

Bolt Beranek and Newman Inc.



①

85 - 0380

Report No. 5215

AD-A222 356

file
DTIC

Combined Quarterly Technical Report No. 27

SATNET Development and Operation
Pluribus Satellite IMP Development
Remote Site Maintenance
Internet Operations and Maintenance
Mobile Access Terminal Network
TCP for the HP3000
TCP for VAX-UNIX

November 1982

Prepared for:
Defense Advanced Research Projects Agency

DTIC
1990
CO

DISTRICT
Approved
Date

000
ZDAAAAA66238800

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Combined Quarterly Technical Report No. 27		5. TYPE OF REPORT & PERIOD COVERED Quarterly Technical 8/1/82 to 10/31/82
		6. PERFORMING ORG. REPORT NUMBER 5215
7. AUTHOR(s) J. F. Haverty		8. CONTRACT OR GRANT NUMBER(s) MDA903-80-C-0353 & 0214 N00039-82-C-0516 N00039-80-C-0664 N00039-81-C-0408
9. PERFORMING ORGANIZATION NAME AND ADDRESS Bolt Beranek and Newman Inc. 10 Moulton Street Cambridge, MA 02238		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS ARPA Order No. 3214
11. CONTROLLING OFFICE NAME AND ADDRESS Defense Advanced Research Projects Agency 1400 Wilson Boulevard Arlington, VA 22209		12. REPORT DATE November 1982
		13. NUMBER OF PAGES 60
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) DSSW Room 1D The Pentagon Washington, DC 20310 NAVELEX Washington, DC 20360		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) Computer networks, packets, packet broadcast, satellite communication, gateways, Transmission Control Protocol, UNIX, Pluribus Satellite IMP, Remote Site Module, Remote Site Maintenance, shipboard communications VAX, ARPANET, Internet.		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This Quarterly Technical Report describes work on the development of and experimentation with packet broadcast by satellite; on development of Pluribus Satellite IMPs; on a study of the technology of Remote Site Maintenance; on Internetwork monitoring; on shipboard satellite communications; and on the development of Transmission Control Protocols for the HP3000 and VAX-UNIX.		

DD FORM 1473
1 JAN 73

EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

Report No. 5215



COMBINED QUARTERLY TECHNICAL REPORT NO. 27

SATNET DEVELOPMENT AND OPERATION
PLURIBUS SATELLITE IMP DEVELOPMENT
REMOTE SITE MAINTENANCE
INTERNET OPERATIONS AND MAINTENANCE
MOBILE ACCESS TERMINAL NETWORK
TCP FOR THE HP3000
TCP FOR VAX-UNIX

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input checked="" type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By _____	
Distribution/	
Availability Codes	
Dist	Special
A-1	

November 1982

This research was supported by the Defense Advanced Research Projects Agency under the following contracts:

N00039-82-C-0516
MDA903-80-C-0353, ARPA Order No. 3214
MDA903-80-C-0214, ARPA Order No. 3214
N00039-80-C-0664
N00039-81-C-0408

Submitted to:

Director
Defense Advanced Research Projects Agency
1400 Wilson Boulevard
Arlington, VA 22209

Attention: Program Management

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

Table of Contents

1	INTRODUCTION.....	1
2	SATNET DEVELOPMENT AND OPERATION.....	2
2.1	Raisting Satellite IMP Installation.....	3
2.2	ARPANET Line #77.....	5
2.3	Software Maintenance Operations.....	7
2.4	Hardware Maintenance Operations.....	8
3	PLURIBUS SATELLITE IMP DEVELOPMENT.....	14
3.1	Network Monitoring by the NU System.....	19
3.1.1	Backbone Processes.....	21
3.1.2	Application Processes.....	23
4	REMOTE SITE MAINTENANCE.....	29
5	INTERNET DEVELOPMENT.....	30
5.1	Introduction.....	30
5.2	Gateway Installations.....	30
5.3	Gateway Throughput Collection.....	31
5.4	Gateway Status Processor.....	34
5.5	UCL Gateway.....	35
5.6	VAN Gateway.....	35
5.6.1	Access Control Table.....	36
5.6.2	X.25 Address Table.....	37
5.6.3	Datagram Processing.....	37
5.6.4	Virtual Circuit Management.....	38
5.6.5	X.25 Usage.....	39
5.7	Exterior Gateway Protocol.....	40
5.8	Gateway RFC.....	40
6	MOBILE ACCESS TERMINAL NETWORK.....	41
6.1	Error Detection on Packet Length.....	42
6.2	C/30 Ruggedization.....	45
7	TCP FOR THE HP3000.....	51
8	TCP FOR VAX-UNIX.....	52
8.1	Software Distribution and Update.....	52
8.2	Kernel Software Development.....	53
8.3	User Software Development.....	55
8.4	Meetings, Presentations, and Papers.....	56

1 INTRODUCTION

This Quarterly Technical Report is the current edition in a series of reports which describe the work being performed at BBN in fulfillment of several ARPA work statements. This QTR covers work on several ARPA-sponsored projects including (1) development and operation of the SATNET satellite network; (2) development of the Pluribus Satellite IMP; (3) Remote Site Maintenance activities; (4) Internet Operations, Maintenance, and Development; (5) development of the Mobile Access Terminal Network; (6) TCP for the HP3000; and (7) TCP for the VAX-UNIX. This work is described in this single Quarterly Technical Report with the permission of the Defense Advanced Research Projects Agency. Some of this work is a continuation of efforts previously reported on under contracts DAHC15-69-C-0179, F08606-73-C-0027, F08606-75-C-0032, MDA903-76-C-0214, MDA903-76-C-0252, N00039-79-C-0386, and N00039-78-C-0405.

2 SATNET DEVELOPMENT AND OPERATION

Tasks we worked on as part of our participation in the Atlantic Packet Satellite Experiment (SATNET) during the last quarter include the installation of the Raisting, West Germany, Satellite IMP; the elimination of ARPANET Line #77; the Satellite IMP software maintenance operations; and the overall SATNET hardware maintenance operations. These tasks are described in the following sections. Some of our other activities are described below.

A new gateway between SATNET and ARPANET was created at DCEC, and the circuit between the DCEC gateway and the Etam Satellite IMP was made operational. As part of the upgrading of all Packet Satellite Project (PSP) terminals in SATNET, we supported COMSAT in replacing the PSP terminals at Etam, Goonhilly, and Tanum with 2nd generation units; unfortunately the upgrade was accompanied by serious problems (see Section 2.4). Tape cassettes were written for loading new Satellite IMP microcode and macrocode into the Raisting and the Clarksburg C/30 Satellite IMPs. The Satellite IMP Loader/Dumper software was revised to accommodate changes made in the gateway access protocol and to direct messages to the ARPANET host INOC instead of the ARPANET host USC-ISID. New paper tapes of the Loader/Dumper software were generated and sent to all Honeywell

316 sites, where the reloading mechanism is via a Teletype ASR-33 terminal.

During the quarter we also prepared for a demonstration of SATNET monitoring at a symposium held in the Supreme Headquarters Allied Powers Europe (SHAPE) Technical Centre in The Hague, Netherlands, and participated in discussions with ARPA and COMSAT on the reliability of the SPADE satellite channel allocated to SATNET.

2.1 Raisting Satellite IMP Installation

A major milestone for SATNET was reached with the installation and successful operation of a C/30 Satellite IMP located at the INTELSAT earth station in Raisting. This installation represents the first new node added to SATNET since 1977, the first Satellite IMP to participate on the channel using a C/30 packet switch processor, and the first time SATNET has operated with heterogeneous packet switch hardware.

Compared with the difficulties encountered with our original installations of Honeywell 316s, this installation proceeded quite smoothly. When we arrived at the site, we found the crate containing the C/30 had arrived undamaged, having survived commercial truck transport in the U.S. and in West Germany and

Military Air Transport between the two countries. The only casualty was a plexiglass baffle installed at the top of the power supply drawer; this baffle was broken when the C/30 was crated in the U.S. Because no suitable primary power outlet was available (the only nearby outlet was on the same circuit powering the PSP terminal), we had Raisting earth station field site engineers run a 220-volt power line from a nearby fusebox to the C/30.

The BBNCC field site installation tests ran into two problems initially. First, some of the C/30 standard test programs failed to run. Eventually the cause was traced to the presence of a blank configuration ROM in the machine (Satellite IMP operation does not require this ROM). Afterwards, a more serious problem was indicated when the memory test program reported numerous memory parity errors. Not until late the first day was the cause identified as an overheating problem due to running the machine with the front panel off and the separator trays out. When the C/30 was properly dressed, the C/30 test programs ran normally.

Because SATNET operation requires accurate time synchronization among all sites, we replaced the 16 MHz master crystal in the C/30 with a more accurate crystal. Measured accuracies of the original and replacement crystals are 25×10^{-5}

and 2.5 10**-5, respectively.

When the Satellite IMP software was loaded into the C/30, some initial problems were traced to a bug in the macrocode and a bug in the microcode. The former caused monitor reports to be generated and sent to the RECORDER/MONITOR program every Hello frame rather than every 64 Hello frames; a patch for this bug was easily found and implemented. The latter prevented correct interpretation of the Testing and Monitoring (T&M) data appended to received packets by the PSP terminal; consequently, the received packets appeared to have checksum errors. As a temporary solution, the generation of T&M data was disabled, pending a new microcode release. Subsequently, Raisting became a participating member of SATNET with data flowing smoothly and network monitoring and control functions operational. Currently, no hosts are connected to the Satellite IMP; the German Aerospace Research Establishment (DFVLR) is waiting for delivery of a Digital Equipment Corporation (DEC) LSI-11/23 system to serve as a gateway into other LSI-11/23 host systems at DFVLR.

2.2 ARPANET Line #77

For the last three years, part of SATNET resources were used to emulate a memoryless, long-delay, point-to-point circuit in order to provide London IMP connectivity to CONUS ARPANET in the

support of NCP traffic from Europe. This service, designated as ARPANET Line #77, was implemented in SATNET using two streams, one for each direction of traffic (datagram service with two-hop delays was unable to provide a reliable IMP-to-IMP connection, while stream service with one-hop delays was satisfactory). In order to ensure a 9.6 Kb/s full-duplex service to European users, almost 50% of the entire satellite channel bandwidth whether used or not was allocated to Line #77.

The constituent elements of Line #77 include a 50 Kb/s terrestrial circuit between the Etam Satellite IMP and the SDAC ARPANET IMP, a 9.6 Kb/s terrestrial circuit between the Goonhilly Satellite IMP and the London ARPANET IMP, and a virtual-circuit equivalent over the satellite channel between the Etam Satellite IMP and the Goonhilly Satellite IMP. Because of mismatched terrestrial circuit components at both ends, difficulties arose in the operation of this circuit. During heavy traffic loads, the circuit became less reliable, as packets arriving from SDAC would have to be discarded at Etam, including packets that are exchanged by ARPANET neighbor IMPs to determine whether the joining circuit is operational.

During this quarter, Line #77 was permanently removed from service, another step in the ARPANET conversion to TCP. The satellite channel bandwidth previously allocated to Line #77 is

now being allocated to internet traffic routed through the gateways. The Satellite IMP memory previously assigned to code space for supporting Line #77 is now allocated to packet buffering. The Etam Honeywell 316 modem interface previously allocated to Line #77 is now assigned to the circuit between the DCEC gateway and Etam.

2.3 Software Maintenance Operations

During this quarter, Satellite IMP versions 3.5:2 for Honeywell 316s and 4.5:2 for BBN C/30s were released. In addition to bug fixes and incorporation of patches developed since the last software release, the following summarizes the changes made in these versions.

- o Support for ARPANET standard HDH protocol was added to C/30 Satellite IMPs.
- o New host configuration tables to include HDH and 1822 options were implemented.
- o Code that unsuccessfully tried to limit allowable priority and delay class of incoming traffic was removed.
- o Usage of priority and delay class by hosts was standardized.
- o Parameters defining the allowable delivery delays and maximum holding time were changed.
- o The number of SATNET nodes supported was increased from four to five.
- o The number of external hosts supported at each node was increased from two to four.

A detailed discussion of the items identified here are presented in the previous Quarterly Technical Report, No. 26.

We began work on Satellite IMP versions 3.5:3 for Honeywell 316s and 4.5:3 for BBN C/30s to process T&M data generated by the 2nd generation PSP terminal. Initially data will be averaged before being displayed on monitoring terminals.

Work on incorporating the Native Mode Firmware System (NMFS) in Satellite IMPs continued (in NMFS, the emulated machine is altered to create one more suited to communications applications, having greater throughput capacity and reduced latency). The SATNET NMFS microcode was designed, written, and successfully assembled. Using a special test program for the NMFS microcode, packets were looped through a C/30 satellite channel interface when internally crosspatched. The Satellite IMP Loader/Dumper macrocode was converted to NMFS and was successfully assembled.

2.4 Hardware Maintenance Operations

During the last quarter, several hardware problems appeared which we helped diagnose and, when they were related to the Satellite IMPs, fixed. In October, accompanying the upgrade of the PSP terminals at Goonhilly and at Tanum was major channel degradation, such that SATNET gave poor service for a week and

was unusable for several days. Symptoms of the problem were that although short packets appeared to get through the network satisfactorily, large numbers of long packets failed. When substantial numbers of packets are lost, the Satellite IMPs declare themselves out of reservation synchronization frequently, causing communications outages a disproportionate amount of time. To eliminate software as the culprit, we loaded into all sites Satellite IMP Version 3.5:1, which was released last April and was operating in the network all summer long. Channel degradation still persisted with the old software, leading us to believe the problems originated in the channel equipment or the channel itself.

In trying to diagnose the problem, COMSAT determined that interference was present on the SPADE pilot signal. The interference was eventually traced to the Greece INTELSAT earth station, which was one of the earth stations causing problems with SATNET last fall. When the interference was removed, performance on SATNET improved somewhat, to the extent that we were finally able to reload the UCL TAC over the network. Nevertheless, significant problems in transmitting large packets over the network were still present.

When COMSAT switched to backup channels in the PSP terminals at Tanum and at Etam (i.e., 1st generation PSP terminal units)

SATNET stabilized at the position where the overall performance vastly improved. The network was satisfactorily usable, although the overall packet lossage was a little higher than normal and the Hello packet reception was a little worse than normal. Because no 1st generation units existed at Raisting and because the 1st generation unit at Goonhilly was completely detached, those two sites continued to run 2nd generation units.

Later, users reported poor service over SATNET, which we were able to substantiate as the MONITOR program indicated large packet lossage on the channel. The problem increased in severity until finally Etam lost all access to the channel. Since the Satellite IMP worked normally in internal crosspatch but failed to hear packets in external crosspatch (immediately into and out of the PSP terminal), we tried our normal PSP terminal corrective procedures, such as checking power voltages, pressing resets, and cycling power. When all these procedures failed to clear the problem, we called COMSAT for assistance. To restore service, COMSAT switched the Etam PSP terminal back to the 2nd generation unit. Only by having the Etam earth station personnel reseal the Linkabit interfaces numerous times was COMSAT able to make the 1st generation unit work satisfactorily.

At another time there was interference for a period of several days on the SPADE satellite channel #619 assigned to

SATNET. When the interference increased to the extent that all SATNET operations were disrupted, COMSAT operations control was informed, but no action was taken until the Etam daytime shift reported to work. Unable to isolate the interference, Etam earth station personnel moved the SATNET frequencies to SPADE channel #784 to restore service. Afterwards, the interference was traced to the Gondule, Senegal, INTELSAT earth station. Subsequently the SATNET frequencies were returned to SPADE channel #619, accompanied by a brief outage in SATNET service.

When a lengthy outage in the Goonhilly site occurred, several things had gone wrong at once, complicating problem diagnosis. Attempts to restart the Satellite IMP or the loader failed, implying the contents of computer memory had been destroyed. Attempts to reload the loader from paper tape failed, because memory location 1 in the the core bootstrap had been changed, presumably by a primary power interruption (locations 1 through 17 cannot be overwritten by software). When memory location 1 was corrected, we were able to load the loader tape; the loader program, however, was unable to access the channel.

After it was verified that the Honeywell 316 and the PSP terminal voltages were within specifications, COMSAT was called for assistance. While COMSAT was on the phone probing the PSP terminal, the blockage cleared, allowing Goonhilly to access the

channel. After a long while and many attempts to dump and load the Honeywell 316, we were eventually able to bring the Satellite IMP up, but packet lossage on the channel was quite high. Following normal channel recovery procedures, we called Etam site personnel for assistance. They were able to detect a spurious interfering source on the SPADE channel and thereafter called INTELSAT, who called the Gondule earth station to clear the channel.

Later, when a severe interference problem totally disrupted all operations on the SATNET satellite channel, Etam site personnel were able to correct the problem in short order by calling the Gondule earth station again.

An air-conditioning failure at Goonhilly occurred and may be responsible for the following problems. Contents of the Goonhilly Honeywell 316 memory were destroyed, requiring a reload of the Satellite IMP starting with the paper tape loader. The primary power supply in the Goonhilly PSP terminal failed (two of the voltages were below specification); under COMSAT's supervision, the backup unit was brought online. A frequency synthesizer failed and the unit was replaced.

Several disruptions to SATNET service were traced by Etam site personnel to Goonhilly's transmit power being low by about 4 dB. Terrestrial circuit problems in the Washington, D.C. area

occurred, isolating Etam from the ARPANET. After an isolation of the SDAC and London IMPs, BBNCC field service personnel were dispatched to SDAC, where the problem cleared after the modem interfaces in the IMP were reseated. A primary power outage at BBN isolated the ARPANET from BBN40 and, because the BBN gateway is connected to this IMP, from all European TCP traffic.

3 PLURIBUS SATELLITE IMP DEVELOPMENT

During the quarter, the three major project activities at BBN centered around Wideband Satellite Network operations and support, ESI and channel testing, and BSAT software development. Network monitoring switched from the TENEX-based system on BBNC to the NU system running under the UNIX operating system on BBN-INOC. The Network Operations Center at BBN began to assume increased responsibility for Wideband Network operations in terms of monitoring sites, reloading PSATs, and logging Western Union field service activities. Network performance was generally good during the quarter with a few exceptions. One interruption was due to Western Union requiring that all sites be looped off of the channel during the last week of August, while the newly installed earth station at Rome, NY, was adjusted.

The installation of the DAQ remote site monitoring equipment at all sites was completed by Western Union on August 11-12. On August 27 and 30, BBN tested the ability of the ESI to read the earth station status from the DAQ equipment at ISI and Lincoln Laboratory. The testing revealed several non-functioning status indicators, as well as some site-to-site inconsistencies. BBN initiated inquiries with Western Union to gain a better understanding of the DAQ equipment's functionality and to rectify the existing DAQ problems.

Another problem during August was the fact that SRI's presence on the channel interfered with the other sites. It was determined that the upconverter's carrier frequency was off by 1600 Hz. and this was corrected by Western Union. The downconverter at ISI failed on August 11 and was replaced by the spare. The RF loopback relay in the uplink traveling wave tube amplifier at Lincoln failed and was repaired. The downconverter at DCEC failed on September 28 and was repaired. The 125 watt high power amplifier at DCEC failed on October 27 and was replaced by the 75 watt unit.

On September 30, the ESI at ISI was tested at 1.5 Mb/s, BPSK, rate 1/2 coding and was unable to acquire gross frequency offset (GFO) in IF loopback. It was found to fail also in baseband loopback at this data rate. However, linkabit found that they could get the ESI to acquire GFO with their test panel installed. It was determined that this ESI was taking longer to acquire than the nominal 90 seconds allowed by the PSAT software. The PSAT's timeout interval was lengthened to 120 seconds and it was found that the ESI could acquire GFO at the 1.5 Mb/s data rate.

The SRI PSAT developed hardware problems during August. Several memory boards were replaced in an effort to track down the problem, but the PSAT would only run for a few hours.

Finally, on September 16, the satellite modem interface was replaced and the machine has been functioning properly since then. A bus coupler in the DCEC PSAT failed on August 16 and it was replaced.

A new version of the PSAT software was released on August 2, 1982. The new version included software to handle stream synchronization, datagram fragmentation, and more flexible configuration of host interfaces. As experimenters' usage of these new features increased during the quarter, several software bugs were identified and fixed in the new host interface and stream synchronization code. Using the steam synchronization algorithm as a model, new software was added to the PSAT during September to provide synchronization of group databases between sites. Groups provide the mechanism for the delivery of messages to multiple destinations within the Wideband Network. A new version of the software that included group synchronization was released on October 8, 1982.

During the month of October, considerable effort was expended testing the PSAT, ESI, and the satellite channel performance at different data rates and coding levels. Preliminary results have identified several system level problems that should be studied further. It was found that a significant number (about 20%) of packets are dropped by the ESI when bursts

are sent at multiple coding rates. This problem can be alleviated if the interburst padding is increased from 640 to 2048 channel symbols. At rate 1 coding (no coding), the number of data packets received with one or more bit errors increased significantly when the data rate was increased to 1.5 and 3.0 Mb/s. At Lincoln Laboratory, for example, the packet error rate went from less than 1% at 772 Kb/s to 5% at 1.5 Mb/s and 18% at 3.0 Mb/s. DCEC had a packet error rate slightly higher than Lincoln, but the two west coast sites had considerably higher error rates. SRI had packet error rates of 47% at 1.5 Mb/s and 64% at 3.0 Mb/s. ISI had a packet error rate of 56% at 1.5 Mb/s and has not been tested yet at 3.0 Mb/s. The performance of the two west coast sites at the higher data rates could be due to their ESI or earth stations being out of adjustment or to their position in the satellite footprint. No bit errors were detected when the data was sent at rate 1/2 coding. We have not finished testing 3/4 and 7/8 rate coding. Channel testing will continue in order to collect more data and to try different combinations of sites and traffic loads. BBN plans to work with Linkabit, Western Union and Lincoln Laboratory to try to get a better understanding of the current measurements and to develop a plan to improve the system's performance.

BBN began interactions with site personnel. Linkabit, and Western Union, in planning for the upcoming equipment

installations at RADC, Ft. Monmouth, and Ft. Huachuca. On August 19, 1982, Mike Bereschinsky of Ft. Monmouth visited BBN to learn more about the PSATs and the Wideband Satellite Network, and to discuss issues relating to equipment installation at both Forts Monmouth and Huachuca. A PSAT Site Planning and Installation Guide was written by BBN and distributed to all existing, and the next three planned, Wideband Network sites.

Considerable progress was made on the BSAT software development during the quarter. Code to create the major data structures and to initialize the various processes was created. The initialization code includes determining the appropriate configuration for each site, allocating memory for message buffers and queues, and getting the processes to start in a coordinated fashion. In addition, a command processor has been written for the Butterfly console terminal. This command processor is still being enhanced, but is designed to allow reconfiguration of a site when hardware fails, and to allow some performance testing of the BSAT.

The program has run on a three processor Butterfly using the Chrysalis operating system. It has already begun to test and exercise the Butterfly hardware and Chrysalis operating system more strenuously than the Voice Funnel application. In the process of getting these BSAT processes to run, several bugs were

found in Chrysalis. On October 28, 1982, the program was loaded into the Butterfly and ran for the first time. The Message Generator, Echo Host, Message Sink, Local Delivery, Initialization, and the command processor TopLevel were in the process of being debugged at the end of the quarter.

3.1 Network Monitoring by the NU System

During August, Wideband Network monitoring switched from a PDP-10 TENEX-based system on BBNC to the NU system running on a BBN C/70 computer under the UNIX operating system. The NU system, which provides more functionality, flexibility and expandability over the old TENEX-based system, is currently running on several BBN C/70 and DEC PDP-11/70 machines at BBN and elsewhere. At BBN, NU systems running on BBN-NOC, BBN-INOC and BBN-NOC2 provide monitoring and control for the ARPANET, SATNET, Internet Gateways, Wideband Satellite Network, BBNNET and a commercial bank's Domestic Data Network. Currently, the Wideband Satellite Network is being monitored only by NU. Network control and the loading of PSAT software over the ARPANET are still performed by the program U. running on the BBNC-TENEX system. These two functions will be assumed by the NU system as soon as the software is available. In this QTR, we will describe the NU software used to monitor the Wideband Satellite Network. It

should be mentioned that the same NU programs are used to monitor the Atlantic Packet Satellite Network (SATNET).

The NU system's flexibility and adaptability is a result of its division of functions into distinct process modules which interact co-operatively using interprocess communication. This structure has several advantages over a monolithic approach as used in the TENEX system in that it is easier to modify (especially in real time), it uses the underlying operating system facilities to enforce inter-module compatibility and synchronization, and it can be tailored to suit requirements ranging from a single-net, single-protocol, passive system to a multi-net, multi-protocol, active one. NU processes are categorized as either backbone or application, where the former exist to perform server functions required by the latter and perform these functions in response to requests from the latter in a manner similar to requests for services made to the operating system. Thus, application processes rely on the operational environment provided by a supporting level of backbone processes. In addition to the backbone and application software, there is a large body of software concerned with maintenance of, and access to, the network database.

3.1.1 Backbone Processes

There are two backbone processes, the External Message Handler (EMH) and the Event Dispatcher (ED).

The EMH handles all message traffic between the processes of the NU system and the networks. The EMH transfers messages to and from the network using the operating system's raw message interface. The format and function of these messages between the EMH and the UNIX interface to a network is governed by that network's interface conventions and by the class of reports, commands, and facilities defined within that network for operating its components. Application programs request that the EMH send them copies of received network messages matching specific selection criteria.

Processes interested in the contents of specific periodic and aperiodic reports from network components submit selection templates which match those messages. The EMH transforms these messages by converting the received network header format into a corresponding NU system internal format. Processes which need to transfer the contents of memory locations to and from network components request the EMH to set up a bidirectional memory transfer path to the required component. The present memory transfer implementation employs an ARPANET or Internet standard packet core protocol.

The ED performs the role of a system-wide clearinghouse for distribution of event messages. Whenever any process in the system detects an occurrence which could be of interest to other processes in the system or to the network controllers, it generates a descriptive event message and sends it to the ED. Event messages may be triggered by simple occurrences, such as the receipt of an alarm report from a network component, a detected change in status of a network component, a network controller action to access a network component, or a database update.

The ED compares incoming event messages with a list of requests submitted by other system processes (including display processes) and forwards event message copies to those processes whose submitted selection criteria match the received event. These filtered event messages can also be appended to files. Information filtering is available over a large number of dimensions of content or origin. For example, a hardware maintainer could view only messages from components of a particular type, or those which appear to be hardware-related, or which exceed a particular level of severity. The ED's filtering capabilities are more elaborate than those of the EMH: it is possible to specify a class of events to receive and exceptions within the class.

3.1.2 Application Processes

There are currently three major applications programs for the satellite networks monitoring system, as well as a number of supporting programs and UNIX shellscripts. The major programs are:

PROGRAM	DESCRIPTION
-----	-----
snp	satellite network poller
status	program to give status information for sites in user readable form from a database maintained by snp
ltbox	lightbox display program to give quick, continuous network status on a CRT terminal

The supporting programs and the UNIX shellscripts are:

PROGRAM	DESCRIPTION
notice	program to put a message into the network status file for display on the lightbox and whenever the status information is printed out
nuon	shellscript to start the primary copy of snp
nuoff	shellscript to stop the primary copy of snp
mon	shellscript to start a user copy of snp in the foreground and to print incoming messages from the net
monitor	shellscript to start a user copy of snp in the background and to print incoming messages from the net
killmon	shellscript to stop the snp started by monitor
recentlog	shellscript to print the most recent network events from the logfile maintained by the primary snp
wideband	shellscript to set up the environment for monitoring the Wideband Satellite Network
satnet	shellscript to set up the environment for monitoring SATNET
matnet	shellscript to set up the environment for monitoring MATNET
bkr-satnet	shellscript to set up the environment for monitoring the SATNET testbed system at BBN.

The program snp is the primary way of collecting and printing monitoring data from the satellite networks. It connects to the EMH to get monitoring messages and to send polling messages, and it prints text versions of these messages on its standard output. The commands snp, nuon, nuoff, mon, monitor, and killmon all run or control the Satellite Network Poller (snp).

Normally, snp prints only interesting events, but this can be modified by some of the switches listed below. snp has a

number of switches which control its exact behavior:

`-net` =<net number> is required to tell `snp` which network it is monitoring (since it must currently be started while connected to directory `net10.`, the ARPANET).

`-primary` indicates that this is the primary (as opposed to user) copy of `snp`. This means that this `snp` will write its output to the network logfile rather than the terminal, error output will go to the file `snp.errorlog`, and it will maintain the statusfile used by the `status` command. Also, it will change logfiles at GMT midnight, renaming the current logfile to "`logfile.old`", and it will also run a shellscript in the net directory called "`midnight.sh`". The midnight shellscript will be used to prepare the daily summary reports of the Wideband Satellite Network's operations.

`-channel` (`-c`) causes channel reports to be printed
`-status` (`-s`) causes site status reports to be printed
`-hosts` (`-h`) causes host reports to be printed
`-all` (`-a`) is shorthand for '`-c -s -h`'
`-tandm` (`-t`) causes T&M messages to be printed

The program `nuon` starts a primary `snp` with the appropriate net number and the `-primary` switch. The program `nuoff` turns off the primary `snp`. The program `mon` starts a user copy of `snp` with specified flags in the foreground. Typing the interrupt character (usually DEL or ^C) stops this copy of `snp`. `monitor` starts a user copy of `snp` in the background, i.e. the terminal is returned to the shell while `snp` is running and printing messages. `monitor` uses an environment variable "`$MONFLAGS`" to determine what flags to pass to `snp`. The default state is all flags off. `killmon` will kill an `snp` started by `monitor`.

The `status` command is used to find out detailed status

information about a site. If given no arguments, it will print status for all nodes in the network. Arguments are assumed to be site names or abbreviations (one-character abbreviations which are the same as used by snp in printouts) for which status is wanted. status will print out a header including the current time and date (GMT), the network notice, if any, and possibly a warning message if the network statusfile appears out-of-date (usually due to the primary snp not running).

status has only one switch, -verbose (abbr: -v), which causes a very verbose printout of site status, with one site per page (formfeeds between each site).

ltbox generates a very limited summary display of network status on any reasonable CRT terminal such as a DEC VT100. ltbox lists each site, its up/loader/not-heard status, the current loop state of the satellite channel interface, and the state of each of (currently) two hosts. Highlighting is done on various things which have changed recently. or are of exceptional interest (i.e. site in loader). The top of the display contains the net name, the current time and date (GMT), and the current network notice, if there is one.

ltbox has three switches:

-interval (abbr: -i) =<refresh interval in seconds>, tells ltbox how frequently to update the screen from the network

statusfile (default 20 seconds). Making this number larger make the display less current, while shrinking it will chew up more CPU time.

-persis (abbr: -p) =<persistence time in seconds>, is the time interval for highlighting recent events on the display.

-verbose (abbr: -v) gives a more verbose (and much larger) display of the network state, including protocol, channel rate, and idle processor values for each site. Hosts are displayed in a different area of the screen in verbose mode.

ltbox runs an infinite loop, and thus ties up the terminal. It may be killed cleanly by using the interrupt character (usually DEL or ^C).

ltbox will beep the terminal bell if any site entered the loader state within the persistence time. The bell will beep once every refresh interval during the persistence time.

ltbox is intended for only the most cursory overview of the network. Normally net problems seen on the lightbox are investigated further using status and/or recentlog, and possibly monitor in extreme cases.

notice puts a network notice string into the status file for use by status and ltbox. It requires a single argument, which is the notice string. The notice string may be null (""), in which case the current notice will be deleted. Notices are always highlighted on the lightbox display. They should be used for indicating important, unusual states of the whole network, such

as network demos, major experiments, repair work, or software testing.

`recentlog` prints out the most recent entries from the logfile maintained by the primary copy of `snp`. If given no arguments, it will print the most recent 23 lines of the status file. Arguments are passed to the the UNIX program `tail`, and the output piped through the UNIX program `more`, with the belief that more of the log will be requested in the arguments. For example `recentlog` would pipe the last 100 lines of the logfile through `more` to the terminal. Because of the way `tail` works, arguments larger than about -200 will not work. The argument `-<number>b` should be used to get the last `<number>` blocks of the logfile instead.

4 REMOTE SITE MAINTENANCE

The Remote Site Maintenance (RSM) Project contract for FY82 expired on September 30. Hardware and software maintenance at the RSM sites continued until that time. The main activity in the final months of the project was completion of documentation of various procedures and software subsystems of the RSMs.

Documentation prepared during the quarter includes:

Graphics Support for BBN-UNIX, BBN Technical Memorandum No. 708.

Compiling Programs with MAKE, BBN Technical Memorandum No. 709.

RSM Accounting System Summary, BBN Technical Memorandum No. 710.

How to Install a BBN-UNIX System, BBN Technical Memorandum No. 714.

RSM Procedures, Practices, and Standards, BBN Technical Memorandum No. 717.

The CRT Text Editor Ned: Tutorial and Reference Manual, BBN Report No. 5174.

A final report summarizing activity in the RSM project during the past year was also completed (Remote Site Maintenance 1982 Final Report, BBN Report No. 5209).

5 INTERNET DEVELOPMENT

5.1 Introduction

The major activity during the past quarter was the continued deployment and maintenance of the Macro-11 gateway. Other important work included enhancing NU gateway monitoring, moving the UCL-TAC onto the UCL gateway, continuing the definition of the Exterior Gateway protocol, and designing the VAN gateway.

5.2 Gateway Installations

A new packet radio gateway was installed at SRI (C3PR). This is the first time that a gateway has run on a non-zero port of a port expander.

A VDH port to SATNET was added to the DCEC gateway causing it to become the second ARPANET-to-SATNET gateway. Having two gateways between SATNET and the ARPANET is providing more reliable service to the European Internet users because when one gateway is down there is a path via the other.

The list of operational gateways is shown in the following table.

Gateway -----	Adjoining Networks -----
BBN	ARPANET - SATNET
BBN-NET	ARPANET - BBN-NET
BBN-PR	BBN-NET - BBN-PR
BRAGG	ARPANET - BRAGG-PR
C3PR	ARPANET - C3-PR
DCEC	ARPANET - SATNET - EDN
NTARE	SATNET - NTARE-TIU - NTARE-RING
SRI-C3P0	ARPANET - SF-PR-2
SRI-R2D2	ARPANET - SF-PR-1
UCL	SATNET - UCLNET - RSRE/NULL - UCL-TACNET

5.3 Gateway Throughput Collection

In the past quarter we have started collecting Gateway throughput data on a 24-hour, seven-day/week basis. We now have the capability to produce daily, weekly, and monthly summaries of gateway throughput. A example of throughput for the week of November 22, 1982 follows:

TOTAL THROUGHPUT

GWY NO.	GWY NAME	RCVD DGRAMS	RCVD BYTES	IP ERRORS	% IP ERRORS	DEST UNRCH	% DST UNRCH
1	RCC	3,812,822	171,635.122	1,280	0.03%	7,201	0.19%
2	BBN	1,216,792	40,289.180	19	0.00%	621	0.05%
3	UCL	1,068,190	47,951,713	0	0.00%	1,160	0.11%
4	NTARE	484.836	16,354.418	1	0.00%	150	0.03%
7	BRAGG	1,027,097	30,427,292	1,014	0.10%	2,562	0.25%
13	R2D2	852,817	22,947,072	1,026	0.12%	196	0.02%
14	C3PO	988,636	26,590,234	996	0.10%	766	0.08%
15	DCEC	1,551,366	74,220.560	1,793	0.12%	3,295	0.21%
23	CRONUS	272,659	8,075,578	0	0.00%	0	0.00%
31	C3PR	1,161,559	84.560,848	1	0.00%	127	0.01%
32	OLDBBN	1,687,546	46,015,680	9,601	0.57%	701	0.04%
TOTALS		14,124.320	569,067,697	15,731	0.11%	16,779	0.12%

GWY NO.	GWY NAME	SENT DGRAMS	SENT BYTES	DROPPED DGRAMS	% DROPPED DGRAMS
1	RCC	3,858.731	165,551,220	84.218	2.14%
2	BBN	1,304,731	41,647,846	449	0.03%
3	UCL	1,154.049	50,230,126	21	0.00%
4	NTARE	560.200	20,337,270	3,435	0.61%
7	BRAGG	1,131.952	30,719,397	20	0.00%
13	R2D2	871,488	22,000,870	5	0.00%
14	C3PO	1,078,763	27,328,373	492	0.05%
15	DCEC	1,565,120	72,651,845	3,916	0.25%
23	CRONUS	310,984	8,270,097	0	0.00%
31	C3PR	1,238,737	83,146,587	513	0.04%
32	OLDBBN	2,162,271	60,748.200	474	0.02%
TOTALS		15,237,026	582,631,831	93,543	0.61%

RATE (per second) and SIZE (bytes per datagram) TABLES

GWY NO.	GWY NAME	RCVD DGRAMS	RCVD BYTES	IP ERRORS	AVG BYTES PER DGRAM
1	RCC	6.38	287.21	0.00	45.02
2	BBN	2.03	67.22	0.00	33.11
3	UCL	1.78	80.00	0.00	44.89
4	NTARE	0.81	27.28	0.00	33.73
7	BRAGG	1.71	50.76	0.00	29.62
13	R2D2	1.59	42.71	0.00	26.91
14	C3P0	1.65	44.43	0.00	26.90
15	DCEC	2.59	124.01	0.00	47.84
23	CRONUS	1.38	40.97	0.00	29.62
31	C3PR	2.17	157.65	0.00	72.80
32	OLDBBN	2.82	76.89	0.02	27.27

GWY NO.	GWY NAME	SENT DGRAMS	SENT BYTES	DROPPED DGRAMS	AVG BYTES PER DGRAM
1	RCC	0.46	277.03	0.14	42.90
2	BBN	2.18	69.48	0.00	31.92
3	UCL	1.93	83.80	0.00	43.53
4	NTARE	0.93	33.93	0.00	36.30
7	BRAGG	1.89	51.25	0.00	27.14
13	R2D2	1.62	40.95	0.00	25.25
14	C3P0	1.80	45.66	0.00	25.33
15	DCEC	2.62	121.39	0.00	46.42
23	CRONUS	1.58	41.96	0.00	26.59
31	C3PR	2.31	155.01	0.00	67.12
32	OLDBBN	3.61	101.50	0.00	28.09

PEAK PERIODS

(Throughput is sum of dgrams/sec, input + output.
Time is time of data collection)

GWY NO.	GWY NAME	TOTAL T'PUT	TIME OF DAY	DROP RATE	TIME OF DAY
1	RCC	52.60	11/22 16:35	20.86%	11/22 03:09
2	BBN	13.04	11/22 08:09	5.63%	11/24 07:50
3	UCL	22.55	11/22 11:24	0.30%	11/24 06:28
4	NTARE	7.33	11/22 12:24	11.07%	11/28 12:25
7	BRAGG	9.87	11/26 09:46	1.16%	11/28 12:55
13	R2D2	5.31	11/23 16:41	0.27%	11/24 07:51
14	C3P0	7.49	11/23 16:41	7.77%	11/22 09:39
15	DCEC	15.60	11/22 11:25	2.87%	11/25 18:46
23	CRONUS	6.26	11/26 18:55	0.00%	11/28 23:55
31	C3PR	22.24	11/23 20:26	5.25%	11/23 15:54
32	OLDBBN	11.78	11/23 11:24	4.45%	11/28 03:10

We are now working on the capability to send daily, weekly, and monthly throughput summaries out, via Electronic mail, to a list of people. When this is in place, the CMCC gateway monitoring program on ISID can be turned off.


5.4 Gateway Status Processor

We have been working on building a gateway status processor in the NU Network Monitoring system, to keep track of which gateways are up, down, or isolated, which gateway's interfaces are up or down, and whether a gateway's configuration changes (e.g., a new version of software is loaded). When the status processor is completed, we will install it in the Network Operations Center (NOC). This will provide the NOC operators

with an important tool to help them start operating the gateways.

5.5 UCL Gateway

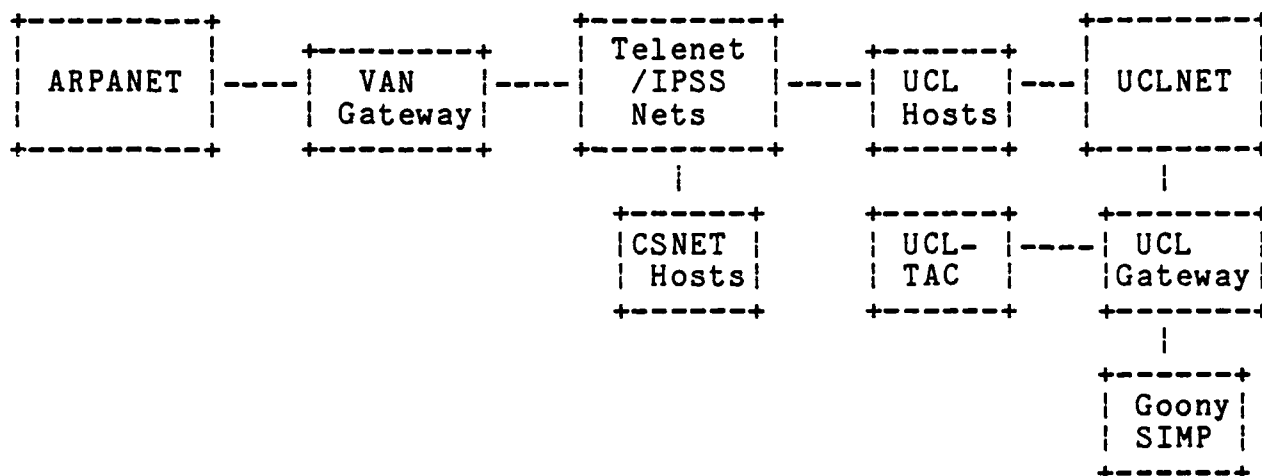
On 1 October, ARPANET Line 77 was permanently removed from service, thereby isolating the London IMP. Prior to the decommissioning of Line 77, the UCL-TAC was installed directly on the UCL gateway, as the final step in the conversion of UCL to TCP-only operation. The current configuration at UCL is as follows:



5.6 VAN Gateway

The VAN gateway, which interconnects the ARPANET and the X.25 Telenet/IPSS networks, will be initially used primarily to provide access to UCL and CSNET hosts. The VAN gateway will also be used as an alternate path for diagnostic purposes to the European SIMP, the UCL Gateway, and UCL-TAC whenever connectivity

over SATNET is absent. Note that the VAN gateway is designed to not have any neighbor gateways on the Telenet/IPSS networks. The overall configuration is as follows:



5.6.1 Access Control Table

Because the VAN gateway is required to restrict the access between the ARPANET and Telenet/IPSS, it must perform access control on traffic originating from both the ARPANET and the Telenet/IPSS. The gateway will implement access control by keeping a table containing information on which group of hosts is allowed to send datagrams to which hosts on or via Telenet/IPSS. Each entry in the table will contain the following items:

- 1) The IP addresses of hosts permitted to exchange datagrams over Telenet/IPSS. Only datagrams containing these IP addresses will be allowed to transit through the gateway. Datagrams with IP source routing options will be scanned by the access control procedure to check that all addresses are

in the table.

- 2) Flags to control under what circumstances Virtual circuits are opened by the VAN gateway. These will be used to control which hosts the VAN gateway will accumulate Telenet/IPSS packet charges.

Flags controlling when circuits are opened have three settings, as follows. 1) Open circuit when none exists. 2) Open Reverse charge circuit when no open circuit exists. This is only valid for hosts in the GTE Telenet network. Reverse charge circuits are not supported for international circuits. 3) Don't open circuit. Consequently datagrams will be sent only when the Telenet/IPSS side first opens the circuit.

5.6.2 X.25 Address Table

The VAN gateway will maintain a table to map from Telenet/IPSS IP addresses and Telenet/IPSS X.25 addresses. The table will contain entries which include the IP address and the X.25 address that corresponds with the IP address.

5.6.3 Datagram Processing

Datagrams sent to or received from the Telenet/IPSS network will first be checked by the VAN gateway to see if their source/destination IP addresses (and Source Route IP addresses if applicable) are in the access table. If they are not, the datagram will be discarded and an ICMP Destination Unreachable message will be sent in response to the datagram. Currently, the

ICMP Destination Unreachable message does not have a code defined that describes datagrams dropped for failing access control restrictions. A code, such as Administratively Prohibited, should be added for this purpose.

If the datagram is to be routed to the Telenet/IPSS network, the VAN gateway looks up the destination X.25 address in the X.25 Address table. It then checks to see if there is an open virtual circuit to the Telenet/IPSS destination address that it found in the X.25 address table. If a circuit is open, the datagram will be sent over it. If there is no open circuit, the VAN gateway opens a virtual circuit based on the control field. If the VAN gateway is permitted to open a circuit (or reverse charge circuit), it will do so and the datagram will be sent over it. If it is not, then the datagram will be dropped, and an ICMP Destination Unreachable, code Host unreachable will be returned.

5.6.4 Virtual Circuit Management

The VAN gateway will always send a datagram on the virtual circuit that was most recently opened to it from a Telenet/IPSS host. It will close any duplicate circuits to the same host. This will reduce the packet charges for the VAN gateway because it favors circuits opened by the Telenet/IPSS host. The VAN gateway will not accept any reverse charge circuits.

The VAN gateway will close virtual circuits that it has opened that idle for more than a fixed length of time. The time limits will be parameters. The initial values of these parameters will be 1 minute and 3 minutes, depending on whether the circuit was a regular or a reverse charge circuit.

5.6.5 X.25 Usage

The VAN gateway will use the IP protocol directly over X.25; there will be no intermediate protocols involved. The IP datagrams will be treated as X.25 data.

The gateway will assume the Telenet/IPSS networks as having a maximum message size of 512 bytes. Datagrams larger than that will be sent as IP fragments. Datagrams larger than the X.25 maximum packet size of 128 bytes will be broken into X.25 packets with the "more data bit" set to indicate the datagram is being sent in pieces.

There will be no data sent in the X.25 packets that are used to establish virtual circuits.

5.7 Exterior Gateway Protocol

The final version of "RFC 827, Exterior Gateway Protocol (EGP)" was submitted and has been relased as a draft standard. We are starting to work on the design of the gateway implementation of the EGP.

5.8 Gateway RFC

The final version of "RFC 823, The DARPA Internet Gateway" was submitted and has been released. This document describes the current Internet gateway and includes detailed descriptions of message formats and gateway procedures.

6 MOBILE ACCESS TERMINAL NETWORK

As part of our participation in the development of the Mobile Access Terminal (MAT) and the MAT Satellite Network (MATNET) during the last quarter, we implemented error protection on the packet length parameter in the packet headers and we designed and implemented a set of ruggedization procedures for BBN C/30 commercial packet switch processors. These two items are described in the following sections. Some of our other activities are described below.

BBNCC delivered to us two C/30 standard packet switch processors, which we installed into the BBN MATNET test rack for an extensive burn-in period prior to any modifications necessary for ruggedization of the units. Because MATNET operation requires accurate time synchronization among all sites, we replaced the 16 MHz master crystals in these C/30s with more accurate crystals. After the ruggedization was completed, both C/30s were submitted to a second burn-in period and were subsequently shipped to E-Systems, ECI Division, in St. Petersburg, Florida, for system checkout.

Satellite IMP version 6.2:2 software was assembled and written on tape cassettes for checkout. As part of software checkout, we created a two-node network in the BBN testbed facility, using the two MATNET C/30s which were being ruggedized

at BBN. Among the changes in this version are new microcode implementing error protection on the length parameter in the MATNET packet header, corrections for bugs found in the last field tests, and macrocode updates from SATNET. The most visible change, however, is that Satellite IMP version 6.2:2 has four-node network capability, where previous releases were limited to a maximum of three nodes.

We visited the carrier USS Carl Vinson in drydock at Norfolk, Virginia, to discuss shipboard tests with MATNET. Despite the size of the ship, room for the MAT equipment is scarce. We processed a request from NAVELEX for a quote on a sixth MAT Satellite IMP.

6.1 Error Detection on Packet Length

One of the conclusions of the preliminary Phase 2 tests in MATNET is that many contention packets would be received with errors in the packet length parameter of the packet header when CPODA (Contention Priority-Oriented Demand-Assigned) channel protocol was used for satellite channel scheduling. Without error protection on the length field, the Satellite IMP microcode would block reception of subsequent packets until enough bits had arrived corresponding to the corrupted length parameter in the original packet. Thus, many packets in sequence could be lost,

forcing the Satellite IMP to go out of reservation synchronization and causing a communications outage a disproportionate amount of time. To circumvent this problem, we implemented coding for error protection on the packet length parameter inserted into the header of every packet in MATNET.

In MATNET, packets can range up to a maximum of 128 16-bit words of data plus 14 words of header; hence, although the packet length field is 12 bits long, the packet length can be represented as an 8-bit number. To incorporate error protection on the packet length, we have converted the remaining 4 bits in the field to parity bits on the length parameter. The error-protection code chosen is a binary linear systematic code having minimum weight 3 (equivalent to a minimum Hamming distance of 3, allowing detection of all 2-bit errors), represented as a (12,8,3) code.

Shown below is the binary generator matrix of the code, in which prefixed to each 8-bit length vector $L=(L_1, L_2, \dots, L_8)$ is a 4-bit parity vector $P=(P_1, P_2, P_3, P_4)$ to create a 12-bit code vector $C=(C_1, C_2, \dots, C_{12})$

PARITY				LENGTH							
P	P	P	P	L	L	L	L	L	L	L	L
1	2	3	4	1	2	3	4	5	6	7	8
0	0	1	1	0	0	0	0	0	0	0	1
0	1	0	1	0	0	0	0	0	0	1	0
0	1	1	1	0	0	0	0	0	1	0	0
1	0	1	1	0	0	0	0	1	0	0	0
1	1	0	1	0	0	0	1	0	0	0	0
1	1	1	0	0	0	1	0	0	0	0	0
1	0	1	0	0	1	0	0	0	0	0	0
1	1	0	0	1	0	0	0	0	0	0	0

The Parity can be represented as

$$\begin{aligned}
 P1 &= L1 \text{ EXOR } L2 \text{ EXOR } L3 \text{ EXOR } L4 \text{ EXOR } L5 \\
 P2 &= L1 \text{ EXOR } L3 \text{ EXOR } L4 \text{ EXOR } L6 \text{ EXOR } L7 \\
 P3 &= L2 \text{ EXOR } L3 \text{ EXOR } L5 \text{ EXOR } L6 \text{ EXOR } L8 \\
 P4 &= L4 \text{ EXOR } L5 \text{ EXOR } L6 \text{ EXOR } L7 \text{ EXOR } L8
 \end{aligned}$$

where EXOR is the Boolean logical exclusive OR operation. The parity vectors associated with the length vectors in the generator matrix have the characteristics that the parity vectors are unique and the weight (total number of 1's) in each 12-bit code vector is 3 or greater. In addition, for greater symmetry, we chose the parity vectors such that when two length vectors are the transpose of each other, their parity vectors are also the transpose of each other. The generation of any code vector requires a linear combination of length vectors in the generator matrix to construct the appropriate length value and the EXOR of their associated parity vectors to construct the accompanying parity vector.

In implementation, the microcode prefixes the parity vector to the packet length parameter via a table look-up mechanism before transfer of the packet to the Black Processor. Upon receipt of a packet from the Black processor, the microcode checks the received parity vector via the same table look-up mechanism. If an error is detected (i.e., the received parity vector disagrees with the value in the table), the microcode aborts processing of the incoming packet and begins searching for the next packet after incrementing an error counter in a microcode register. If no error is detected, the microcode sets the parity vector to all zeroes.

6.2 C/30 Ruggedization

Commercially manufactured C/30 packet switch processors are designed for installation into benign areas suitable for most digital computing equipment; i.e., into air-conditioned rooms free from excessive dust and moisture. However, in Phase 3 of the MATNET program, MATs are to be installed on Naval ships, where significant vibration and mechanical shock are present. Equipment survivability of the commercial design in such a shipboard environment was seriously questioned, inasmuch as some chips, including all memory chips, are not soldered but inserted into sockets and the power supply is inadequately mounted to

resist abuse. Furthermore, the large printed circuit boards in the commercial design are supported only on their sides by the guides in the chassis; vibration of the chassis could dislodge these boards, causing an accidental short circuit through inadvertent contact between boards. Therefore, we have established a set of minimal modifications to the standard C/30s to enhance survivability for the approaching sea-going MATNET demonstration tests. These modifications include the following items:

- increased spacing between printed circuit boards;
- printed circuit board stiffening;
- securing chips to the printed circuit boards;
- securing daughter boards to the mother boards;
- securing printed circuit boards to the chassis;
- securing the power supply to the chassis;
- securing the front panel to the chassis;
- modifying the air intake;
- installing a high temperature 110-volt power cutout;
- isolating primary power with an RFI filter;
- securing cables to their chassis connectors.

The C/30 ruggedization design philosophy assumes that shock mounts will be installed at the base of the entire C/30 rack. These mounts in conjunction with the large total mass of the rack dampen high-frequency vibration and mechanical shock. Modifications to the C/30 are directed toward raising the resonant frequencies of the printed circuit boards away from the low-frequencies not damped by the shock mounts. A design principle to which we firmly adhered is that it is absolutely

imperative that the introduction of these items must not adversely affect airflow and cooling and must not modify any of the electrical equipment (printed circuit boards and power supply) to the extent that the manufacturing facility would be unable or unwilling to provide commercial hardware maintenance service. Specific details on these modifications are presented below.

Because available space in the small (3-slot standard) C/30 chassis is severely limited and thereby restricts implementation of the necessary modifications, the large (7-slot standard) C/30 chassis was procured for the MATNET Satellite IMPs. Unused slots of the chassis are filled in the manufacturing process with empty trays for ducting cooling air over those slots filled with printed circuit boards. However, we rearranged the standard machine by placing the printed circuit boards in every even numbered slot {2, 4, 6} and the empty trays in every odd numbered slot {1, 3, 5, 7}. The new arrangement is processor board top, I/O board middle, and memory board bottom, so that the register display lights can show through the window in the front panel.

By bolting each printed circuit board to the standard tray beneath it (using specially fabricated spacers), we were able to increase board rigidity without any modification to the board itself. No new holes were drilled into the boards; bolts were

inserted into existing holes formed during the manufacturing process. The only printed circuit board modifications included soldering corner pins of the plug-in chips to the sockets attached to the board and moving the board stiffener from the top of the board to the bottom of the board. In the latter position, the board stiffener seats against a brace we installed on the tray for increased board rigidity. To direct cooling air over the bottom of the printed circuit boards, we milled cutouts in the sides of the trays.

To provide positive support, we added braces to the front of the chassis and registration support pins on the printed circuit board/tray assembly. During insertion of the board/tray assembly into the chassis guide channels, the registration support pins seat into holes drilled in the front braces.

All I/O daughter boards are secured in place by bolting their connectors to the mating connectors on the mother boards. To provide additional support, insulated standoffs are bolted between the mother board and the I/O daughter boards on the edge opposite where the connector of the daughter board is located. Bolt down caps are placed over the 1822 Local Host interface cable connectors to secure them in place on their mating connectors on the mother boards.

To provide a positive screw-down attachment of the front panel to the chassis, holes were drilled in the sides of the front panel and the sides of the chassis, and screws were inserted. Brackets were added to the rear of the chassis to allow the chassis to be supported both front and back in the rack.

In order to secure the power supply, we raised the 12-volt power supply module so that the tops of the 12-volt and 5-volt power supply modules are the same height and placed a brace over both modules in such a fashion as to secure the 12-volt module to the 5-volt module. The entire assembly is then attached to the air ducting side wall in the drawer. To provide positive support, we added registration support pins on the front of the power supply drawer for seating into holes drilled in a front brace.

Because the MAT shipboard installation is expected to create hot-spots at the back of the rack, the air intake at the rear of the power supply drawer was blocked to prevent the cooling fans from drawing heated air into the chassis. A new cooling air intake was created by milling slots in the front panel. Weather stripping was placed on the front panel to eliminate air leaks around the sides. To improve air cooling around the power supply, we installed an additional baffle in the power supply

drawer. We also added a high-temperature cutout thermostat to disconnect the 110-volt electrical power to the power supply while leaving the cooling fans running during overheated conditions.

To minimize RFI interference, we installed an RFI filter. To secure cables, we placed positive screw-down attachments over the plug of the 110-volt primary power cable and over the console terminal/cassette loader connector. We believe all other cables are satisfactorily secured to the C/30 during normal equipment manufacturing.

All the above items are specifically directed for equipment survival in an environment where a significant amount of vibration and mechanical shock are present. If salt spray and fungus growth are environmental problems sufficiently severe to require protection, then application of a protective coating may be added later in the equipment testing stages. However, it is preferable to avoid coatings, if at all possible, to facilitate hardware maintenance on the printed circuit boards.

7 TCP FOR THE HP3000

During the last quarter of the HP3000 Internet project no effort was expended in maintaining or improving the HP3000 Internet software. The only effort on the project occurred late in September when the LSI-11 Front End experienced hardware problems. The problems were serious enough to require BBN to send a person to DARPA in Washington to diagnose them. Diagnosis revealed a broken LSI-11 memory board, which was subsequently replaced by Digital Equipment Corporation field service.

No software delivery was made because DARPA does not have the requisite hardware needed to connect an HP3000 directly to the ARPANET.

8 TCP FOR VAX-UNIX

The VAX TCP project continued with activity in the areas of kernel and user software development, distribution and bug fixing of the 4.1BSD version of TCP, and attendance at the DARPA Internet and Berkeley Steering Committee meetings.

8.1 Software Distribution and Update

The BBN VAX TCP/IP Software Distribution has now been sent to 30 sites. New orders for the software have slowed, owing to the coming availability of the Berkeley 4.1cBSD and 4.2BSD systems. We expect an increase in orders, however, as the January 1, 1983 ARPANET TCP transition date nears.

We have had very useful feedback from a number of sites, especially USC Information Sciences Institute, MIT Laboratory for Computer Science, and Purdue University, on problems with the kernel and user software. Their help was valuable in tracking down and fixing a number of bugs in the distribution.

Among the major bugs that were found and fixed during the quarter were the following. A reversal of two lines of code in one of the TCP send routines caused data to be held and not transmitted under certain circumstances, until another TCP send was issued by the user. The main effect of this bug was to cause

improper closure of TCP connections that was especially apparent with SMTP mail, and caused multiple copies of messages to be transmitted. Another problem was caused by an off-by-one error in TCP FIN processing when a FIN was present in a segment containing data. This bug caused FTP connections to not close properly under certain conditions. A bug that caused the network buffer allocation limits to slowly rise over time until the maximum number of network buffers was allocated was traced to improper lowering of the allocation for network control connections. Finally, a number of bugs were found in the FTP user and server programs, which affected operation with TOPS-20 hosts. These bugs included improper handling of IACs on the FTP control connection, use of the wrong line terminator in ASCII transfer mode, improper handling of the argument to the TYPE LOCAL command, bugs in the handling of arguments to the LIST and NLST commands, and a problem in properly handling the invocation of user FTP with no destination host argument.

8.2 Kernel Software Development

There was much new software development for the VAX TCP kernel. Most of the effort centered on making the software interrupt version of the network code more reliable, and developing network interface device drivers. A new internal host

map scheme was developed that uses hashing for better performance in lookups, and that places the host entries in network buffer memory, rather than a fixed size table. The internal host map entry format was modified to hold Ethernet address mappings for use with the new InterLan Ethernet controller driver.

Driver development work continued for the InterLan Ethernet controller and the ACC IF-11/HDH ARPANET interface. The InterLan controller was completed, using an address translation scheme developed at MIT Laboratory for Computer Science. This scheme is similar to the one we proposed (which was reported on in Quarterly Technical Report No. 26, BBN Report No. 5529), but differs by implementing the address resolution packets in the Ethernet local network layer, rather than in the IP layer. It was adopted at the DARPA Internet meeting (see below). Initial throughput tests of the InterLan controller (looping on the Ethernet) indicated a throughput of approx. 300 Kb/s. By increasing the network buffer size from 128 to 256 bytes, this figure was improved to 400 Kb/s with the same test. The major limiting factor here seems to be interrupt load on the processor.

Much time was spent on debugging the driver for the ACC IF-11/HDH interface. Hardware problems and errors in documentation of the IF-11 were the main cause of delay in getting the driver to work. Problems included a bug which prevented the device from

working on the VAX-11/780 UNIBUS, a number of errors in the device documentation which caused much wasted effort in driver debugging, and problems with the ROM software provided by ACC that implement the 1822 HDH protocol. As of this writing, the host interface part of the device has been debugged, and the driver works with the IF-11/HDH in internal loopback mode. However, the device does not yet work when connected to an HDH IMP, a problem that ACC is attempting to fix. Once this problem is corrected, the driver should work correctly with the device.

Other changes to the TCP kernel included a modification to the pseudo-teletype driver to allow the TELNET server to close its TCP connection when a user does a logout, and an optimization to reduce the number of I/O operations needed to transmit small messages (such as for TELNET interactive use).

8.3 User Software Development

Aside from the bug fixes to FTP described above, the main work in the higher level protocol user programs centered on the host address translation library and the name server. A number of fixes were made to the network library software. In addition, documentation on the new network library was completed.

An implementation of the Internet Name Server described in IEN-116 was completed and installed on the BBN VAX. This name server works with the new network library software, and allows name/address/service queries over User Datagram Protocol (UDP) connections. The name server and related user software was tested against the Network Information Center (NIC) name server implementation. Further work in this area requires modifications to the IEN-116 name server specification, which is not complete in several areas. In addition, software to allow caching of frequently requested translations in client hosts is needed to allow operational use of the name server software.

8.4 Meetings, Presentations, and Papers

We participated in the DARPA Internet Meeting held at DFVLR in Oberpfaffenhoffen, Federal Republic of Germany, on September 14-15. At the meeting we gave a presentation on the state of the VAX TCP implementation and some performance measurements of the system, as well as a presentation on the scheme we proposed for Ethernet/IP address mapping.

We also participated in the Berkeley Steering Committee meeting that took place at DARPA in Washington on September 20-21. The main topic of discussion related to the VAX TCP Project was the role of the BBN VAX TCP software in the coming 4.2BSD

system. Various options for reconciling the BBN and Berkeley versions of the software were discussed, but no final decisions were reached.

We wrote a paper in which the VAX TCP implementation was used as a model for discussing issues in host network software. The paper. "Systems Engineering Issues in Network Protocols" by J. Haverty and R. Gurwitz, will appear in a future issue of Data Communications.

DISTRIBUTION

ARPA

Director (3 copies)
Defense Advanced Research Projects Agency
1400 Wilson Blvd.
Arlington, VA 22209
Attn: Program Manager
R. Kahn
R. Ohlander
D. Adams
B. Leiner

DEFENSE DOCUMENTATION CENTER (12 copies)

Cameron Station
Alexandria, VA 22314

DEFENSE COMMUNICATIONS ENGINEERING CENTER

1860 Wiehle Road
Reston, VA 22090
Attn: Lt. Col. F. Zimmerman

DEPARTMENT OF DEFENSE

9800 Savage Road
Ft. Meade, MD 20755
R. McFarland C132 (2 copies)

DEFENSE COMMUNICATIONS AGENCY

8th and South Courthouse Road
Arlington, VA 22204
Attn: Code 252

NAVAL ELECTRONIC SYSTEMS COMMAND

Department of the Navy
Washington, DC 20360
B. Hughes, Code 6111
F. Deckelman, Code 6131

BOLT BERANEK AND NEWMAN INC.

1300 North 17th Street
Arlington, VA 22209
E. Wolf

MIT Laboratory for Computer Science

545 Technology Square
Cambridge, MA 02138
D. Clark

DISTRIBUTION cont'd

BOLT BERANEK AND NEWMAN INC.

10 Moulton Street
Cambridge, MA 02238

S. Blumenthal
M. Brescia
R. Bressler
R. Brooks
P. Carvey
P. Cudhea
W. Edmond
L. Evenchik
G. Falk
J. Goodhue
S. Groff
R. Gurwitz
J. Haverty
F. Heart
J. Herman
R. Hinden
D. Hunt
E. Hunter
S. Kent
W. Mann
A. McKenzie
D. McNeill
W. Milliken
M. Nodine
R. Rettberg
R. Reynolds
J. Robinson
E. Rosen
P. Santos
J. Sax
S. Storch
R. Thomas
B. Woznick
Library

Report No. 5215

Bolt Beranek and Newman Inc.