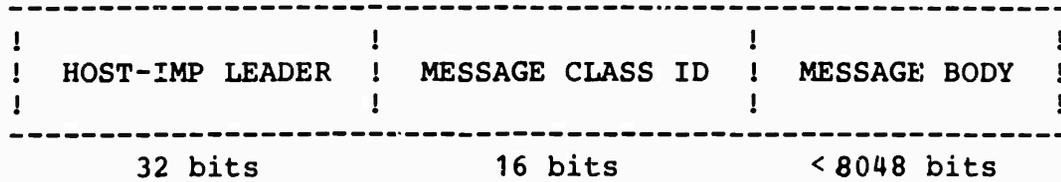
Bits	Contents	Description
0000-0031	Field 1	Control Characters
0032-0143	Field 2	LASA Signal Arrival Queue File Entry
0144-0179	Field 3	LASA Time (36-bit abbreviated format: DDDHHMMSS)
0180-0235	Field 4	LASA LP Status and Repeat Indicators
0236-1051	Field 5	LASA LP Data
1052-1087	Field 6	ALPA Time (36-bit abbreviated format: DDDHHMMSS)
1088-1151	Field 7	ALPA LP Status and Repeat Indicators
1152-2063	Field 8	ALPA LP Data
2064-2271	Field 9	LASA Off-line Results
2272-2367	Field 10a	Program Coordination Data
2272-2343	Field 10b	SP Data Request
2368-2399	Field 11	Control Characters

Figure 7. Real-Time Message Format to NORSAR

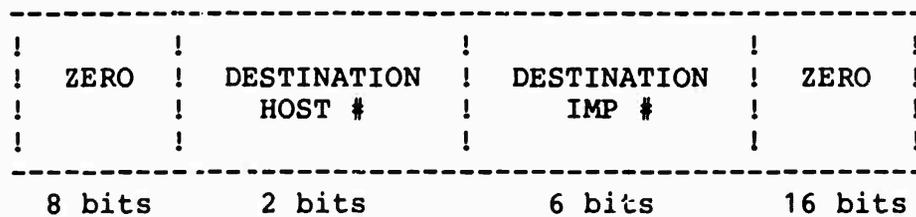
3. Format for Control Messages

The ARPANET message format for communication of commands and operator messages between the CCP and the Sites is shown in figure 8. This format was designed so that the protocol will use a Host-level acknowledgement scheme for reliable communication between the CCP and the Sites. It consists of a special 32-bit Host-IMP leader [2], a 16-bit message class identifier, and a variable length message body. This Host-IMP leader is special in so far as only its Destination Host Number field (bits 9-10) and Destination IMP Number field (bits 11-16) are ever nonzero. These control messages are distinguishable from real-time data messages by the fact that the message ID in the Host-IMP leader of the real-time messages is nonzero, while the message ID for control messages is always zero. The formats for each of the various classes of messages, shown in figures 9 and 10, are described below. The final sixteen bits in the message body contain a checksum for the entire message. A message is acknowledged only if it is successfully error-checked by the receiver. If an end does not receive such a Host-level acknowledgement within a specified time-out period after it sends a message, it will retransmit the message.

- * Class 0 messages are sent from the CCP to either the KSRS SPLI or the SITE II SPLI containing command data to those Sites. The first forty-eight bits in each Class 0 message body contain a 44-bit time code identifier. The structure of this time code is identical to that used by the station controllers [4] and



(a) Format of CCP-SPLI Control Messages



(b) Host-IMP Leader for Control Messages

Figure 8. Format of CCP-SPLI Control Messages

Message Class ID	Definition of Fields in Message Body	Field Size (bits)	Message Interpretation
0	Undefined	4	Commands to Station Controller
	Time Code identifier	44	
	Message to KSRS or SITE II	-	
	Checksum	16	
1		0	HELLO
2		0	I-HEARD-YOU
3	Undefined	4	Message Acknowledgement
	Time Code of acknowledged message	44	
	Checksum	16	
4	Undefined	4	IRAN SPLI to CCP Operator Message
	Time Code identifier	44	
	ASCII Text	-	
	Checksum	16	

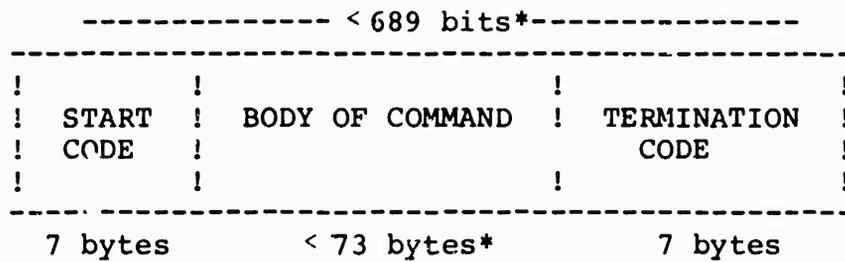
Figure 9. Classes of Control Messages from the CCP to the SPLI

Message Class ID	Definition of Fields in Message Body	Field Size (bits)	Message Interpretation
1		0	HELLO
2		0	I-HEARD-YOU
	Undefined	4	Message
3	Time Code of acknowledged message	44	Acknowledgement
	Checksum	16	
	Undefined	4	
	Time Code identifier	44	CCP to
4	ASCII Text	-	IRAN SPLI Operator Message
	Checksum	16	

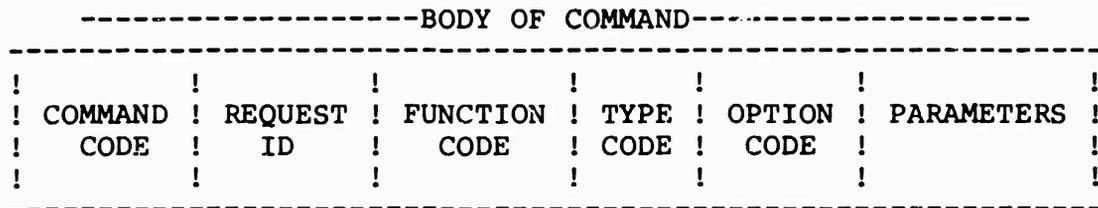
Figure 10. Classes of Control Messages from the SPLI to the CCP

consists of eleven BCD 4-bit subfields (year-tens, year-units, day-hundreds, ..., second-units). Following the time code identifier is the command message which is described in section 3.1 and shown in figure 11.

- * Class 1 (HELLO) and Class 2 (I-HEARD-YOU) messages have no body and are exchanged periodically to detect data link outages as described in section 4.
- * Class 3 (Acknowledgement) messages are sent whenever an end successfully receives command or operator messages. Each Class 3 message includes a 48-bit field which specifies the 44-bit time code identifier of the message being acknowledged. The final sixteen bits contain a checksum for the preceding forty-eight bits of this Class 3 message which is being sent. If an end does not receive a correct Class 3 message within a specified time-out period after it sends a message, it retransmits the message.
- * Class 4 messages are used to send operator messages between the IRAN SPLI and the CCP. As for Class 0 messages, the first forty-eight bits in each Class 4 message body contain a 44-bit time code identifier and the final sixteen bits contain a checksum.



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



(ASCII CHARACTERS)

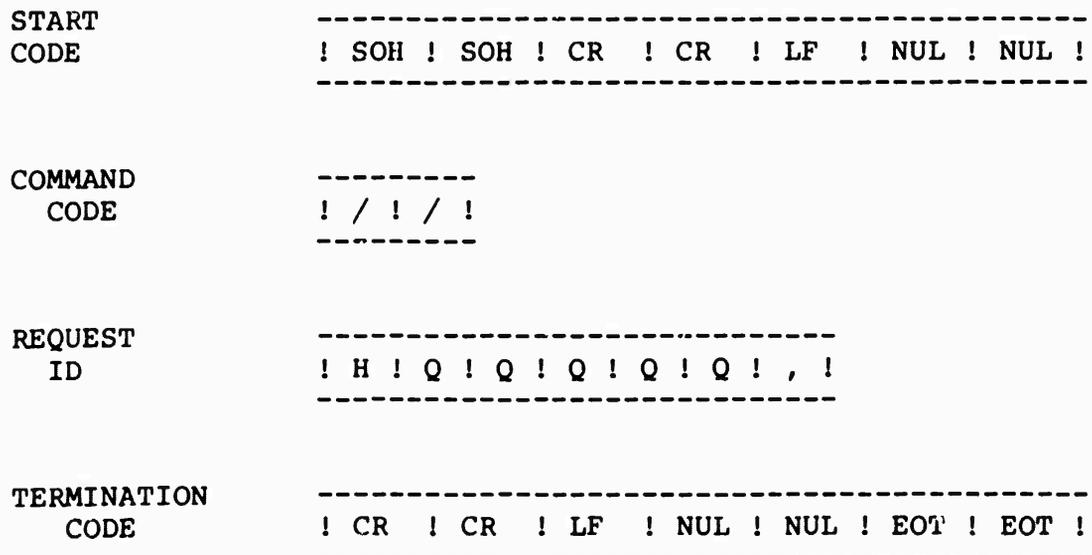


Figure 11. Command Message Format to KSRS and SITE II

3.1 Commands to KSRS and SITE II

The command messages to the station controllers at KSRS and SITE II have the same format and are described in detail in [7]. The length of each command will be less than 86 ASCII characters except for the "MESS" (message) command. The "MESS" command can contain a text message of up to 150 lines of 72 characters each, which is equal to 86400 bits. The restriction that an ARPANET message must be no longer than 8095 bits implies that lengthy "MESS" commands will be transmitted by using a sequence of as many as eleven Class 0 messages. The Host-level acknowledgement protocol, however, prevents these messages from arriving out of order. Each command will be sent as a stream of ASCII characters. The SPLI to station controller interface will be consistent with the current 75 bps modem data link.

As shown in figure 11, each command consists of a 7 character start code, a body, and a 7 character termination code. The request ID in the command body is a 7 character field. The first 6 characters of this field are included as part of the block data portion of the station controller to SPLI data frames (see section 2.1), so that a command and its corresponding response can be associated with one another. The first of the six characters must be "H". The remaining five are arbitrary, but must be alphanumeric.

The function code field specifies the operation to be performed as a result of the command. The type and option code fields may not

be present in a command. However, if they are present, they specify the operation of the command in more detail. The parameter field of a command either specifies parameters to be modified by the command or, in the case of the "MESS" command, includes the text message that is to be transmitted. Like the type and option code fields, the parameter field may not be present for every command.

3.2 Operator Messages between the IRAN SPLI and the CCP

As mentioned earlier, a Class 4 message body will contain an operator message from either the IRAN SPLI to the CCP or vice versa. Taking into account the Host-IMP leader, the message class identifier, the time code identifier, and the checksum (112 bits in all), each Class 4 message may contain up to $8095 - 112 = 7983$ bits, i.e. 997 characters of ASCII text. Longer operator messages can be sent as a succession of Class 4 messages. These messages will be received in the proper order on account of the Host-level acknowledgement protocol (see section 3.1).

4. Detailed Operation of the Communications Link

4.1 Transmission of Real-Time Data

The seismic sites will generate real-time data blocks at a rate of one per second. These blocks are buffered in a 10 second buffer. If this buffer is full when a new data block arrives, the oldest record in the buffer is replaced by the new data. Output from this source buffer to the network should be initiated whenever there is data to be sent. Normally, this will be at a rate of once per second. When a RFNM (Ready For Next Message) for a previously sent message is received at the source, the buffer allocated to that data block can be released (if it is still assigned). If an Incomplete Transmission is received in response to a previously sent message, that message should be retransmitted by the site if it is still available in the source buffer.

At the destination (the CCP) there are two analogous processes which manipulate the destination buffer. The process which receives real-time data from the network discards messages which arrive late, i.e., that cannot be delivered with the required 10 second delay. Data with a bad software checksum may also be discarded. The checksum is included primarily for end-to-end error detection, but would also allow easy implementation of end-to-end acknowledgement in the future (see section 4.3). Time-of-day clocks at the array sites and the CCP are both calibrated to an absolute time standard (such as WWV) so that the source and destination can remain

synchronized.

The destination buffer is filled with 10 seconds of data when the data link is initialized. The source buffer only fills up due to network delays resulting from retransmissions, channel outages, lost messages (or acknowledgements), etc. When more than one message is available for transmission at the source, several message IDs will be used so as to clear out the source buffer, i.e., catch-up, as quickly as possible. This "pipelining" procedure improves the network throughput for real-time data. The IMP subnetwork, however, imposes a constraint on the throughput which is possible. In particular, only 4 messages can be transmitted prior to receiving a RFNM for the first message.

If the network delays are short, the 10 seconds of buffered data will be sufficient for lossless real-time communication with a fixed delay. Some types of errors, however, can cause gaps to occur in the data. Loss by the ARPANET of a part of a message, for example, can cause a gap of up to 60 seconds to occur while the source IMP times out the associated acknowledgement. The only way around this problem is to buffer up to one minute of data both at the source and the destination. This would, of course, imply that all data would be delivered 60 seconds late rather than 10 seconds late. Since this type of error is expected to be rare (probably less than once per day), the added delay and cost of buffering were considered inappropriate. This data should always be available on tape and could be accessed via the network if desired.

4.2 Flow of Commands and Operator Messages

In contrast to the real-time messages, the commands and operator messages employ a reliable communication protocol which requires message acknowledgement or timeout and retransmission of unacknowledged messages. The Class 3 Acknowledgement message for a particular message is sent to the sender by the receiver to indicate that the received data arrived correctly as evidenced by recomputing the associated checksum. If the data is lost or arrives with a detectable error, no Host-level acknowledgement is returned, the sender eventually does a timeout, and the message is retransmitted. A 60 second timeout would probably be appropriate here.

In addition to the specification of reliable transmission for commands and operator messages, these control messages also require a simple implementation of flow control not required for real-time data. In each direction, control messages are sent over a single logical channel (the Host-IMP Leader message ID is zero--see section 3). This implies that no pipelining is possible for these messages. In particular, a RFNM must be received for the previous message (in a given direction) before the next message (in that direction) can be transmitted. Although this scheme limits the throughput for transmission of commands and operator messages, it also has the positive effect of minimizing the buffer space required for messages in transit.

The CCP or Site will eventually receive either a RFNM or an Incomplete Transmission message from its IMP or SIMP for each of the

messages which it transmits [2]. These Host-IMP messages are discarded because of the Class 3 Host-level acknowledgement procedure. Although Incomplete Transmission messages could be used to improve throughput of control messages, the added complication was judged inappropriate.

4.3 The Question of End-to-End Acknowledgement of Real-Time Data

The communication protocol described above does not include any mechanism for Host-to-Host acknowledgement of real-time data messages. A message is assumed to be delivered correctly by the source when a RFNM from the destination is received. What the RFNM actually implies is that the first packet of the associated message has been transmitted to the Host-IMP interface by the destination IMP. Whether the rest of the message made it through the interface and was received correctly by the destination Host is not known. To add error checking to the overall communication path, a Host-to-Host acknowledgement scheme for real-time messages similar to that used for control messages would be required. A Host-level acknowledgement message would be used for this purpose. The time code on the real-time messages could provide the unique identifier required for handling messages arriving out of order and duplicate messages.

One unfortunate side effect of the end-to-end acknowledgement of real-time data, however, is that it increases the data link vulnerability to long outages due to lost messages. If a Host-level

acknowledgement or the RFNM for a Host-level acknowledgement were lost, the system would be stuck for up to 60 seconds while the CCP IMP timed out the RFNM. Since the positive acknowledgement creates this problem and adds to the complexity of the real-time data protocol, it was not included in the design. Such a mechanism could be added at some later time if it were deemed appropriate.

4.4 Error Recovery and Initialization

To decide that the data link between the CCP and the Site is working properly, a scheme is used analogous to that used with a very distant Host, the VDH-to-IMP circuit test procedure [2]. Every R seconds both the CCP and the Site independently send each other a Class 1 (HELLO) message. Whenever the CCP or the Site receives a HELLO message, it must respond with highest priority by sending the other a Class 2 (I-HEARD-YOU) message. The I-HEARD-YOU is a Host-to-Host acknowledgement of the corresponding HELLO. If either end of the data path sends more than T consecutive HELLOs without receiving an I-HEARD-YOU acknowledgement, it declares the Site-CCP data path to be dead. After declaring the path dead, it does not send or receive any messages for $2RT$ seconds to allow the other party also to declare the path dead.

After an end waits $2RT$ seconds, it attempts to reinitialize the path. This is done by sending only HELLO messages every R seconds until X consecutive HELLOs have been acknowledged by I-HEARD-YOUs. After X HELLOs in a row have been acknowledged, the data path is

declared alive. After the data link has been declared alive, regular messages can be sent in addition to the periodic HELLO/I-HEARD-YOU pairs.

The value of R is initially 2, the value of T is 5, and the value of X is 5. The Network Control Programs for the CCP and the Site should be designed so that it is easy to change these parameters.

Since the CCP can be on either the SDAC IMP or the SDAC TIP, the Site must be prepared to try both Host addresses when trying to bring up the data path. It can do this by sending the HELLOs to both Host addresses simultaneously.

At system startup the data path will be assumed to have been declared dead, and the procedure for waiting 2RT seconds before sending HELLO messages will be used for initialization.

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] BBN Report No. 1822, Specifications for the Interconnection of a Host and an IMP, April 1973.
- [3] McKenzie, A., Host/Host Protocol for the ARPA Network, NIC 8246, January 1972.
- [4] Texas Instruments Inc., Specification for Station Controller to HDT Communication Formats, Drawing No. 138658.
- [5] Texas Instruments Inc., Iranian Long Period Array Design Review Addendum No. 1, April 1, 1974.
- [6] Teledyne Geotech, REF 110S,N -ISRSPS Program Reference Manual: Section 1-3-2, Rev. 8/31/71; Section 1-4-3, Rev. 6/08/72.
- [7] Texas Instruments Inc., Specification for Station Controller Command and Response Formats, Drawing No. 140521.

Table of Contents

1. Overall Link Operation.....	2
2. Data Formats.....	6
2.1 Waveform Data and Status Information.....	6
2.2 File Directory Control Message.....	12
3. CCP-SIP Messages.....	16
4. SDAC Backup Recording.....	23
5. Detailed Operation of the Data Link.....	24
5.1 Normal Operation.....	24
5.2 Error Recovery and Initialization.....	26
6. References.....	28

List of Figures

- 1. CCP-SIP Data Link.....4
- 2. Logical Form of the One Second Status/Data Message.....7
- 3. Detailed Structure of the Time Code Field.....8
- 4. Format of Status Information in One Second Status/Data Message.10
- 5. Format of Waveform Data in One Second Status/Data Message.....11
- 6. Format of the File Directory Control Message.....13
- 7. CCP-SIP Message Format.....17
- 8. Classes of Messages from the CCP to the SIP.....18
- 9. Classes of Messages from the SIP to the CCP.....19

1. Overall Link Operation

A protocol for the on-line transmission of seismic data between the Communications and Control Processor (CCP) at SDAC and the Seismic Input Processor (SIP) at CCA is described in this document. This protocol is consistent with both the initial specification of the Datacomputer files presented by R. Lacoss [1] and the subsequent file descriptions developed by E. B. McCoy [2].

Both the CCP and SIP are interfaced as Hosts on the ARPANET and communicate at the Host-to-Host level although not via the standard Host-Host Protocol [3]. As shown in figure 1, both the SIP and the Datacomputer are Hosts on the CCA TIP. The SIP will provide buffering to handle outages of the Datacomputer and will permit higher rate transfers to the Datacomputer. The buffering at the SIP is provided by a large disk with a capacity of on the order of 24 hours of data. Normally the SIP will buffer on the order of one hour of data before transferring it to the Datacomputer at a rate of about 100 kbaud. Larger amounts of data will be buffered, of course, when the Datacomputer crashes or when it is scheduled for preventive maintenance.

When the SIP successfully transfers the data from its disk to the Datacomputer it will send a message that will be stored in a table in the CCP. These transfers will occur typically every hour and so a table of SIP-to-Datacomputer transfer status for a day or two will be quite small.

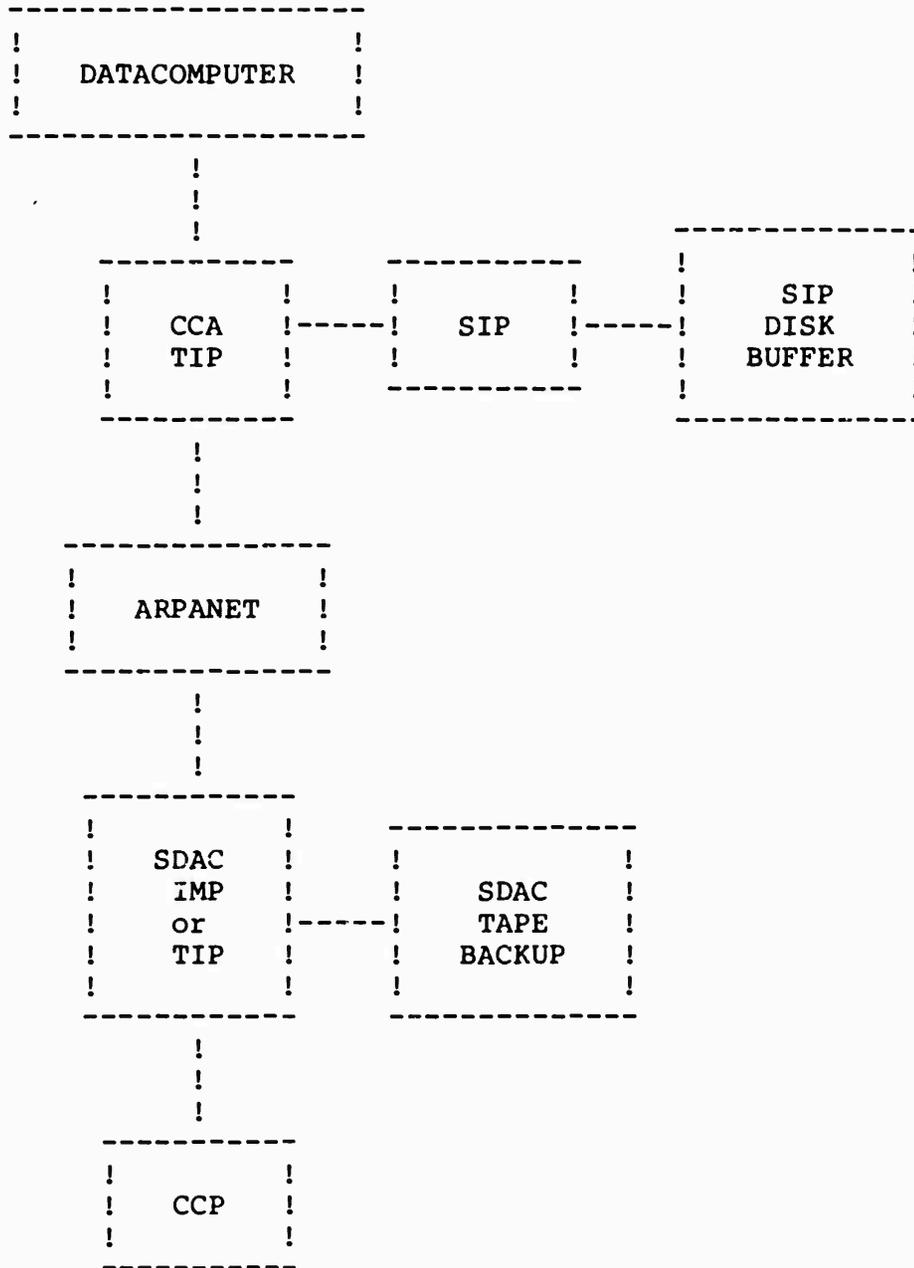


Figure 1. CCP-SIP Data Link

At the SDAC end of the data link, buffering must be provided by the CCP to handle the uneven data flow caused by the asynchronous nature of the ARPANET. Tape backup also must be available at SDAC in case of a network outage or failure of either the SIP or Datacomputer. It is recommended, however, that a continuous backup recording capability be provided by SDAC. Everything which is transmitted to the SIP would be recorded at SDAC and kept for a week before the tapes are recycled. These tapes would be used off-line at the Datacomputer to fill in any data missed during data link failures.

Three kinds of data are to be sent from the CCP to the SIP:

- * Waveform data consisting of Long Period (LP) and Short Period (SP) real-time seismic data and beams will be sent in one second data blocks arranged by site.
- * Status information concerning the data transmitted from each site will be sent to the SIP prefixed to the waveform data.
- * A file directory control message, taking effect either immediately or at a preset time (e.g. midnight) and of course at initialization, will specify the interpretation of subsequently received data blocks (e.g. which channels are being sent and the files into which they should be placed).

A detailed description of the formats of these messages will be presented in the following section.

2. Data Formats

2.1 Waveform Data and Status Information

Once each second, the CCP sends to the SIP a message containing one second of seismic waveform data preceded by status information concerning that data from each of the sites. A variable number of raw data channels and beams may be sent from each site, so the message must be interpreted according to the most recent file directory control message (discussed below) received by the SIP from the CCP. Each one second data block is formed with a given delay for each of the sites, but the differential delay for each channel and beam within a site must again be interpreted according to the current file directory control message.

The format in which the waveform data is sent to the SIP by the CCP is shown in figure 2. The data is arranged by site so that this data format can be used to send specified data from the CCP to the SDAC 360/40. For sites not sending both long period data and short period data the corresponding data fields (described below) will not be used. Within each site field this status/data message logically consists of five fields: the time code; the status of the long period data; the status of the short period data; the long period data; and the short period data.

* Each one second data block from a site is prefixed with a time code. The structure of this field (shown in figure 3) is identical to the 44-bit time code used by the station

!	!	!	!	!	!	!	!					
!	LASA	!	NORSAR	!	KSRS	!	SITE II	!	IRAN	!	ALPA	!
!	!	!	!	!	!	!	!	!	!	!	!	!

(a) One Second Status/Data Message Compiled by Site

-----SITE-----										
!	!	!	!	!	!					
!	TIME CODE	!	LP STATUS	!	SP STATUS	!	LP DATA	!	SP DATA	!
!	!	!	!	!	!	!	!	!	!	!

(b) Structure of the One Second Status/Data Message from a Site

Figure 2. Logical Form of the One Second Status/Data Message

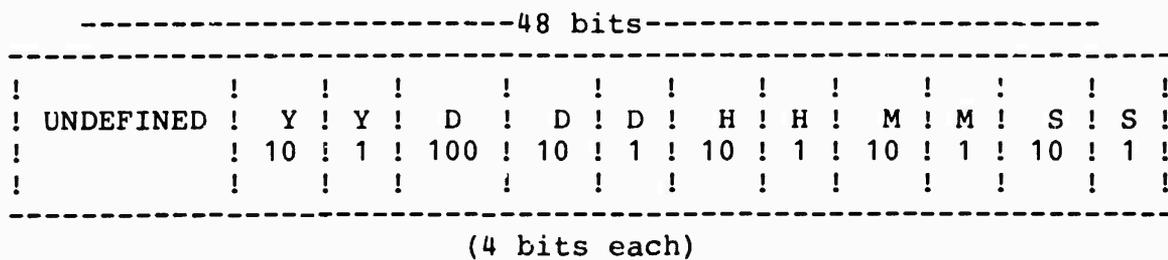


Figure 3. Detailed Structure of the Time Code Field

controllers [4] and consists of eleven BCD 4-bit subfields (year-tens, year-units, day-hundreds, ... , second-units).

* The status information for the one second block of long period data used from each site will specify:

** the availability of data from an individual channel or beam

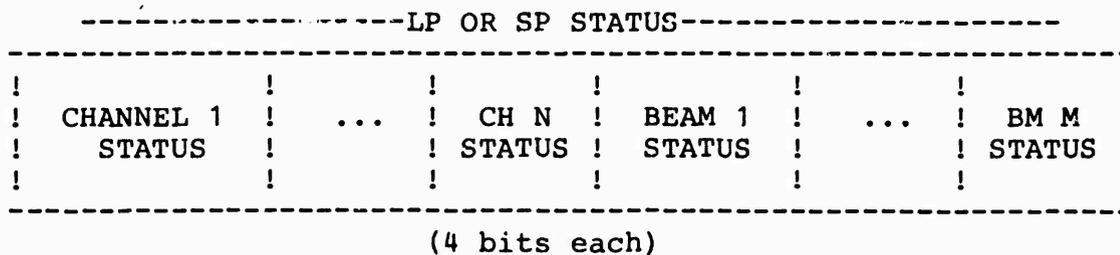
** that the data for a particular channel or beam is bad

** that a particular seismometer is either being calibrated or otherwise tested

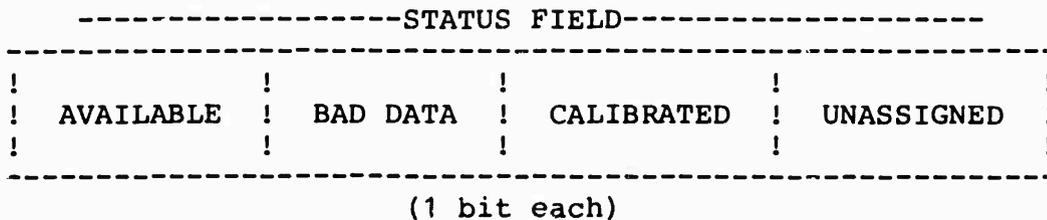
To indicate these conditions, a 4-bit status field (see figure 4) is assigned to every beam and channel that is being used from a site as specified in the file directory control message described below. Three bits are used as flags to specify the above conditions, and the fourth remains unassigned.

* The status information for the one second block of short period data used from each site will be given in the same format described above for the long period data status information.

* The format of the one second block of Long Period data used from a site is shown in figure 5. It consists of one 16-bit data sample from every selected seismic channel, followed by specified 16-bit beams formed for that site. The

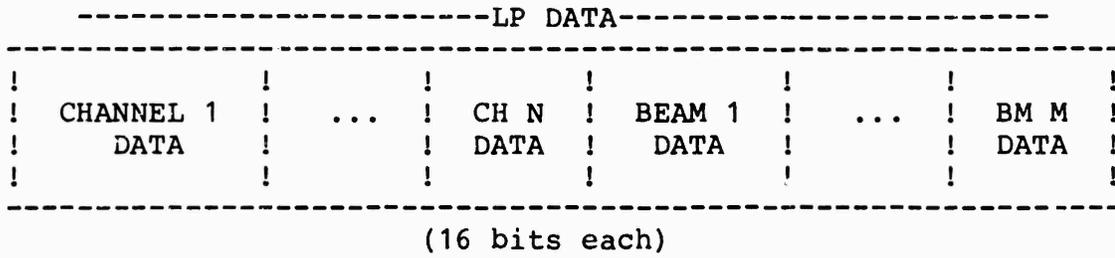


(a) The Ordering of the Status Fields given for either Long Period or Short Period Channels and Beams

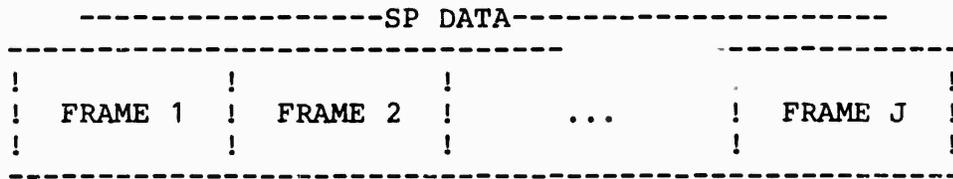


(b) Definition of an Individual Channel or Beam Status Field

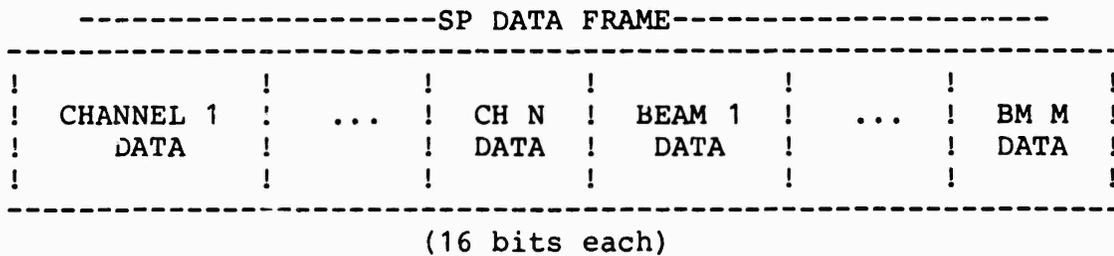
Figure 4. Format of Status Information in One Second Status/Data Message



(a) One Second of Long Period Data from an Individual Site



(b) One second of Short Period Data Frames from an Individual Site
(J samples per second, where J=20 for all sites other than LASA
where J=10)



(c) One Frame of Short Period Data from an Individual Site

Figure 5. Format of Waveform Data in One Second
Status/Data Message

interpretation of the ordering of these data samples is done using the file directory control message described below.

- * The format of the one second block of Short Period data used from a site is also shown in figure 5. It is divided into J time frames, where J is the number of samples given each second by that site. J is equal to 20 for all sites other than LASA where J equals 10. Each frame consists of one 16-bit data sample from every selected seismic channel at the site, followed by specified 16-bit beams formed for that site. The interpretation of a frame of data from a site is accomplished using the file directory control message described below.

2.2 File Directory Control Message

During initialization of the data link or at any time at the start of a new file, a file directory control message is used for three purposes: identifying the channels and beams contained in subsequent waveform data messages; specifying the differential delays of these channels and beams with respect to the standard times of their sites; and designating the Datacomputer file into which an individual data sample is to be entered.

- * Each channel and beam in the entire seismic network is assigned a unique 16-bit identifier. (Using four 4-bit BCD fields, for example, this would accommodate 10,000 channels and beams.) The file directory control message contains identifier-labeled 48-bit subfields (see figure 6) corresponding to each of the

!	!	!	!	!	!	!	!
!	LASA	!	NORSAR	!	KSRS	!	SITE II
!	!	!	!	!	!	!	IRAN
!	!	!	!	!	!	!	ALPA
!	!	!	!	!	!	!	!

(a) Logical Form of the File Directory Control Message
Compiled by Site

-----SITE-----			
!	!	!	!
!	TIME CODE	!	# OF SUBFIELDS
!	!	!	FOR THIS SITE
!	!	!	!
!	!	!	LONG PERIOD
!	!	!	SUBFIELDS
!	!	!	!
!	!	!	SHORT PERIOD
!	!	!	SUBFIELDS
!	!	!	!
			(one for each LP channel and beam)
			(one for each SP channel and beam)

(b) Structure of the File Directory Control Message from
a Single Site

!	!	!	!
!	CHANNEL/BEAM ID	!	DIFFERENTIAL DELAY
!	!	!	DESTINATION FILE #
!	!	!	!
			(16 bits each)

(c) Definition of the File Directory Control Message Subfield
given for each Channel and Beam

Figure 6. Format of the File Directory Control Message

channels and beams included in the status/data message. For Long Period data, subfields are given corresponding to the channel and beam samples within each site field of the Long Period waveform data. For Short Period data, subfields are given for all the channels and beams sampled within one time frame from each site. This ordering is the same, of course, for every frame.

- * Each identifier-labeled subfield also contains a 16-bit differential delay for that channel or beam with respect to its site. (These delays should probably be given in twentieths of a second, and will most likely not exceed thirty seconds.)
- * The sixteen bits following the delay in each subfield are used as a file number by the SIP, specifying which Datacomputer file is to receive the data from this identified channel or beam. (Only ten conceptual seismic files at the Datacomputer have as of yet been specified [1].)

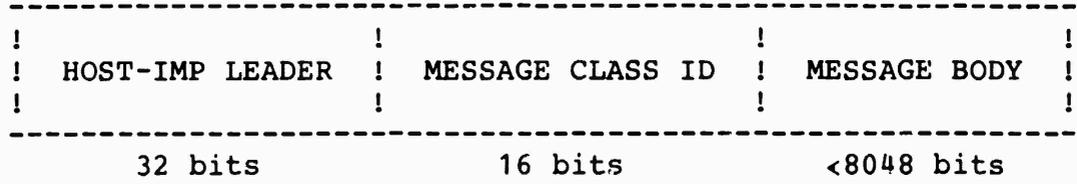
As for the status/data message described above, the file directory control message is arranged by site (see figure 6a). In order for the SIP to interpret this message correctly, the number of 48-bit subfields given for each site is specified in a 16-bit field following the time code in each site field (see figure 6b). The time code of when this message was sent is given in a 48-bit field (see section 2.1). Using this ordering-correspondence scheme, the number of seismic channels or beams referenced from each site may be easily altered. The choice of files into which specified data is

entered is also readily controlled. These changes should probably take effect at a preset time (e.g. midnight), but this protocol also provides for such changes to be made at any time.

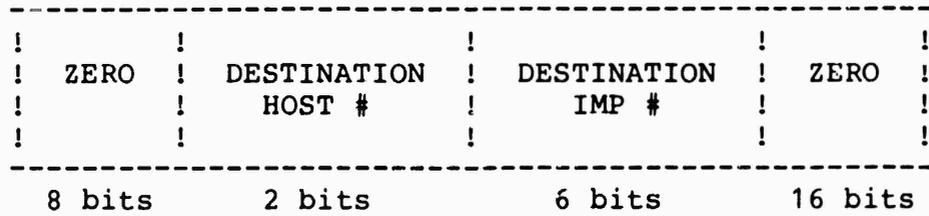
3. CCP-SIP Messages

The ARPANET message format for communication between the CCP and the SIP is shown in figure 7. This format was designed so that the protocol will use a Host-level acknowledgement scheme for reliable communication between the CCP and the SIP. It consists of a special 32-bit Host-IMP leader [5], a 16-bit message class identifier, and a variable length message body. This Host-IMP leader is special in so far as only its Destination Host Number field (bits 9-10) and Destination IMP Number field (bits 11-16) are ever nonzero. The formats for each of the various classes of messages, shown in figures 8 and 9, are described below.

The restriction that an ARPANET message must be no longer than 8095 bits necessitates the use of some scheme for partitioning longer messages that are sent from the CCP to the SIP. A straightforward approach is to divide each logical message into acceptably-sized pieces, giving all pieces a message label and a numerical ordering. This protocol employs the 44-bit time code (see section 2.1) of the message as the label, and uses the preceding 4-bit field to represent the piece number. The number of pieces which the SIP must wait to receive before reassembling the message is determined from the Class 4 message (see figure 8) previously received from the CCP. (The control messages themselves specify in each of their pieces the total number of pieces of which they consist.)



(a) CCP-SIP Message Format



(b) CCP-SIP Host-IMP Leader

Figure 7. CCP-SIP Message Format

Message Class ID	Definition of Fields in Message Body	Field Size (bits)	Message Interpretation
1		0	HELLO
2		0	I-HEARD-YOU
	Message Class ID of acknowledged message	16	
3	Undefined	4	Message
	Time Code of acknowledged message	44	Acknowledgement
	Checksum	16	
	Number of this piece	4	File Directory
	Time Code when this message was sent	44	Control
	Number of pieces in this message	8	Message for
4	Number of pieces in Status/Data message	8	Reinitialization
	Immediate File Directory Control Message	-	or Starting
	Checksum (last piece only)	16	New Files
	Undefined	4	Close Files---
			Use Next
5	Time Code	44	Class 4 Message
			To Start
	Checksum	16	New Files
	Number of this piece	4	File Directory
	Time Code when this message was sent	44	Control
	Number of pieces in this message	8	Message for
6	Number of pieces in Status/Data message	8	Starting
	Future File Directory Control Message	-	New Files
	Checksum (last piece only)	16	at Preset Time
	Number of this piece	4	
	Time Code	44	One Second
7	One Second Status/Data Message	-	Status/Data
	Checksum (last piece only)	16	Message

Figure 8. Classes of Messages from the CCP to the SIP

Message Class ID	Definition of Fields in Message Body	Field Size (bits)	Message Interpretation
	Undefined	4	Datacomputer
0	Time Code of last data transferred	44	Transfer
	Checksum	16	Completed
1		0	HELLO
2		0	I-HEARD-YOU
	Message Class ID of acknowledged message	16	
3	Undefined	4	Message
	Time Code of acknowledged message	44	Acknowledgement
	Checksum	16	

Figure 9. Classes of Messages from the SIP to the CCP

The final sixteen bits in the last piece of each message contain a checksum for the entire message. Each message is acknowledged by the SIP only after the entire message has been reassembled, error-checked, and correctly stored on the SIP disk. If the CCP does not receive such a Host-level acknowledgement within a specified time-out period, it will retransmit all the pieces.

Figures 8 and 9 depict the eight classes of messages used in CCP-SIP communication.

- * Class 0 messages are sent to the CCP by the SIP whenever it has successfully transferred the data from its disk to the Datacomputer. A Class 0 message contains the time code of the last data sample that was transferred to the Datacomputer by the SIP.
- * Class 1 (HELLO) and Class 2 (I-HEARD-YOU) messages have no body and are exchanged periodically to detect data link outages as described in section 5.
- * Class 3 (Acknowledgment) messages are sent to the CCP by the SIP whenever it successfully receives and incorporates messages (other than HELLO and I-HEARD-YOU) from the CCP. Each Class 3 message includes a 16-bit field which specifies the message class identifier of the message being acknowledged. The following forty-eight bits contain the 44-bit time code (see figure 3) of the acknowledged message. The final sixteen bits contain a checksum for the preceding sixty-four bits of this

Class 3 message which is being sent. If the CCP does not receive a correct Class 3 message within a specified time-out period after it sends a message, it retransmits the message.

Class 3 messages are likewise sent to the SIP by the CCP to acknowledge the receipt of Class 0 Datacomputer transfer status messages.

- * Class 4 messages are used for either reinitializing the data link or defining a new format for Datacomputer files. Following an outage or a Class 5 Close Files message (see below), the CCP must send a Class 4 message to the SIP. This message is labeled with the time code of when it was sent and contains information as to the size of subsequent status/data messages, as described above. It also includes the immediate file directory control message (see section 2.2) being used by the SIP. Only after the SIP successfully acknowledges this Class 4 message does the CCP continue data transmission.
- * Class 5 (Close Files) messages are sent by the CCP to inform the SIP of a change in either which data is being sent or the files into which data should be placed. When the SIP receives a Class 5 message it must immediately close all files. It will open new files to accept future data in a format specified by the Class 4 file directory control message it next receives. To insure reliability, the Host-level acknowledgement procedure requires that the CCP receive a Class 3 Acknowledgement message from the SIP before it sends the file directory control

message. Furthermore, this Class 6 message must be acknowledged to have been successfully incorporated by the SIP before the CCP continues data transmission according to this file directory control message.

- * Class 6 messages have exactly the same format as Class 4 messages, but are used by the SIP at a preset time (e.g. midnight) to update the current file directory control message being employed by the SIP. A Class 6 message takes effect at the beginning of a new file, and thereafter is used by the SIP in interpreting incoming messages.

- * Class 7 messages contain the one second status/data messages (see section 2.1), each piece being labeled with the time code specific to the one second data block which, together, they contain.

4. Backup Recording

Tape backup must be provided by SDAC, at least to handle possible outages of the CCP-SIP data link. It is recommended, however, that all the data sent to the SIP be saved at SDAC for a week on tape. This would insure the integrity of the data files despite any losses by the Datacomputer. These tapes could be mailed to CCA in order to update the Mass Store.

A "lost data scenario" might go as follows: Each second, all data sent to CCA is also recorded on tape at SDAC. When some data is lost due to a network outage, an indication of the lost data (start time, stop time) is presented to the CCP operator and is also recorded in a CCP lost data table. This table has sufficient space for the maximum number of outages that could possibly occur during one tape's worth of data. (If a tape holds 5 hours of data, one can impose the restriction that outages less than one minute apart cause the entire minute's worth of data to be marked, thus no more than 300 entries are required in the lost data table.) When a full tape is unmounted, the CCP operator can request to see the old lost data table and decide whether or not to send the recorded data (more likely a copy of it) to CCA.

5. Detailed Operation of the Data Link

5.1 Normal Operation

A one second status information and waveform data message, or a file directory control message, is produced by the CCP and placed in the output buffer to the SIP. It is held there for retransmission by the CCP if a correct acknowledgement for it is not received from the SIP within a specified time-out period. In the case of a data link outage, it will be stored on tape to prevent the loss caused by the overflowing of the output buffer. It is recommended that all the data also be saved on tape at SDAC for a week in order to protect against any losses incurred at the SIP end of the data link.

A checksum is computed for the message, and then it is divided into pieces (see section 3) which are sent over the ARPANET. To accomplish this, the Network Control Programs of the CCP and SIP must both implement the standard Host-IMP protocol described in detail in BBN Report No. 1822 [5]. When all the pieces are received by the SIP it reassembles the message. A checksummed acknowledgement (i.e. one containing a checksum on itself) for the message is sent to the CCP by the SIP only if no error has been detected by the message checksum and the data is safely stored on the disk. If the CCP does not receive a correct acknowledgement for the message within a specified time-out period, it retransmits all the pieces. Occasional duplicate messages received by the SIP do not cause any problem since if it is a data message, the data will

simply be recorded over the first copy.

The CCP will eventually receive either a RFNM (Ready for Next Message) or an Incomplete Transmission message from its IMP or TIP for each of the message pieces which it transmits [5]. These Host-IMP messages are discarded because of the Class 3 Host-level acknowledgement procedure. Although Incomplete Transmission messages could be used to improve throughput, the added complication was judged inappropriate. RFNM and Incomplete Transmission messages are likewise discarded by the SIP.

As indicated in section 4, the CCP keeps a table specifying those messages that are not successfully transferred to the SIP. This occurs when the data link is down for a period for which the CCP output buffer is not sufficient. The CCP process which writes messages to the output buffer is in this case required to overwrite the oldest message stored. Tape backup, however, prevents any loss.

When the SIP successfully transfers the data on its disk to the Datacomputer it sends the CCP an appropriate Class 0 message. These transfer status messages give the time code of the last data sample transferred and are recorded in a table in the CCP. It is probably not necessary to store more than a day or two of this kind of information, so the table will be quite small.

5.2 Error Recovery and Initialization

To decide that the data link between the CCP and SIP is working properly, a scheme is used analogous to that used with a very distant Host, the VDH-to-IMP circuit test procedure [5]. Every R seconds both the CCP and the SIP independently send each other a Class 1 (HELLO) message. Whenever the CCP or SIP receives a HELLO message, it must respond with highest priority by sending the other a Class 2 (I-HEARD-YOU) message. The I-HEARD-YOU is a Host-to-Host acknowledgement of the corresponding HELLO. If either end of the data path sends more than T consecutive HELLOs without receiving an I-HEARD-YOU acknowledgement, it declares the CCP-SIP data path to be dead. After declaring the path dead, it does not send or receive any messages for 2RT seconds to allow the other party also to declare the path dead.

After an end waits 2RT seconds, it attempts to reinitialize the path. This is done by sending only HELLO messages every R seconds until X consecutive HELLOs have been acknowledged by I-HEARD-YOUs. After X HELLOs in a row have been acknowledged, the data path is declared alive. After the end declares the path alive, a Class 4 (reinitialization control) message (see section 3) followed by regular data messages can be sent by the CCP, in addition to the periodic (every R seconds) HELLO messages. If a control message is received by the SIP before it has declared the link alive, the message should not be acknowledged.

The value of R is initially 2, the value of T is 5, and the value of X is 5. The Network Control Programs of the CCP and SIP should be designed so that it is easy to change these parameters.

Since the CCP can be on either the SDAC IMP or the SDAC TIP, the SIP must be prepared to try both Host addresses when trying to bring up the data path. It can do this by sending the HELLOs to both Host addresses simultaneously.

At system startup the data path will be assumed to have been declared dead, and the procedure for waiting 2RT seconds before sending HELLO messages will be used for initialization.

6. References

- [1] Lacoss, R., Main Conceptual Seismic Files on the Datacomputer, SDPCC, 19 March 1974.
- [2] McCoy, E. B., SDAC On-line Processing with Datacomputer, 29 April 1974.
- [3] McKenzie, A., Host/Host Protocol for the ARPA Network, NIC 8246, January 1972.
- [4] Texas Instruments Inc., Specifications for Station Controller to HDT Communications Formats, Drawing No. 138658.
- [5] BBN Report No. 1822, Specifications for the Interconnection of a Host and an IMP, April 1973.