LEVEL II

(2)

Final Report

COMSAT PARTICIPATION IN ARPA PACKET SATELLITE PROGRAM (PSP)

Sponsored by

Defense Advanced Research Projects Agency
ARPA Order No. 3287

Monitored by SAMSO/YAPC Under
Contract No. F04701-C-0240

October 1979

COMSAT Laboratories
Clarksburg, Maryland 20734

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE (19)	READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER (18) SD TR-80-32 (6)	2. GOVT ACCESSION NO. AD-A089 224
	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Final Report; COMSAT Participation in ARPA Packet Satellite Program (PSP).	5. TYPE OF REPORT & PERIOD COVERED Final, ~~5 October 76 to~~ 1 July 79
	6. PERFORMING ORG. REPORT NUMBER (14) TSLAP-79-2
7. AUTHOR(s) (10) L. C. Palmer, D. Mills, S. Rothschild	8. CONTRACT OR GRANT NUMBER(s) (15) F04701-76-C-0240.
9. PERFORMING ORGANIZATION NAME AND ADDRESS Communications Satellite Corporation 950 L'Enfant Plaza, S.W. Washington, D.C. 20024 (COMSAT Laboratories)	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS ARPA Order No. 3287
11. CONTROLLING OFFICE NAME AND ADDRESS Department of the Air Force Headquarters Space Division (AFSC) Los Angeles Air Force Station, P.O. Box 92960 Worldway Postal Center Los Angeles, California 90009 (12) 287	12. REPORT DATE (11) October 1979
	13. NUMBER OF PAGES 292
14. Monitoring Agency Name & Address Defense Advanced Research Projects Agency (DARPA) 1400 Wilson Blvd. Arlington, Virginia 22209	15. SECURITY CLASS. (of this report) Unclassified
	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE N/A

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited

(9) Final rept, 1 Jun 78 - 1 Sep 79,

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

Same as 16.

18. SUPPLEMENTARY NOTES

None

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

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

This final report summarizes the last year of the Communication Satellite Corporation's (COMSAT's) participation in the Atlantic Packet Satellite Program (SATNET), an experimental program involving packet data transmission via a single 64 Kb/s SCPC channel in the SPADE transponder of INTELSAT IVA (Atlantic Primary) between standard A earth stations at Etam, W.Va., Goonhilly U.K., and Tanum Sweeden. A fourth, smaller (G/T = 28.5 dB/K) earth station at COMSAT Laboratories, Clarksburg, Md. also participated in the experimental network.

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

406867 DM

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

Atch #1

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

The report describes new packet satellite program (PSP) terminals which were fabricated at COMSAT Laboratories for installation at the large INTELSAT earth stations. These terminals contain the necessary interfaces and modems to interconnect the SIMP (Satellite Interface Message Processor) and the earth station intermediate frequency (i.f.) equipment. Also described is the command and monitoring (C&M) module which is a part of the PSP terminal.

Channel monitoring data collected during 1979 are summarized. Preliminary tradeoff result comparing packet satellite transmission to other alternatives are also included in the report.

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

2

COMMUNICATIONS SATELLITE CORPORATION
950 L'Enfant Plaza, S.W.
Washington, DC 20024

TSLAP-79-2

FINAL REPORT

1 June 1978 to 1 September 1979

Issued 1 October 1979

Sponsored by
Defense Advanced Research Projects Agency
ARPA Order No. 3287

Monitored by SAMSO/YAPC under
Contract No. F04701-76-C-0240

Period of Performance

5 October 1976 to 1 July 1979

L. C. Palmer
Program Manager

301-428-4575

ACKNOWLEDGEMENTS

Sections of this final report were prepared by various members of COMSAT Laboratories. Section 4 was prepared by Dr. D. Mills. S. Rothschild supplied the PSP description in Section 5 and W. Redman prepared the CM&M description in Appendix A. J. McCoskey provided Section 5.4, contributed to the data compilation in Section 6, and developed parts of the tradeoff model in Section 7 and Appendix B.

The material in Sections 2 and 3 is updated from an earlier interim report prepared by Dr. M. Mohajeri, who is now with INTELSAT.

Table of Contents

Table of Contents (Continued)

Table of Contents (Continued)

List of Illustrations

List of Illustrations (Continued)

List of Illustrations (Continued)

List of Illustrations (Continued)

List of Illustrations (Continued)

List of Tables

List of Tables (Continued)

1. INTRODUCTION/SUMMARY

This final report summarizes the recent work performed by the Communications Satellite Corporation (COMSAT) on the Atlantic Packet Satellite (SATNET) program under Contract No. F 04701-76-C-0240. This final report covers the time period from the third quarter of 1978 until the end of the contract in the third quarter of 1979. Although COMSAT's participation in the SATNET program dates back to the Fall of 1976, the work during that earlier period is summarized adequately in References [1]–[4]. This report will not cover those earlier phases of the contract.

Early in 1978, COMSAT started work under a fourth modification to the contract to design, develop, fabricate and install new packet satellite program (PSP) terminals at the Goonhilly and Tanum earth stations. A third PSP terminal for the Etam earth station was to be fabricated and installed under separate funding. The new PSP terminals were planned to replace the temporary equipment configurations at each earth station which utilized SPADE channel units, and prototype SPADE/SIMP Interface (SSI) modules, to conduct the SATNET experiments. Additionally, the new terminals would contain more complex (Phase-II) interface modules, allowing a variety of coded/uncoded and QPSK/BPSK packet-transmission-mode combinations to be selected when a small earth station must participate in SATNET. Finally, the new terminals would include a new command and monitoring module (CMM) to allow the SIMP to both change the configuration of the PSP terminal, and to extract measurements from the modems and interfaces to monitor status and assess performance.

The development, fabrication, and installation of the new PSP terminals has turned out to be a much more complex job than originally estimated. As a result, the time and resources required to complete the last modification to the contract have exceeded those original estimates by more than a factor of two. It follows that almost all of COMSAT's effort during the latter part of the contract has been devoted to completing and installing the PSP terminals. This final report reflects that emphasis.

Sections 2 and 3 of this report summarize earlier descriptive and technical material relative to the UET (unattended earth station) implementation, and early study activities in support of the SATNET experiment. This material is extracted (with some updating) from the second interim progress report [4].

Section 5 of this report includes a description of the new PSP terminal as delivered to the three large earth stations. Each subassembly is described briefly and the formats of the various packet transmission and reception modes are defined. Test data on the microprocessor-based QPSK modems included in the terminals are also included in this section. Finally, Section 5 concludes with a functional description of the CMM; more detail on this module is contained in Appendix A to this report.

In addition to the work directly related to the new PSP terminals, three other activities have been pursued by COMSAT, as time and resources permitted, during the last year of the contract. First, work has continued on SATNET experiments and the software related to initiating and monitoring these experiments. A summary of these activities through early 1979 is contained in Section 4 of this report. Secondly, a sustained effort has been necessary to monitor SATNET performance, to respond to reports of link problems, and to document the results of the necessary corrective actions. This work has required close coordination with the network control center at BB&N in Boston, network operations at COMSAT Headquarters, and, in some cases, the earth stations. The data collected and the problems experienced since January 1979 are summarized in Section 6 of this report. Finally, Section 7 of this final report contains the preliminary results of a tradeoff study comparing packet transmission techniques to the alternatives for data transmission by satellite. This tradeoff study constituted a major task in COMSAT's original contract. Unfortunately, work on this study was always preempted by higher priority work. Consequently, the material included in Section 7 is the result of only a brief activity to develop an economic tradeoff model to assess the effectiveness of packet transmission by satellite. The approach to modeling is described in Section 7 of this report with details of the computer program given in Appendix B. The preliminary results are illustrative of the tradeoffs that could be examined with further development and refinement of this model.

In the remainder of this section, a brief summary will be given of the three major areas of activity by COMSAT during the last year of the SATNET contract. These areas are: PSP installation, SATNET channel monitoring and maintenance, and the preliminary tradeoff study. The intent is to summarize the major events and conclusions in each area.

1.1 PSP DELIVERY AND INSTALLATION

The general milestones in PSP delivery and installtaion are shown at the top of Figure 1-1. The terminal was delivered to the Etam earth station early in the week of 27 June and installation and loopback tests were successfully completed during

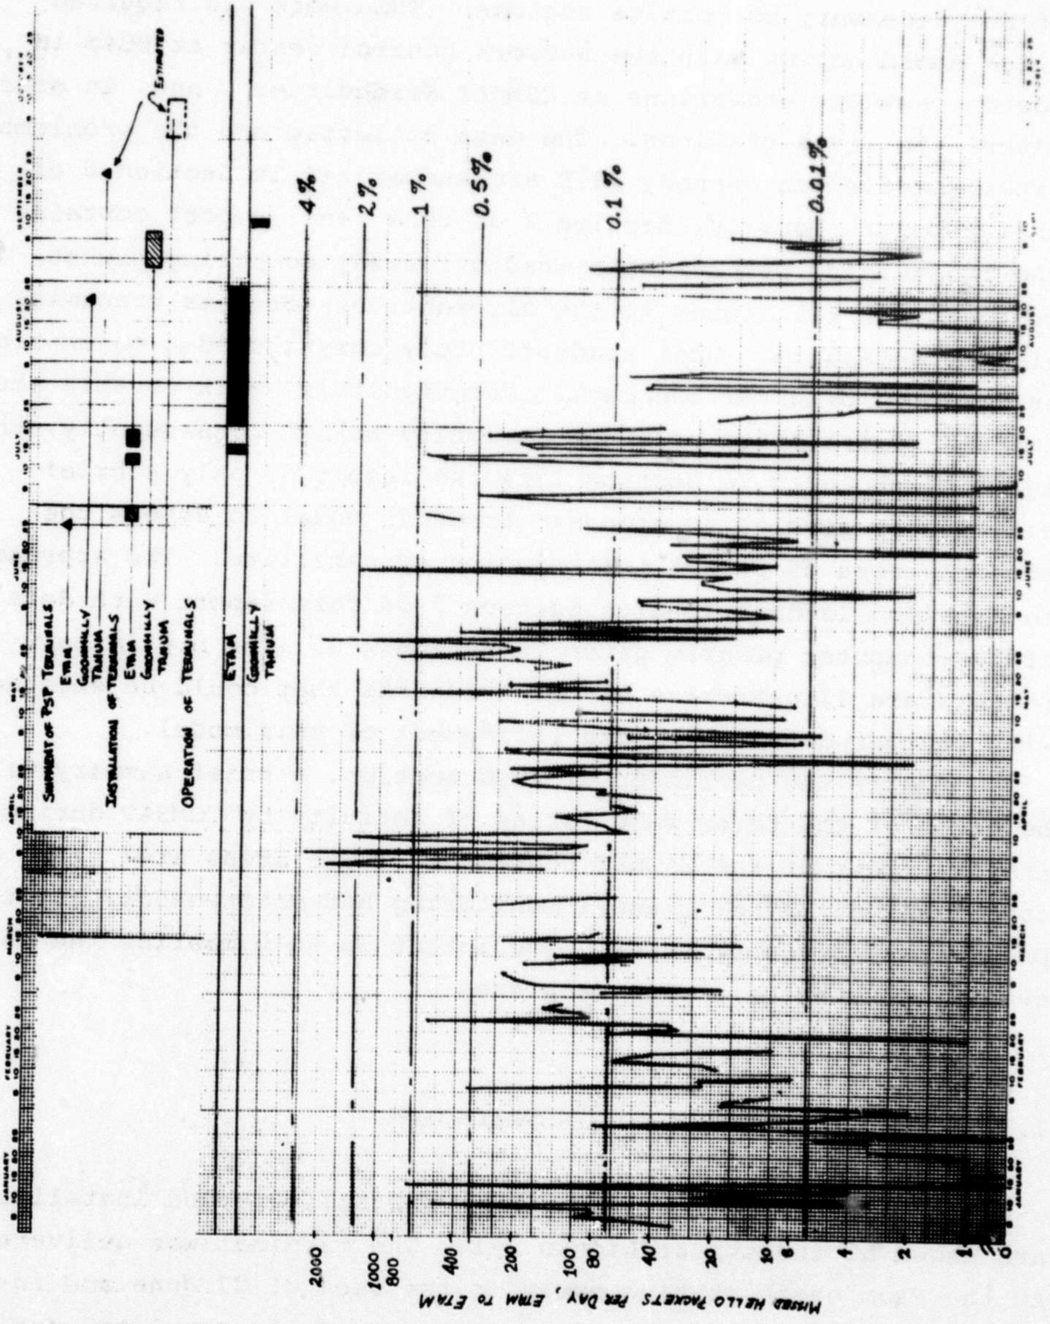

Figure 1-1. Missed Packets Per Day, Etam to Etam

that week. Operation in the network was not successful due to an incompatability between the SIMP/COMSAT/Linkabit interfaces which made the PSP terminal unable to receive the transmissions from the European stations. A hardware change was made to the interfaces at all three earth stations and the PSP operated in the net successfully for a few days. An outage of several days was experienced on 15 July when the PSP terminal stopped and it was necessary to restore the network using the SPADE backup system. The problem was traced to the microprocessors in the interfaces which could not be reset on SIMP command via the CMM. A change was made to the EPROMS in the CMM and the PSP was restored to the network. Note in Figure 1-1 that the PSP operated normally for a period of 36 days. Near the end of August, SIMP software and hardware changes at Goonhilly again caused the Etam terminal to miss a large number of the bursts transmitted from the European stations. The SPADE backup system was once again activated to restore the network.

The lower part of Figure 1-1 gives the number of hello packets missed per day on the link from Etam to Etam. The data covers the time from the first of the year to date and includes the periods before PSP installation, during the interval of more than a month when the PSP was operating at Etam, and finally the short recent period when the PSP terminal was removed from the network. Early in the year, this link experienced a highly variable miss rate on the approximately 66,000 hello packets transmitted and received each day. On certain days, more than one percent of the hello packets were missed; on other days, the miss rate is lower than one-tenth percent. The long term average of the miss rate over the first half of the year is in the range 0.2–0.3 percent. This is approximately the same as the other links among the three large earth stations.

There is some indication from the data in Figure 1-1 that the PSP terminal improves the missed hello packet rate by as much as 10-to-1. Note that during that part of the month of August when the PSP was in the network, the miss rate for the Etam-to-Etam link never exceeds 0.1 percent. Unfortunately, the data collection interval is too short to draw any firm conclusions; but the preliminary indications do show some improvement.

The milestones at the top of Figure 1-1 show that the second PSP terminal was shipped in late August to Goonhilly with installation taking place in the first two weeks of September. The third terminal is scheduled for shipment to Tanum in the fourth week of September with installation taking place in early October.

1.2 SATNET MONITORING

A sustained effort was started in January 1979 to monitor and summarize the daily SATNET channel summaries and to compare the results to a model of the performance to be expected in the SPADE transponder of the INTELSAT IV-A satellite. These detailed results are given in Section 6 of this report. This monitoring has concentrated on the missed-hello-packet data for each of the nine links among the three large earth stations. These links utilize a single 64-kbit/s carrier in the relatively sparsely occupied "SCPC" segment of the SPADE transponder. Each SATNET link is set up with a 17.5-dB received carrier to noise ratio as measured at Etam. The setup includes an equalization of the transmitted power levels in the bursts from each of the three stations so that they are equal, to within 1 dB or better, as measured at Etam. This balancing of the burst levels is an important requirement. For example, antenna tracking problems at Goonhilly in February 1979 caused its bursts to be low, by at least 2 dB, relative to the level of the other bursts. Such departures resulted in a large increase in the missed packet rate for transmissions originating from, and received at, Goonhilly. On several days in February these rates exceeded 10 percent on one or two links—such performance causes serious disruption to the network. A more subtle problem, although it was resolved quickly, was noted at Etam on two occasions since the beginning of the year. The symptom was a tendency for Etam to miss hello bursts from certain stations depending on the position of the bursts in the hello subframe. The miss rate was much higher for bursts preceded by "silence," i.e., bursts that were not immediately preceded by another burst in either the hello or information subframes. This problem was resolved by alignment of the VCO in the SPADE channel unit to be within a few hundred Hz of the nominal frequency (the VCO was not actually out of specification limits, but was off nominal by 1200 Hz). This problem was encountered twice, each time building up over a period of several days. Once the Etam personnel were notified of the problem, corrective action was taken quickly.

Other than the power level problems at Goonhilly, and the oscillator adjustments at Etam, no other satellite-link-related problems were noted during the first eight months of 1979. On certain occasions, missed hello packet performance has seemed to be sensitive to time of day, and correlated among links (indicating possible dependence on traffic activity and possibly interference in the satellite transponder) but such occurrences are transitory and infrequent. When such events are observed in "yesterday's data" it is too late to search for a cause/effect relationship in the transponder.

Although satellite-link-related problems have been relatively rare, one must conclude from the data in Section 6 that the missed hello packet data is highly variable from day to day and from link to link. The reason for this is partly explained by the analysis in Section 6. First, it is useful to define a "missed hello packet" more carefully. As the data is presently collected, these misses include those situations where a packet is acquired and demodulated, but contains one or more bit errors, as well as those cases where the receive modem fails completely to detect or acquire synchronization on a received packet. Since the hello packets are reasonably long, and contain almost 1000 bits that are protected by the CRC error detection implementation, the miss rate is highly dependent on the probability of bit error in the receive modem. Very roughly, one can relate missed packet probability to bit error rate (p), as $\Pr\{miss\} \cong Kp$, where K is a large number. Prior to a change in the format of the hello bursts in early April 1979, the constant K was approximately 10^3. Thus, for the early data in 1979, $\Pr\{miss\} \simeq 1000\ p$. This relationship gives a rough correspondence between missed packet rates of 0.1, 1.0, and 10 percent and bit error rates of 10^{-6}, 10^{-5}, and 10^{-4}, respectively. The 64-kbit/s carriers operating in the SPADE transponder, including the SATNET carrier, would be expected to experience this range of bit error rates. Most of the time, the error rate would be 10^{-6} or better; degradation to 10^{-5} would be expected to occur occasionally, and further degradation to 10^{-4} would occur rarely. Such a variation would be perfectly acceptable on these channels where the 64-kbit/s carriers convey PCM speech. On the SATNET channel, the difference between a 0.1-percent miss rate (66 hello packets missed per day, at $p = 10^{-6}$), and 1.0 percent (660 misses per day at $p = 10^{-5}$) is significant. Performance levels with 1.0-percent misses is usually cause for alarm. Without some type of error correction, bit errors will occasionally occur in the long hello packets, and these events, along with the cases where modem fails to acquire the packet altogether, will both contribute to the packet miss rate. Based on this line of reasoning, which is supported by the analysis in Section 6, the seemingly erratic missed packet behavior observed on SATNET is not that surprising. As noted earlier, there are cases where the miss rates are excessive, and equipment alignment or adjustment was necessary to remedy the situation.

1.3 TRADEOFF STUDY

The tradeoff study contained in Section 7 of this report utilizes a model developed to examine the relative cost effectiveness of three alternative transmission techniques by communication satellite. The techniques are packet transmission, circuit switching, and circuit switching with activity concentration. A nominal scenario is selected where 1000 network users generate a total traffic volume of 8×10^{12} packet/year (where a "packet" is assumed to consist of 500 information bits). The users are assumed to consist of three classes; high, medium, and low rate users. The 1000 total users are divided so that a certain fraction is of each class, and likewise each user class generates a certain fraction of the total annual traffic volume. Full connectivity is required among users of the same class.

Two important parameters of each user class are the utilization (fraction of the day that each user is "off-hook"), and the "activity" of his data source, given that he is "off-hook" or utilizing the network. With pure circuit switching, users are assumed to be assigned to satellite circuits and then retain this circuit for the duration of the "call" or session regardless of their activity. An alternative requiring more equipment complexity and cost at the earth stations is to concentrate "busy" circuits with low activity. Such techniques are well known for speech traffic with digital TASI and DSI* being examples. The final technique is packet transmission where the blocks of data are addressed to the recipient without the need for establishing a "circuit." As expected, the utilization factor, u_K, ($u_K \leq 1$), and the activity factors, a_K, ($a_K \leq 1$) have a major impact on the relative effectiveness of the three techniques.

Variables of the satellite network include the satellite beam coverage and the number of earth stations in the network. Network users are assumed to be confined to CONUS, and these users are equally divided among 20, 40, or 100 earth stations. The satellite is assumed to provide either single CONUS-beam coverage, or spot beam coverage. In the latter case, narrow beams are assumed that can, in effect, point at the individual earth stations. Operation at either 6/4 or 14/11 GHz is assumed, although for the level of detail presently included in the model (availability due to rain loss is not included), the results are not significantly different for the two bands.

Three different mixes of utilization/activity among the users have been examined to date. The first is a baseline (to calibrate the model) where all utilizations and activities = 1.0. Using cost-per-packet as the effectiveness measure (actually $ per 10^6 packets) circuit switching is the most cost effective technique for this situation.

For another extreme situation, where the users have relatively low utilization, but high activity when busy, circuit switching is also the most effective technique.

*TASI = Time Assignment Speech Interpolation; DSI = Digital Speech Interpolation.

For an intermediate case with moderate utilizations, and moderate to low activity factors, the packet transmission technique is more cost effective than the other techniques. This advantage is more pronounced for the spot beam satellite, when spacecraft resource cost is at a maximum. Generally, the advantage of packet transmission relative to circuit switching with activity concentration is in the range of a 50-percent reduction in cost per megapacket for the former technique. This advantage is moderately sensitive to the assumed costs of the earth station equipment to accomplish the packet transmission. For example, if the models used underestimate such costs by a factor of 2-to-1, then the advantage of packet transmission is essentially eliminated.

As noted above, the model developed for the tradeoff study is the result of a limited effort and thus should be considered as preliminary. Suggested areas for extension and refinement of the model are given at the end of Section 7.

2. SMALL STATION IMPLEMENTATION ACTIVITIES

COMSAT is implementing a "small" UET* packet satellite terminal located at COMSAT Laboratories in Clarksburg, Maryland, to participate in the DARPA Atlantic Packet Satellite Experiment. This earth terminal is small relative to the Standard A INTELSAT earth terminals (e.g., Etam, Goonhilly, and Tanum) in that its figure of merit, G/T, is 11 dB less. Alternative strategies for this station's participation in the experiment have been studied and documented in the Technical Operating Report, "UET Implementation Planning," ASAP-77-1, January 7, 1977. The First Interim Progress Report, ASAP-77-4 [1], describes the chosen alternative.

In the planned final equipment configuration shown in Figure 2-1, the item labeled UET represents the earth terminal and the equipment involved in the RF-to-IF and IF-to-RF translation. The item labeled modem represents the modem and the associated support electronics, and the MSI denotes the modem/SIMP interface discussed in Section 5. The SIMP is the same minicomputer located at the large earth stations (Honeywell 316 with 32K words of memory) and will use essentially identical software. The gateway, a PDP-11 computer, will be used to perform and support certain measurements. Figure 2-2 shows the planned physical layout of the equipment.

The main activities associated with the UET support task have included the following:

a. testing the modem which was developed for this application;
b. verifying that the IF and RF equipment are operating properly by loopback testing over the INTELSAT IV-A satellite;

*Unattended earth station.

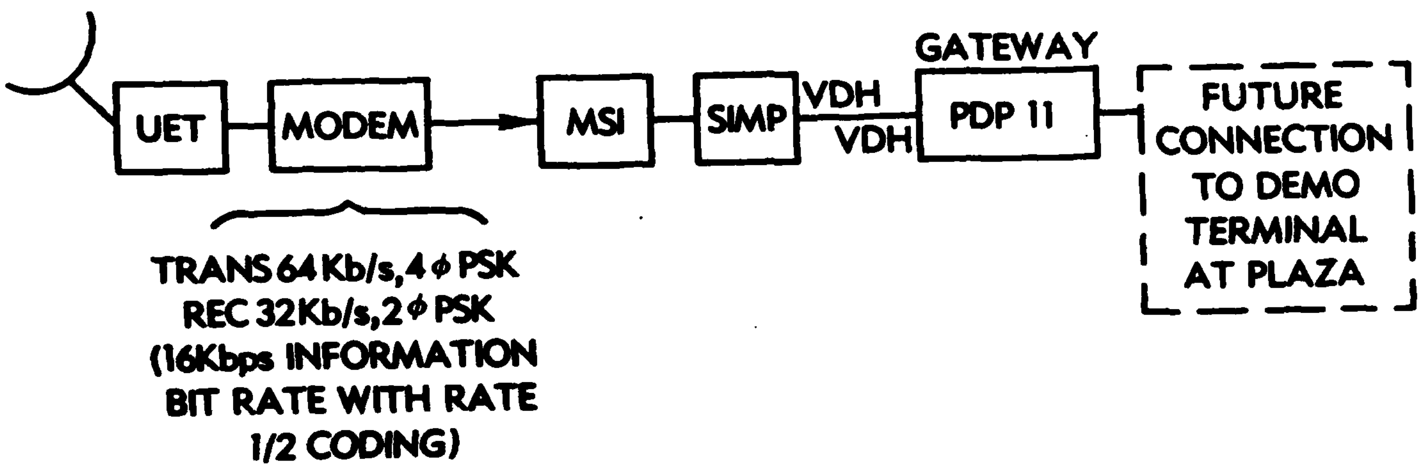

Figure 2-1. Small Station Configuration at Clarksburg, Maryand

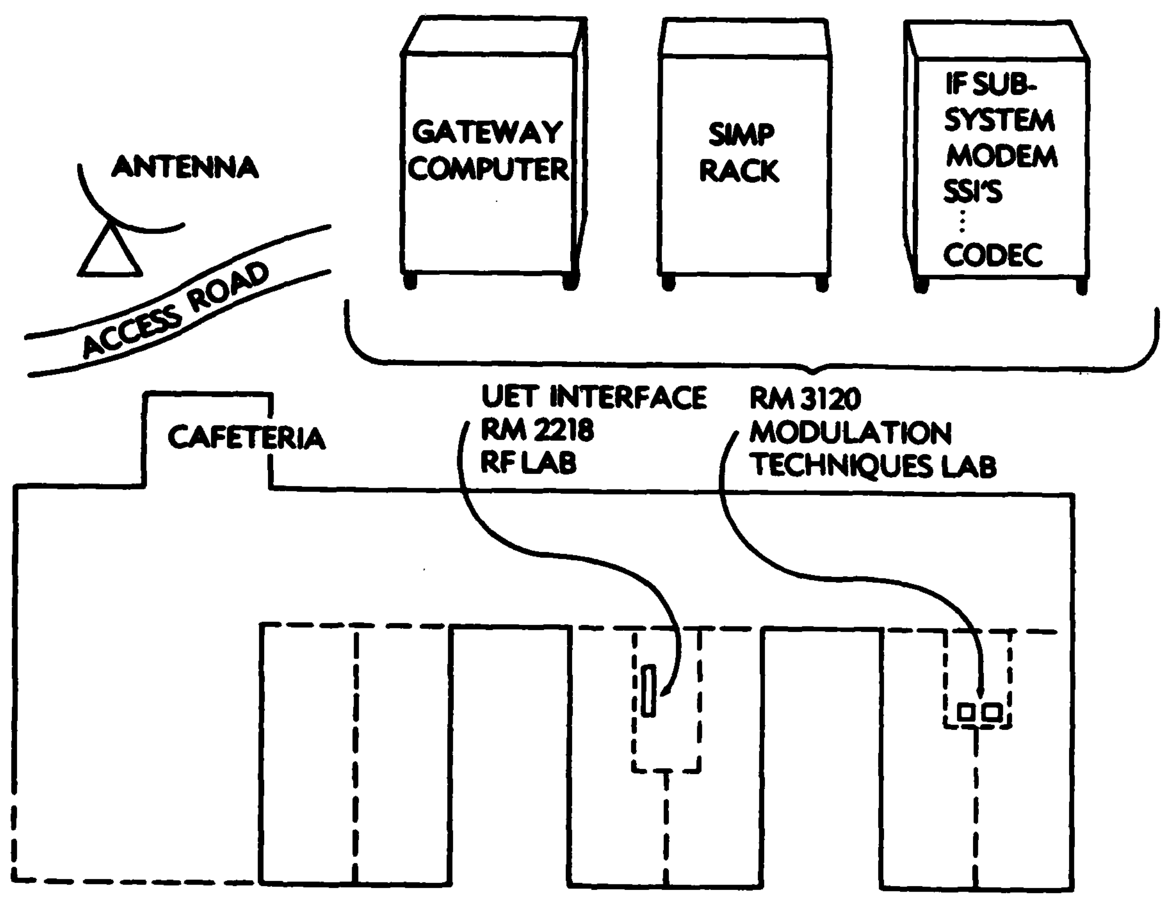

Figure 2-2. Small Earth Station Physical Layout

c. completing the development of the MSI and the associated test equipment (the data test set);

d. interfacing the MSI and the SIMP; and

e. achieving the final configuration and verifying the operation of the gateway-SIMP data path.

2.1 MODEM TESTING

The modem to be used at the small station was developed by the Linkabit Corporation under contract to COMSAT and was delivered in May 1977. The specification for this modem is included in the Technical Operating Report, "UET Implementation Planning."

The modem is based on a digitally-controlled microprocessor implementation and has been specifically designed to minimize the time required for burst mode acquisition in a low signal-to-noise ratio environment. SPADE compatibility in terms of its IF characteristics is required with the capability of transmitting and receiving any SPADE channel. However, the initial version of this modem, which has been tested, has only a BPSK receive capability.* The modem provides a 3-bit quantized output and can perform a hard decision start-of-message (SOM) detection internally. After SOM detection, the phase ambiguity at the output of the BPSK demodulator is resolved. When the modem is utilized with the MSI, SOM detection is performed in the MSI, and the SOM detection internal to the modem is disabled by replacing a read-only memory (ROM) element.

The modem is able to minimize the time required for burst acquisition because the same preamble symbols are used to

*Thus, the modem is capable of operating at a maximum burst information bit rate of 32 kbit/s in the SPADE channel.

acquire both carrier and bit timing. The processing is completely digital after the second frequency translation, including the use of parallel digital filters which resolve frequency uncertainty. Modem operation has been described in the technical literature [5]. Later versions of this modem will support QPSK reception and incorporate the monitoring capabilities discussed in Subsection 5.4. These modifications are scheduled to be made in August 1979.

To establish the basic capabilities of this modem, which serves as the prototype for the QPSK modems developed for the PSP terminal, extensive testing was performed in back-to-back, IF loopback, and satellite loopback configurations. The tests involved bit error rate performance as a function of signal energy per information bit to noise spectral density, (E_b/N_o), burst acquisition performance as a function of E_b/N_o, and sensitivity of the modem to the absence of bit transitions within a packet.

For adequate testing, a modem test set was required which could emulate packet transmission and reception (i.e., simulate the operation of the MSI and SIMP combination) and could provide sufficient flexibility to support the tests. Moreover, a significant effort was necessary to calibrate the measurements, particularly the value of E_b/N_o because of the unique modem implementation which made accurate E_b/N_o measurements difficult when the tests were performed through the satellite link.

2.1.1 MODEM TEST SET AND EXPERIMENTAL CONFIGURATION

A fairly elaborate test set, which is depicted in Figure 2-3, was developed by the Modulation Techniques Department

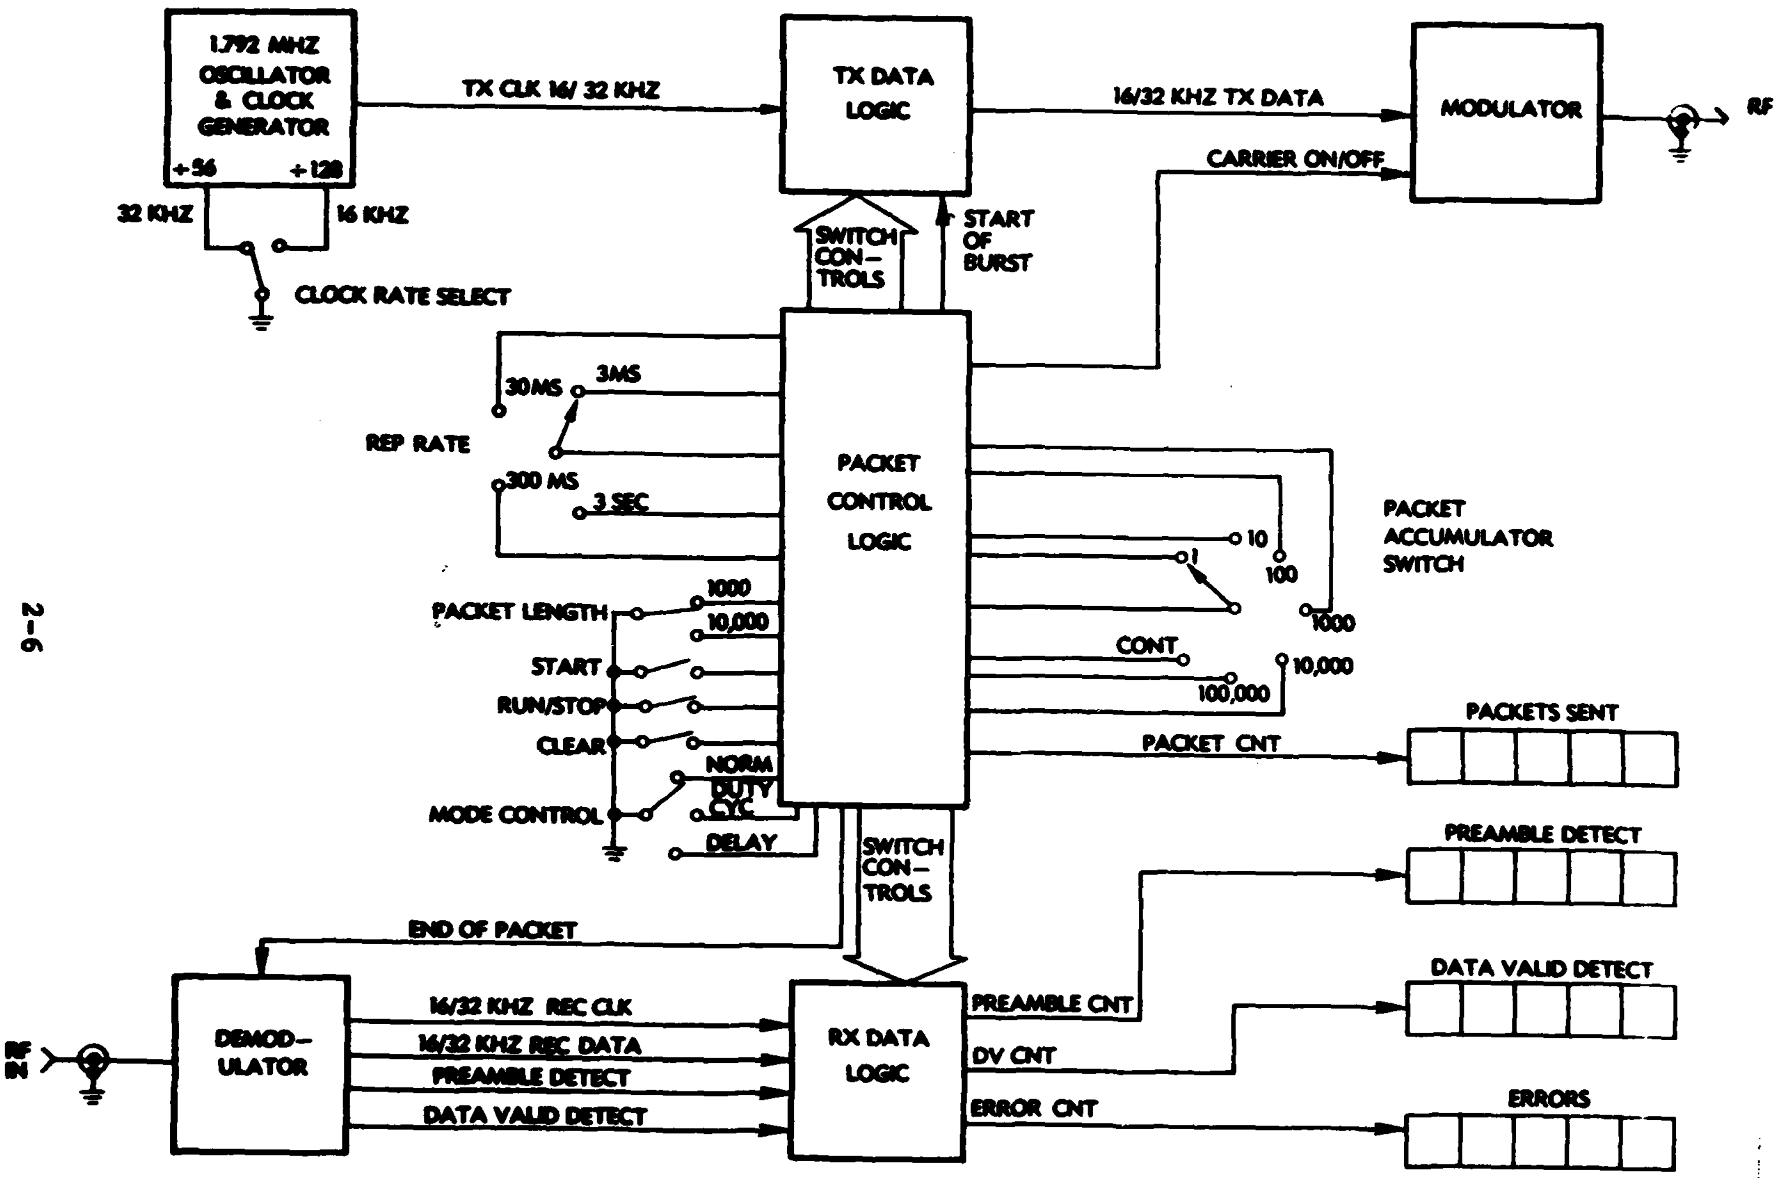

Figure 2-3. Block Diagram; Linkabit Modem Test Stand

at COMSAT Labs to evaluate the Linkabit modem. It was necessary to build this test set internally because an adequate instrument was not available on the market, and an existing test set could not be modified to yield satisfactory performance. The TDMA environment necessitated the generation and the input of packets or bursts of data to the modem. Furthermore, control signals, such as carrier on/off and end-of-packet signals, were necessary to interface the test set and the modem. Thus, a transmit modulated carrier could be synchronized with the data sequence, and the demodulator could be notified upon packet termination so that reinitialization could occur.

One parameter of interest is the off-time between packets (i.e., interpacket interval). The minimum time required to reinitialize the modem to prepare for a subsequent packet is approximately 0.5 ms; therefore, varying the off-time from 0.5 ms to about 5 s is desirable for the detection of any discrepancies. The duration of the packet can be either of two fixed lengths, 1 kbit or 10 kbit, with a 128-bit preamble. The data consists of a 1023-bit pseudorandom sequence truncated at 1 kbit; to accommodate the longer length, ten of these sequences are concatenated together. Other repetitive sequences, which can be substituted for the data, range from an on/off square wave at the fastest achievable rate, R/2, to very low duty cycles, including a DC level. In the latter case, performing tests such as bit error rate sensitivity to the stability of the recovered clock is feasible. Another feature is the delayed data mode, in which the observation of the data sequence can be started at some prescribed delay interval after the preamble; therefore, bit error rate tests may be conducted as a function of the interval over which the recovered clock remains stable. However, the bit error rate sensitivity measurements proved adequate in testing clock recovery so that the delayed (gated) error rate measurement capacity has not been used to date.

To determine the percentage of packets lost, the test set is also equipped with a burst counter which makes it possible to send a packet individually or to accumulate as many as 100K packets for long term tests. During this measurement phase, it became apparent that the fraction of packets lost was a function of the frequency offset in the overall transmission link. Hence, a variable frequency offset was introduced.

Both bit error rate and packet loss probability measurements were performed in three progressively more complex configurations, back-to-back, looped through an IF subsystem (as shown in Figure 2-4), and by means of the satellite link. These results are shown in Figures 2-5 and 2-6.

2.1.2 MODEM TEST RESULTS

Figure 2-7 summarizes the bit error rate results of the various modem tests over an operating range of E_b/N_o = 0–6 dB. Performance is slightly degraded (0.2–0.3 dB) when the communications path is extended to include the IF subsystem and when the satellite path is added (~0.2 dB). At a 3.5-dB nominal operating point, modem performance is within 1.2 dB of ideal even for the satellite loopback case. Additional deviation at larger values of E_b/N_o is expected; however, some of this additional degradation may be due to the method of setting the AGC level in this particular modem. Optimizing performance at large values of E_b/N_o was not attempted because operation in conjunction with an encoder/decoder at a nominal E_b/N_o value of 3.5 dB will be the nominal operating point for the UET modem.

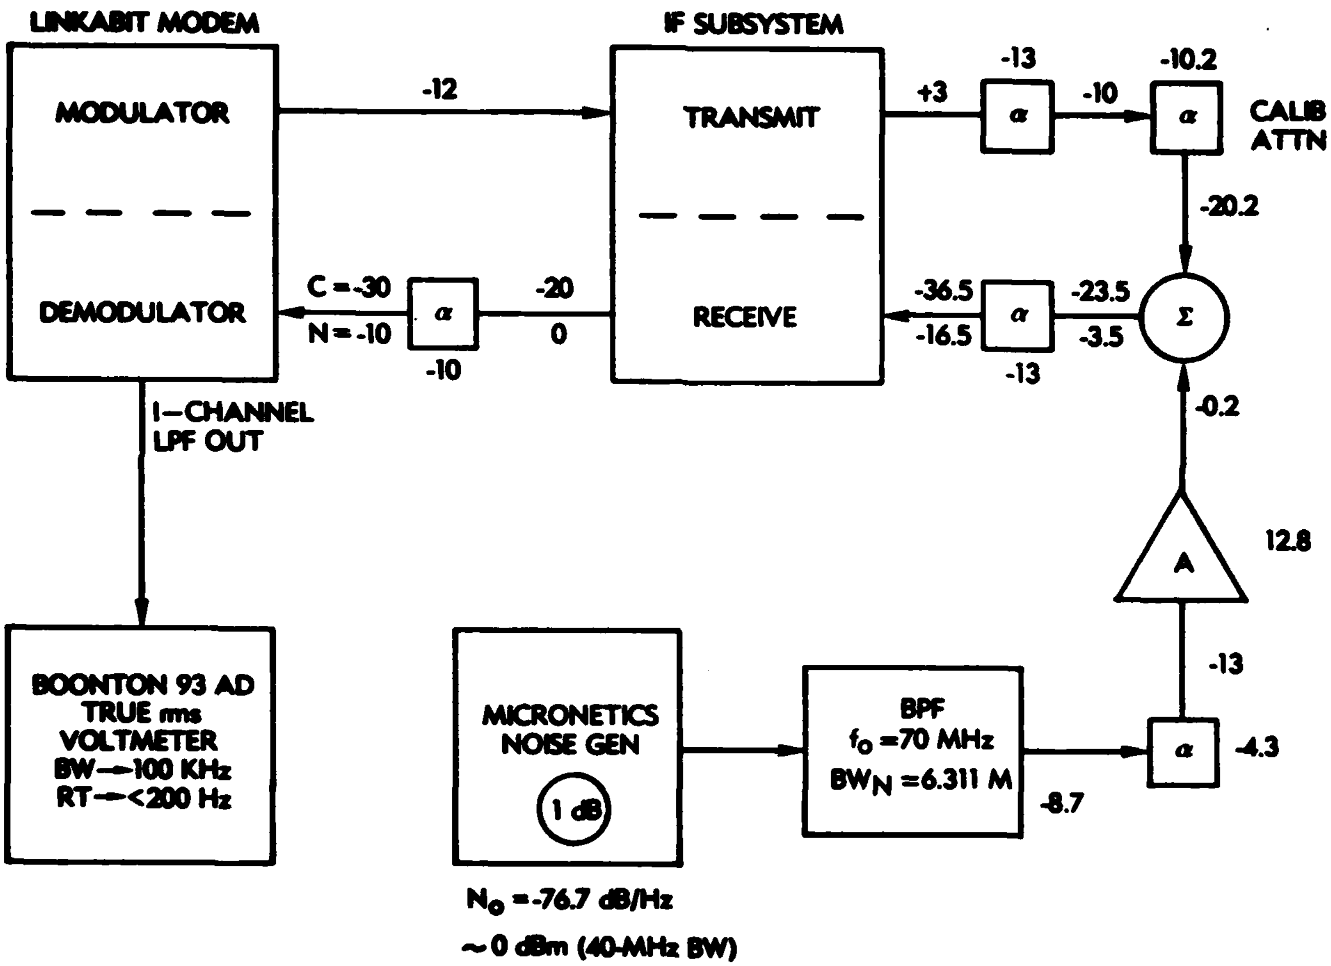

Figure 2-4. Packet Test Setup With IF Subsystem (all levels are dBm)

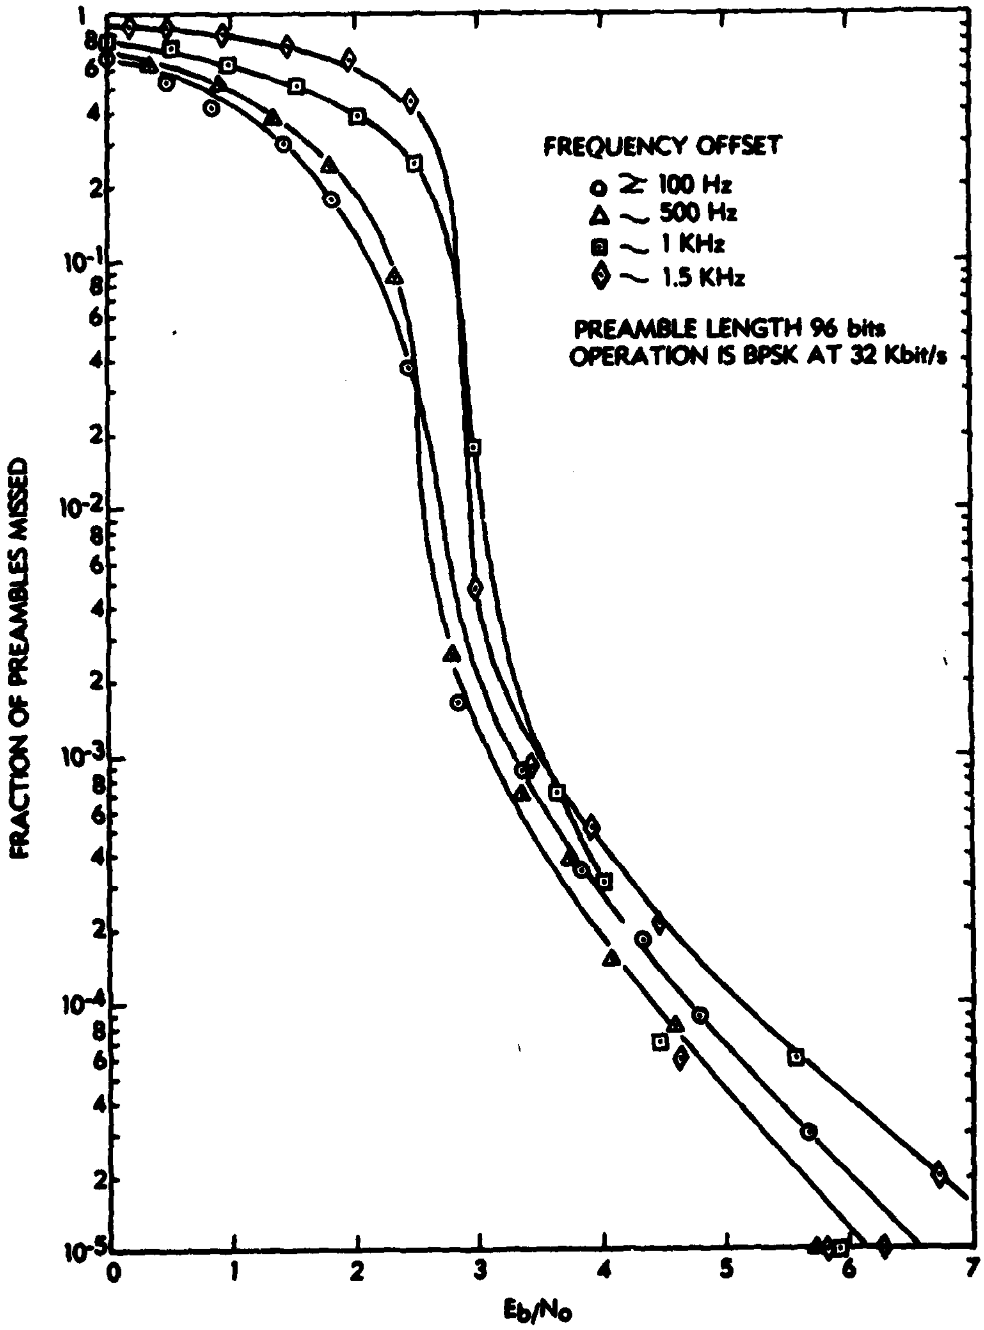

Figure 2-5. Preamble Loss Probabilities vs E_b/N_o and Frequency Offset

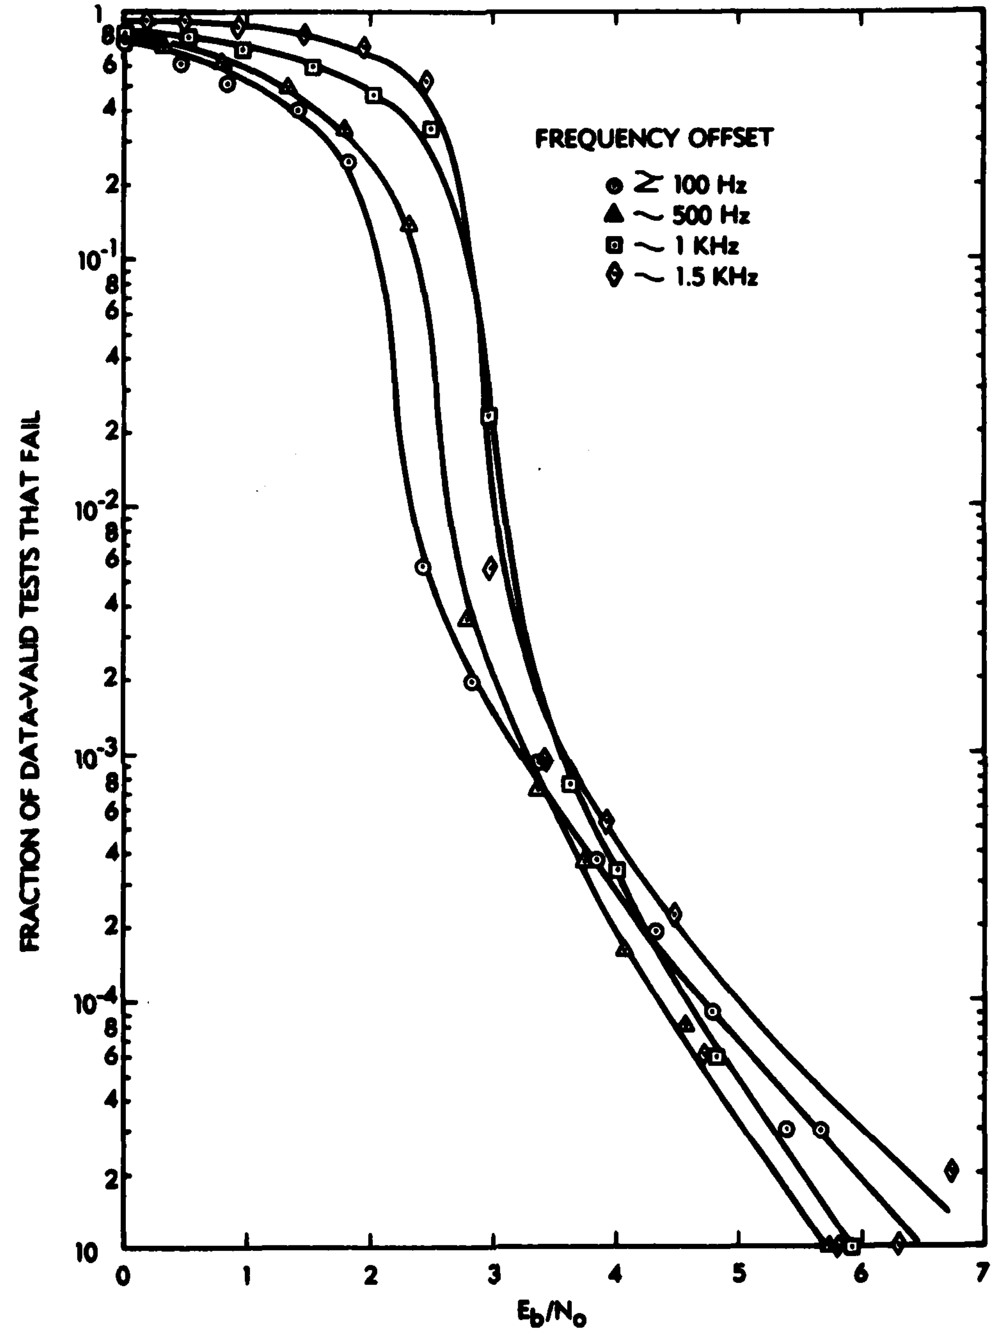

Figure 2-6. Data Valid Loss Probabilities vs E_b/N_o and Frequency Offset

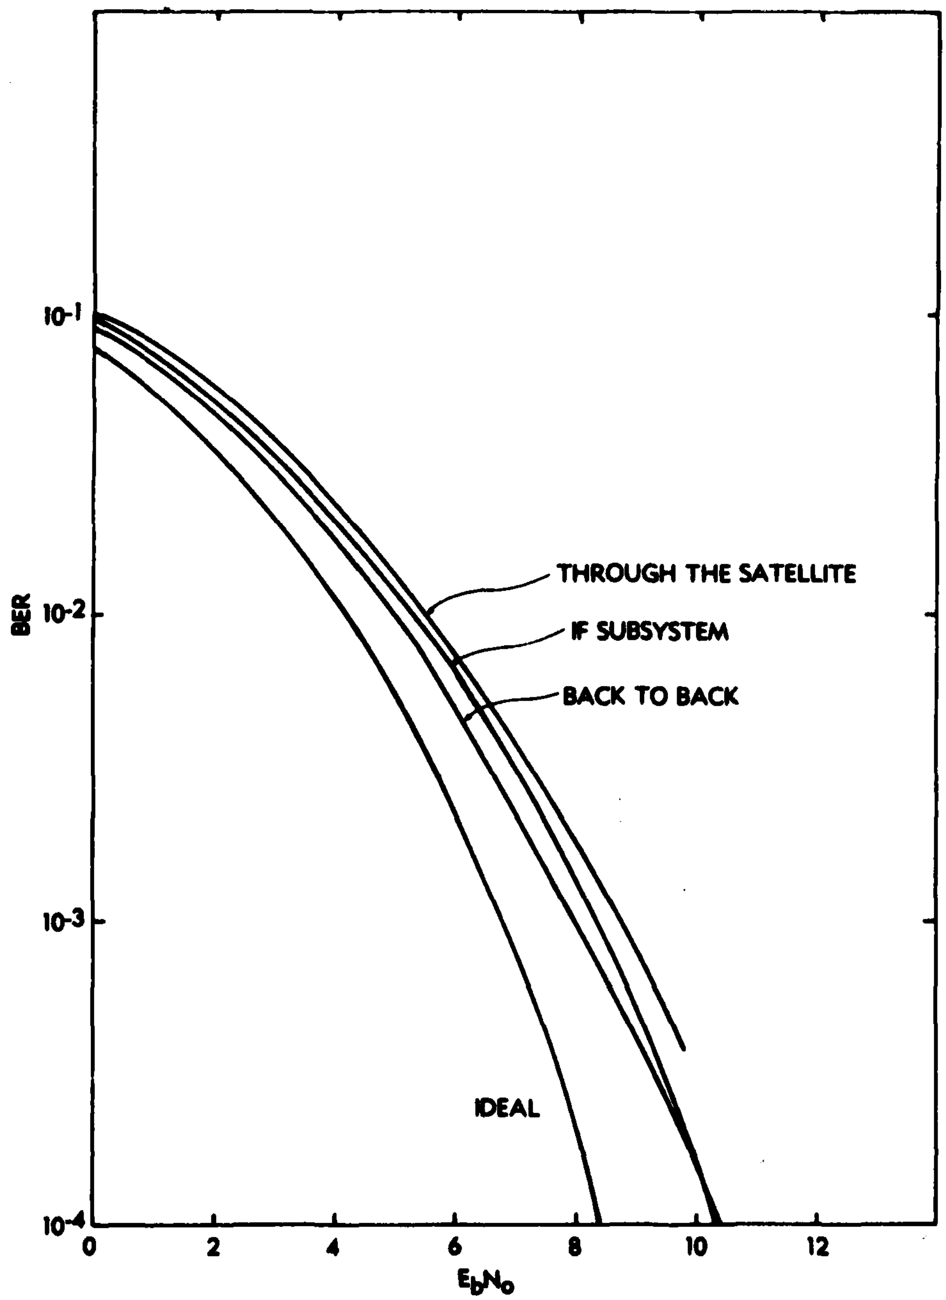

Figure 2-7. Bit Error Rate Results

handled its coding and decoding functions. The tests are significant because of the unique and previously untested manner in which the MSI communicates with the SIMP interface (using a 128-kHz burst clock).

The interface of the MSI to the modem, which was not accomplished during the May tests, was delayed until the backplane wiring on the small station terminal could be modified to reduce the level of interference caused by the 1.796-MHz signal used in the system operation. This interface is not expected to present difficulty because of its simplicity. After a few additional features are incorporated in the MSI microcode at the Linkabit Corporation (SIMP-I capability, retransmit mode, mode 7, duty cycle control, and contention dummy packet) and the small station chassis is rewired to minimize interference problems, the MSI will be permanently installed.

3. STUDY ACTIVITIES IN SUPPORT OF THE EXPERIMENT

During the course of the SATNET program, a number of study activites have been performed, and most of the protocol specifications and techniques to be implemented and tested in the Atlantic Packet Satellite Experiment have been finalized. COMSAT has been an active participant in this process. In addition to these group activities, certain specific studies have been conducted, often in direct support of the development of experimental capabilities, including studies of contention detection and techniques to permit small station operation at a 16 kbit/s data rate with minimal or no hardware modifications required at the large stations. A major problem in the latter area is the design of an appropriate SOM sequence. Finally, a preliminary study of satellite networking alternatives has established some interesting relationships between network realizations and the type of multiple-access techniques employed.

3.1 GROUP STUDY AND DEFINITION ACTIVITIES

A number of small working group meetings have been held during SATNET development to define various aspects of the SATNET protocols and implementation details. Regular participants at these meetings were Linkabit, BBN, and COMSAT with the other experiment participants attending on topics of particular interest. The major topics were the host-SATNET protocol, SATNET flow and congestion control strategies, modifications to the PODA algorithm to improve its performance, and mixed station control strategies.

The host-SIMP protocol specifies the rules for passing data and control information between the SIMP and the host computer, which may be a gateway or a computer attached directly to the satellite network. The discussions concerned both the final definition of this protocol for block traffic, and the specification of an appropriate version of the protocol for stream traffic. Specifying SATNET capabilities from the user viewpoint was particularly important. These discussions were documented in Packet Satellite Program Working Notes (PSPWNs), which are normally authored by one of the participants as a record of the issues and decisions. The relevant PSPWNs are PSPWN 100, "SATNET Host Access Protocol, Version 1," by Richard Binder of BBN, January 26, 1978; and PSPWN 104, "Host/SATNET Stream Access Protocol, Version 1," by Richard Binder of BBN, April 1978.

The flow and congestion control strategies for SATNET remain under discussion, and definitive decisions await the availability of simulation results from Linkabit. Some of the ideas which have been discussed are contained in PSPWN 82, "A Possible Approach to Flow Control, Buffer Management and Acknowledgment for a Broadcast Packet Satellite Network," by Lin-Nan Lee and Nicolas Abel of Linkabit, September 11, 1977.

The PODA algorithm, originally suggested by Linkabit, is a demand-assigned multiple-access algorithm being developed for the PSP to accommodate various types of traffic (block, stream, and voice), various delay constraints, and various priorities. Two versions of the algorithm have been implemented. One, termed CPODA, uses random access techniques (slotted ALOHA) in a reservation subframe; the second, FPODA, uses fixed assignment TDMA. Early design considerations are partially documented in PSPWNs 6 and 38. The current implementation is discussed in several papers [6], [7] and documented in PSPWN 86, "Channel

Protocol Module in SIMP II," by Nai-Ting Hsu of BBN, October 27, 1978; and PWPWN 87, "PODA channel Scheduling Synchronization," Nai-Ting Hsu of BBN, October 27, 1978.

The results of the recent PODA design discussions are summarized in the paper entitled "General Purpose Packet Satellite Networks," coauthored by Dr. Jacobs of Linkabit, Mr. Binder from BBN, and Dr. Hoverstein from COMSAT, which appeared in the November 1978 issue of the *Proceedings of the IEEE*. Major modifications include the use of scheduled acknowledgments and multiple packet reservations and burst, as opposed to packet scheduling. After implementation, these modifications will reduce SIMP storage requirements and increase efficiency. This paper also briefly presents the conclusions of the mixed station control discussions.

3.2 SOM DESIGN FOR 32-Kbit/s OPERATION OF THE SMALL STATION

The Phase 2 SSI data modules and the MSI data module have been designed to include certain transmission modes which would allow transmission at a 32-kbit/s data rate to the small earth station.* This rate can be achieved by using coded QPSK transmission. (The channel rate would be 64 kbit/s, and rate one-half convolutional encoding would be employed.)

The available channel signal-to-noise ratio will probably be adequate to support this data rate at the small station at least during periods of good propagation. However, satisfactory operation can only be achieved if the modem acquires the burst preamble and the MSI reliably detects the SOM. The former should not be a problem because the 96-symbol preamble used for

*Note that generally, transmission to the small station would be rate-one half encoded BPSK at an information bit rate of 16 kbit/s. It is possible that coded QPSK transmission could be received (information bit rate 32 kbit/s) under certain conditions.

transmissions to the small station has been proven adequate. However, a longer SOM sequence is required to achieve frame synchronization and to resolve the QPSK phase ambiguity at the low signal-to-noise ratio available at the small station.

The interfaces were designed to transmit two basic QPSK SOMs. SOM-1, which is used for QPSK transmission to a large station, can be varied in length; the available lengths are 16, 20, 24, and 28 bits per channel. SOM-3, which is used for QPSK transmission to a small station, has a single length of 32 bits per channel; this is considered to be the minimum required. However, the SOM detection circuitry used at the large stations, which was designed and implemented as part of the Phase I SSI data module, can accommodate only QPSK SOMs which are less than or equal to 28 bits per channel. More importantly, this circuitry can only be configured to detect a single QPSK SOM. Adding the ability to search for either SOM-1 or SOM-3 would significantly increase circuitry and cost.

The use of SOM-3 for all QPSK transmissions, regardless of destination, would create channel inefficiencies; that is, when the transmission is sent to a large station where SOM-1 would be adequate, the difference in length between SOM-3 and SOM-1 would correspond to wasted channel time. Thus, the design of a SOM-3 which included an appropriate SOM-1 as part of the sequence (in particular, as the last bits of the sequence) was attempted. The use of such a SOM word would provide increased detectability at the small station and require no modification to the SOM detection circuitry used at the large stations. Only the SOM detection circuitry at the small station would require modification.

Two 32-bit words were obtained by block encoding the pair of 16-bit SPADE SOMs through the (31, 16, 7) code generator polynominal

$$g(x) = 1 + x + x^2 + x^3 + x^5 + x^7 + x^8 + x^9 + x^{10} + x^{11} + x^{15}$$

After encoding, the sequences become

$$AW_1 = x^{15}(SW_1) + g(x)\ Q_1(x)$$
$$= 0010\ 0000\ 0111\ 0101\ 1010\ 1101\ 1001\ 1101$$

$$AW_2 = x^{15}(SW_2) + g(x)\ Q_2(x)$$
$$= 0001\ 0100\ 1101\ 1100\ 1001\ 0110\ 0000\ 1001$$

where $Q_1(x)$ and Q_2 are the quotient polynomials which result from dividing SW_1 and SW_2 by $g(x)$. The pair of 16-bit SPADE SOMs are denoted by SW_1 and SW_2. The two 32-bit words to be used for SOM-3 are AW_1 and AW_2. The last 16-bits of AW_1 and AW_2 are identical to SW_1 and SW_2, respectively.

With both preceding and trailing sequences consisting of alternating 1010 . . . patterns, the autocorrelation of AW_1, the autocorrelation of AW_2, and the cross correlation of AW_1 and and AW_2 are shown in Figures 3-1—3-3, respectively. A more detailed discussion is given in Reference 8.

3.3 CONTENTION DETECTION

The Phase 2 SSI data module should be able to declare the existence of a channel contention condition and notify the SIMP using the normal data path. Sensing contentions is potentially useful for at least two reasons: to provide the information input to a stability control algorithm for the slotted ALOHA (control) portion of the CPODA frame, and to permit a station, particularly a disadvantaged (small) station, to reinitiate with

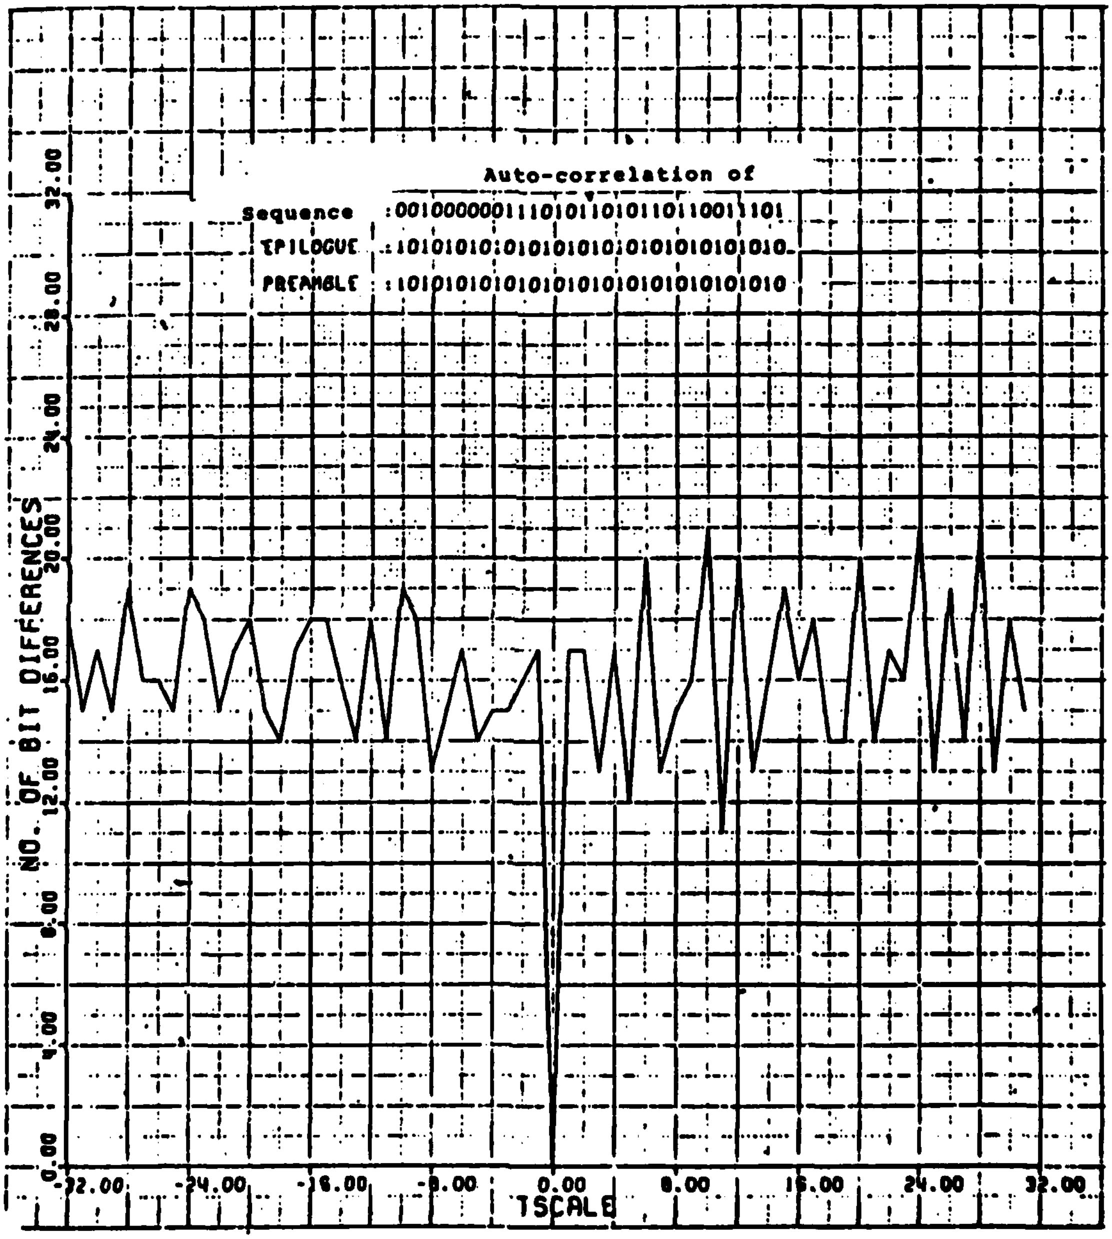

Figure 3-1. Autocorrelation of AW_1

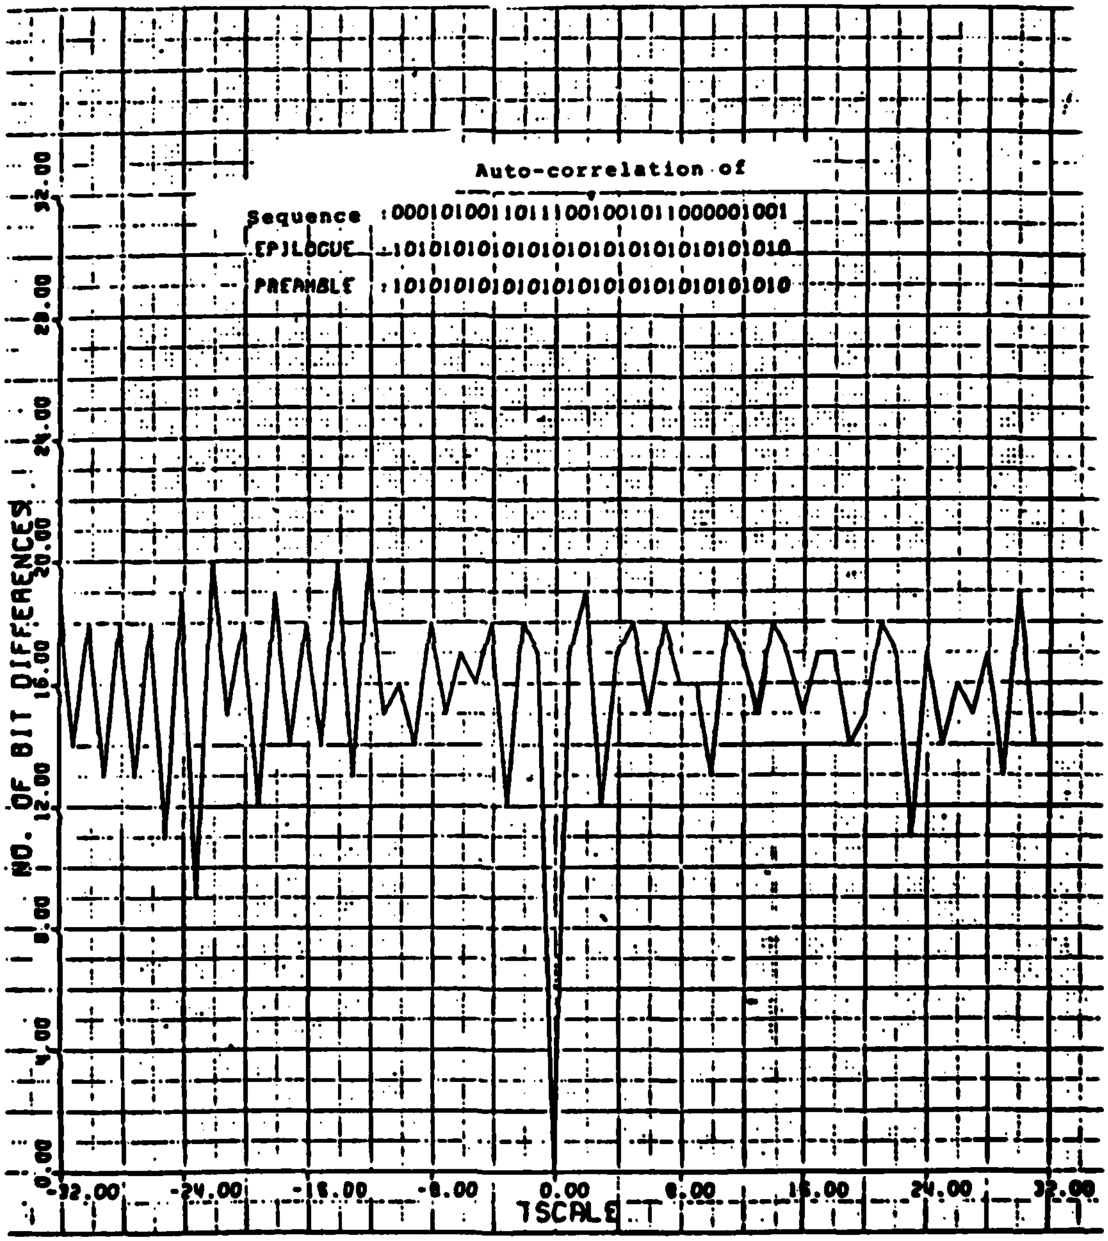

Figure 3-2. Autocorrelation of AW_2

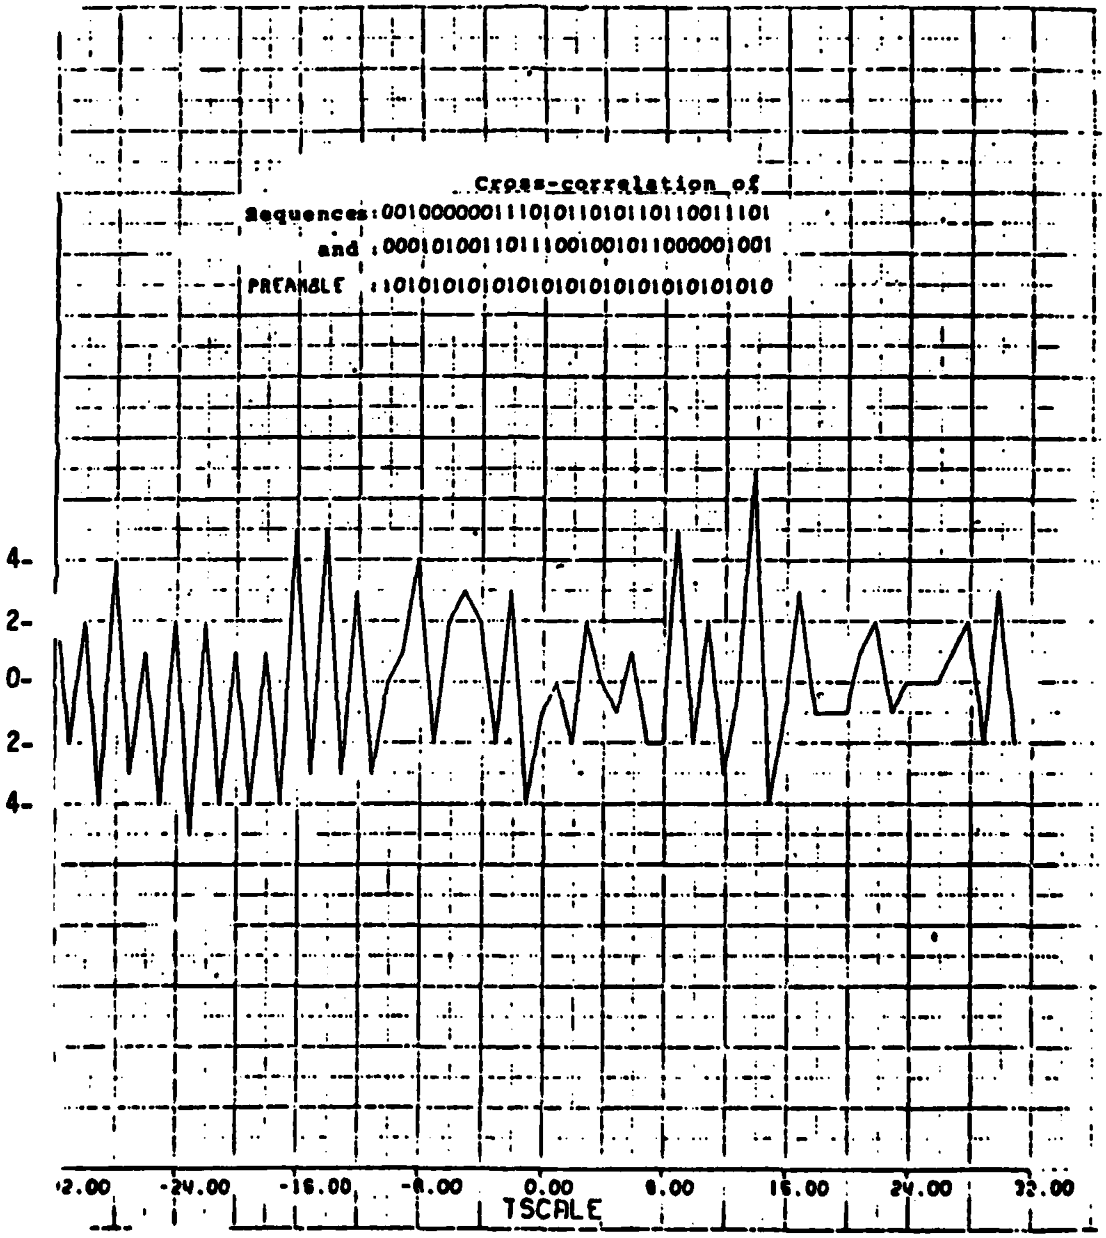

Figure 3-3. Cross Correlation of AW_1 and AW_2

minimal delay, the transmission of a packet that has suffered a collision. Collisions in the information subframe of the CPODA algorithm are unlikely; however, if they occur, reliable and simple detection is improbable. Thus, attention will be concentrated on the control subframe, which uses the slotted ALOHA algorithm and contains transmissions from two or more stations that are attempting to use the channel simultaneously, i.e., whose transmissions overlap in time when they reach the satellite.

To provide maximum information for stability control, a contention should be sensed at all stations. However, prompt reinitiation of the transmission of two or more control packets that have collided requires that the transmitting stations recognize the existence of the contention. Also, certain stability control algorithms may be able to usefully employ the more limited information of only recognizing collisions when they are transmitting. The fundamental difference between the two situations is the detection requirement on the modem.

The collision of two control packets should be considered. (The extension to a collision involving three or more packets is clear.) In the most important case, these packets will both be mode 1 type (64-kbit/s uncoded QPSK) transmissions. The packets will arrive at the satellite at essentially the same time. The starting times will differ by only a small fraction of a symbol time unless offsets are specifically introduced. Each packet will begin with a square wave modulated (101010. . .) preamble. In principle, the phasing of the two RF carriers could effectively cancel the two signals (experience destructive interference) during the time occupied by the preamble. The transmitters should be within 100 Hz; during the 1- to 3-ms interval (32- to 96-bit preamble at a 32 kilosymbol/s rate), the maximum relative phase change may be 0.63–1.89 rad or significantly smaller. Therefore, a collision may increase or decrease the power available in the preamble for detection. In particular, the existence of a collision may or may not lead to a preamble detection with the present modem algorithm. Because of the cost associated with implementing a different form of energy detection, this situation will be accepted, and techniques to provide an adequate level of contention detection will be examined.

Since a collision does not reliably initiate a modem preamble detection (or indication of channel activity), collisions are not always being detected at all stations. However, this situation need not reduce the effectiveness of contention detection for the stations whose control packets collided. These stations know when they transmitted and hence when they should hear their own transmission. If all the stations had the same reception capability, no special contention detection would be required because the stations can detect (with high reliability) a control packet collision by listening. This is true because there would be a single transmission mode which all the transmitting stations could hear; if no packet or a damaged packet is received in the expected slot, the occurrence of a collision can be assumed. A station's decision that its control packet did not get through (collided or was otherwise damaged) will be wrong only when the packet is received in error due to local noise. The probability of a collision in the control subframe should normally dominate the probability of a packet error due to local noise.

Some of the stations are disadvantaged (small) and cannot normally receive their control packet transmissions. The control packets under a central or buddy controller strategy would normally be transmitted at the so-called buddy rate, which is matched to the more capable stations for channel efficiency (see Subsection 3.4). For the reliable detection of these collisions, the transmission and detection process must ensure that the disadvantaged station receives certain indicators (information) about its control packets when they do not collide and does not receive these indicators when they do collide. However, in either case (collision or noncollision) the disadvantaged stations would not normally be able to 'hear' the data in the packet.

For simplicity, a network composed of two types of earth stations, small and large, i.e., those that support 64-kbit/s QPSK reception and those that do not, will be considered. (Other situations, which will not be discussed, will tend to follow similar principles.) It is also assumed that the control packets can be "heard" by the large stations but not by the small, i.e., the transmissions are essentially mode 1. The three situations of interest involve the collisions of two or more large stations, a set of large and a set of small stations, and two or more small stations.

Large transmitting stations can determine the occurrence of a collision with high probability by listening in the relevant slot of the received control subframe. Therefore, contention detection at small transmitting stations presents the problem. For simplicity, it is assumed that a single small station was transmitting in the slot where a collision occurred. (For the technique being proposed, it does not make any difference whether a single small station or several small stations are in a collision.)

As explained previously, the control packet transmissions must contain some information that can be heard by the small stations to distinguish collisions from noncollisions. This can be accomplished by introducing two new modes for the transmission of control information in the control subframe. (Although a single new mode would be sufficient, two are considered for clarity. The availability of both modes also provides greater flexibility for the present experiment.) Mode 1B, which would be used by the large stations for transmitting their control packets, would differ from the current mode 1 only in terms of the choice of preamble and SOM; that is, the data (hardware and software generated portions of the packet) would be transmitted at 64 kbit/s. In particular, this mode would use preamble 2 and SOM-3. (Table 3-1 shows the two new modes and the previous modes.) The small stations would send control packets in the control subframe using mode 7. This mode is also a variation of mode 1 and involves a change in preamble and the use of two contiguous SOMs following the preamble; the basic packet data, both hardware and software generated bits, are transmitted at 64 kbit/s. Table 3-1 and Figures 3-4 and 3-5 show the characteristics of this mode.

The use of these modes enables a small station to distinguish between a noncollision situation involving one of its control packets and a collision with one or more large station's control packets. It is assumed that, if a collision occurs, the packet transmissions start essentially simultaneously relative to a symbol time. (The global timing accuracy seems consistent with this assumption, since a symbol time is 31.25 μs and the global timing is accurate within 10 μs). If there is no collision, the small station will lock onto the preamble. If the slot is occupied by a single large station control packet, the small station will also detect SOM-3, the QPSK SOM designed for small station reception. If the slot is occupied by a single small station control packet, the small station will detect SOM-2, the BPSK SOM. Thus, in the noncollision situation, the small station (and the large station) will detect both a preamble and a SOM.

Table 3-1. Mode Chart

Mode	F/S	Coding	CL, PL	Pre	1st SOM	2nd SOM	1st Info Rate[a]	2nd Info Rate[a]	1st Info Rate[a]	2nd Info Rate[a]
1A	F	None	CL = 0	1	1	X	64	X	64Q	X
1B	S	None	CL = 0	2	3	X	64	X	64Q	X
2	S	All	CL = PL	2	2	X	16	X	32B	X
3A	F	All	CL = PL	1	1	X	32	X	64Q	X
3B	S	All	CL = PL	2	3	X	32	X	64Q	X
4	S	Part	CL < PL	2	2	1	16	64	32B	64Q
5A	F	Part	CL < PL	1	1	X	32	64	64Q	64Q
5B	S	Part	CL < PL	2	3	X	32	64	64Q	64Q
6A	S	All	$CL + CL^1 = PL$	2	2	1	16	32	32B	64Q
6B	S	All	$CL + CL^1 = PL$	2	2	3	16	32	32B	64Q
7	S	None	CL = 0	2	2	1[b]	64	X	64Q	X

[a]Rates are in kbit/s, Q = QPSK, B = BPSK.
[b]Second SOM is contiguous to first SOM.

If a collision occurs between a large station control packet and the small station control packet, a preamble either will or will not be detected; however, the small station (and also the large station) will not detect either SOM (neither SOM-3 nor SOM-2). Therefore, this conflict situation is distinguished by either no preamble detection or a preamble detection without a SOM detection. However, since the absence of a preamble detection could also correspond to an empty slot, the best strategy seems to be to notify the SIMP when a control packet is heard without conflict.

PL ≡ Packet Length

CL ≡ Coded Length

CL^1 ≡ 2nd Coded Length for Mode 6, $CL^1 = PL - CL$

F/S ≡ Signal to Designate Preamble (sometimes referred to as L/S)

(F → transmission to large station)

(S → transmission to small station)

Pre ≡ Preamble

(1 ≡ large station preamble - nominal 60 baud times at 32 kilosymbol/s)

(2 ≡ small station preamble - nominal 96 baud times at 32 kilosymbol/s)

SOM Designation

	Nominal Value
(1 ≡ large station QPSK SOM,	32 bits at 64 kbit/s)
(2 ≡ small station BPSK SOM,	32 bits at 32 kbit/s)
(3 ≡ small station QPSK SOM,	64 bits at 64 kbit/s)

Info Rate

Information rate in kbit/s

Chan Rate

Channel rate in kbit/s

Figure 3-4. Parameter Definitions

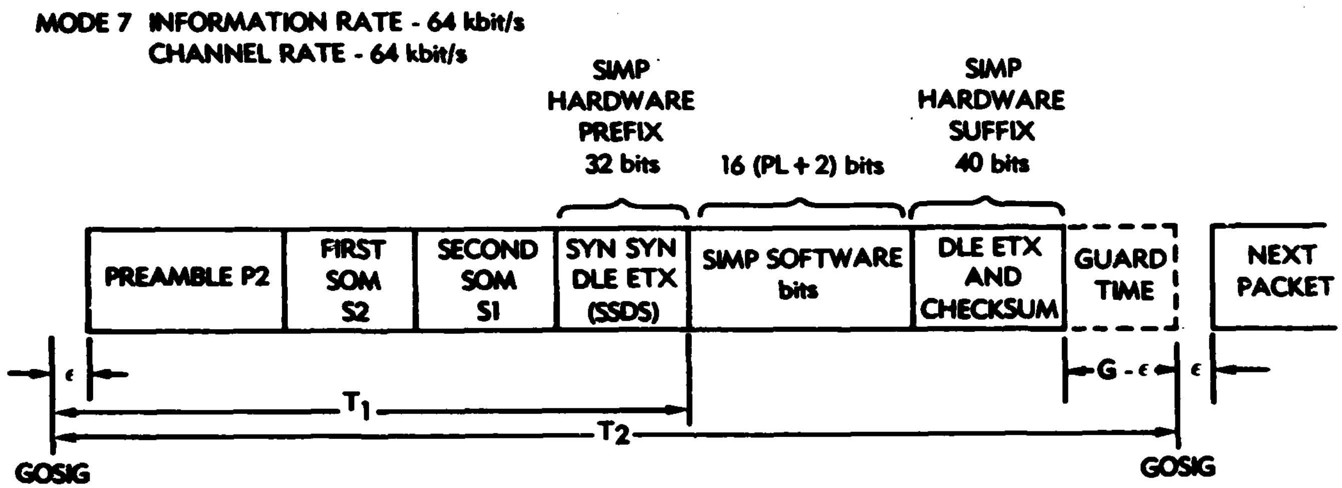

Figure 3-5. Channel Packet Format for Transmission Mode 7

If a small or large station detects both a preamble and SOM but no SDS (SYN DLE STX) [mode 1B or QPSK SOM (mode 7)], it announces that no conflict is present by sending a suitable dummy packet into the SIMP. A small station will send this dummy packet for each slot into the control subframe which is occupied by a single packet. A large station will send this packet only when there is no conflict in the slot and a missed SDS occurs, a low-probability event. Large stations will normally receive the control packet in the absence of a conflict because they will reliably detect an appropriate QPSK SOM (SOM-3 for mode 1B and SOM-1 for mode 7) and an uncoded SDS. If a station has transmitted in a particular slot in the control subframe and upon reception of the control subframe receives neither the transmitted packet nor the dummy packet, it can, with high probability, assume that a conflict has occurred and take appropriate action.

The remaining problem is to distinguish between the absence of conflicts and a conflict between two or more small stations. The necessary information can be supplied by offsetting the allowed transmission times of the small stations. For simplicity, it is assumed that the number of small stations is less than four. If each small station is assigned a unique starting time in a control slot and the starting times differ by an integer number of symbol times, then a collision between two or more small stations transmitting mode 7 packets will result in either no preamble detection or a preamble detection with no SOM detection. Thus, when a collision occurs, none of the transmitting small stations will receive the dummy packet and the occurrence of a collision can be assumed. The proposed mode of operation is summarized in Table 3-2. The dummy packet will resemble an ordinary packet except it will contain only two words in the normal software generated portion of the packet. The format is shown in Figure 3-6.

Table 3-2. State Assumed by SIMP for Control Slot vs Receive Status*

Event in Slot	Stations That Transmitted in Slot		Stations That Did Not Transmit in Slot	
	Large	Small	Large	Small
No Preamble Detection	Conflict	Conflict	Empty	Empty
Preamble but no 1st SOM Detection	Conflict	Conflict	Empty	Empty
Preamble and 1st SOM Detection but no SDS or 2nd SOM event caused by local noise in case of large stations)	No Conflict Dummy CPH Packet Transmitted to SIMP	No Conflict Dummy CPH Packet Transmitted to SIMP	No Conflict Dummy CPH Packet Transmitted to SIMP	No Conflict Dummy CPH Packet Transmitted to SIMP
Packet Received (may contain bit errors)	No Conflict	--	No Conflict	--

*Assumes unique transmission modes, e.g., time offsets.

If offsets are assumed in the start of transmission times of all the stations, which is not costly for a small number of stations, these basic ideas could be accomplished by introducing a single new mode, either mode 1B or mode 7. Two new modes are included herein primarily because it is not clear that reliable SOM-3 detection will be possible at the small station;

the additional mode will also provide increased experimental flexibility, i.e., with a single small station, no transmission offsets are required if two modes are used. In actual operation, only one of the two modes may be used.

Checksum	ETX DLE	Header Word 2	Header Word 1	STX DLE SYN	First Bit Transmitted

The formats of header words 1 and 2 should be the following:

Header Word 1

$$0\;0\;0\;0\;0\;0\;\vdots\;\underbrace{0\;0\;0\;0\;0\;0\;0\;0\;0\;0}_{\text{PL}}$$

Header Word 2

$$\underbrace{0\;0\;0\;0\;0\;0\;0\;0\;0\;0}_{\text{CL}}\;\vdots\;\underbrace{0\;1\;0\;0\;0}_{\text{Mode}}\;\vdots\;\underset{\substack{\uparrow \\ \text{LSB}}}{0}$$

All bits in these two words are zeros except the SOM bit, which is a one. The checksum is 24-bits long and will have a bit pattern determined by the preceding 6 words.

Figure 3-6. Format of the "Control Packet Heard" (CPH) Dummy Packet

The above conflict detection technique imposes certain requirements on the interface operation. Table 3-3 specifies the actions to be taken by the receive interface upon the reception of modes 1B and 7. The SIMP software must also be modified to specify the appropriate transmission modes and to regulate the input of the conflict information for utilization of the conflict detection capability to minimize the time required to transmit control information.

3.4 SATELLITE NETWORKING

The multiple access and broadcast capabilities of a satellite can be used to form a network, i.e., provide connectivity among a set of cooperating earth stations. The required connectivity can be realized by various methods including fixed capacity allocations on links connecting the earth stations and dynamic sharing to provide the capacity and connectivity required by the actual traffic. In systems which utilize dynamically changing connectivity or capacity, the response time necessary to change the connectivity or capacity allocations strongly influences the properties of the network realization. The economics and properties of various network realizations are also influenced by whether the network occupies an entire transponder or shares the transponder with other users.

Three basic types of satellite network realizations, which illustrate some of the options and tradeoffs, are dedicated channel, dedicated up-link/broadcast down-link, and shared up-link/shared down-link. In a dedicated channel network realization, fixed data rate point-to-point links connect each pair of stations in the network (assuming a fully connected network).

Table 3-3. Receive Interface Actions for Modes 1B and 7 Reception

Event	Possible Cause(s)	Interface Action
No preamble detection	1. Empty slot 2. Missed preamble 3. Conflict	None
Preamble detection but not first SOM detection	1. Missed SOM 2. Conflict 3. Preamble false alarm	1. Send EOP signal to modem τ seconds after preamble detection; τ set by T&M module 2. Declare error condition; output appropriate trap data, could put out a 'conflict' dummy packet
Preamble detection, first SOM detection, no SDS detection (mode 1B) or no QSOM detection (mode 7)	1. Missed SDS (mode 1B) 2. Preamble and SOM false alarm 3. Missed QSOM (second SOM in mode 7) 4. No conflict but transmission mode above station capability 5. Conflict between large stations if transmissicn not offset	1. Send EOP signal to modem 2. Send dummy packet to SIMP (If large station transmissions are not offset, the large stations which transmitted in this slot should assume a conflict.) 3. Small - assume no conflict 4. Large - assume no conflict, i.e., local noise error assume offset
Preamble detection, SOM detection, SDS detection (mode 1B) or QSOM detection (mode 7)	1. No conflict and station can 'hear' packet	1. Mode 1B - send received packet to SIMP 2. Mode 7 - send received packet to SIMP if SDS is detected or declare error and send EOP signal to modem if no SDS is detected

For an N node network, N(N - 1) full-duplex links are required with each link sized to handle the average point-to-point traffic between the two connected stations while satisfying a suitable utilization or traffic delay constraint.

In a dedicated up-link/broadcast down-link realization, each station in the network has a dedicated fixed data rate up-link which can be received by all the other stations in the network. This realization requires N broadcast simplex links for an N station network with each link sized (subject to a suitable utilization or traffic delay constraint) to handle the average traffic originating at the transmitting station. Each station multiplexes its own originating traffic on its simplex link for transmittal to the other stations in the network. In the shared up-link/shared down-link realization, a single broadcast communication channel is dynamically shared among all the stations using demand assignment techniques to provide the required connectivity and link capacity.

The choice between the use of TDMA and FDMA depends strongly on the network traffic matrix and the method of realization. A shared up-link/shared down-link network, for example, is highly constrained by FDMA. The choice is also influenced by any other traffic carried by these stations and by any transponder sharing. As network size increases, TDMA becomes more cost effective than FDMA in providing one hop connectivity because less equipment (e.g., modems) is required at the ground stations; however, each piece of equipment is more expensive. TDMA allows more flexibility than FDMA in establishing and adjusting individual link data rates to the average traffic demands.

With TDMA techniques, a direct tie is established between the flexibility available in the TDMA frame and the possible network realizations. If the bursts within the frame are fixed or slowly vary in both length and assignment to specific earth stations and the bits within each burst are also essentially permanently assigned to specific earth stations (corresponding to fixed capacity simplex links from the station transmitting the burst to the various destinations), only dedicated channel networks can be formed. If burst lengths within the frame are fixed and allocated to specific stations but the bit assignments to various destinations within the individual bursts can be dynamically varied to match the real-time traffic demands of the transmitting station, then dedicated up-link/ broadcast down-link types of network realizations can be supported. If the entire frame can be flexibly demand assigned, i.e., fixed bursts are not allocated to specific stations, then the shared up-link/shared down-link realization is possible. Networks realized via up-link/shared down-link TDMA techniques are economically attractive because of their efficient use of satellite capacity.

FDMA network realization techniques readily accommodate networks which must share transponders. The shared up-link broadcast down-link network realization is a demand assigned TDMA system in an FDMA derived channel. (The PODA algorithm could, with minor modification, utilize several FDMA derived channels.) Networks can also be realized in a subchannel of a higher data rate TDMA system, e.g., a portion of each frame could be allocated to the network. If the transponder is shared via FDMA techniques, an interference penalty and a loss of efficiency may be experienced because of the need for linear transponder operation. If the transponder is shared via TDMA techniques, there is some loss of system flexibility and a potential cost impact because operation at the higher burst rate is necessary.

4. EXPERIMENTAL PROGRAM

4.1 INTRODUCTION

The COMSAT experimental program was designed to complement and enhance the programs of the other participants in the SATNET community. The principle goal in the COMSAT program was to measure the system performance under a wide variety of conditions likely to occur in an operational network and to determine to what degree the system meets its design objectives. There are two issues inherent in those activities. The first is the measurement of system throughput, delay and other measures of customer service with respect to resource utilization within the satellite channel and supporting ground equipment. The second is the validation of stability-control and fault-recovery mechanisms built into the system to insure robustness and fairness.

In contrast to these issues judged of primary importance in the COMSAT experimental program, the programs of the other participants, namely BBN, NDRE, UCL and UCLA, emphasized other issues. BBN, the principal implementor of the ground-support software, was primarily concerned with implementation, debugging and validation. Many of the tools developed by them, namely EXPAK and MONITOR, were useful to the other participants as well. The COMSAT experimental program complemented the BBN program by exploring issues of fairness and robustness and in detecting problems in implementation. NDRE and UCL were active in speech experiments, end-user applications and gateway performance validation. The COMSAT program did not address these areas directly, since the measurement data activities involved only the SIMP's and the channel itself. UCLA was concerned with a wide spectrum of channel measurements and several protocol variations, including

CPODA, and with the detailed evaluation of CPODA throughput and delay throughout the regime of stable system operation. The COMSAT program complemented this program by exploring configurations including three stations operating in potentially unstable regimes and in cases where stability control and fairness principles could be tested.

The COMSAT program also included performance measurements involving the Unattended Earth Terminal (UET) at Clarksburg. Operation of the UET with the CPODA protocol was possible only near the end of the reporting period and then only in loopback mode to itself. Experiments, reported later in this section, confirmed that the software and hardware was compatible and that the interface operated correctly in the BPSK-coded mode with the LINKABIT modem.

In the following subsections, the structure of the network is described with reference to the experimental data paths used for command, data collection and performance monitoring of the network. The structure of the experiment-support software, including portions implemented by BBN, UCLA and COMSAT, is described, with special emphasis on the COMSAT contribution. This software, which operates in host machines on the ARPANET, interacts with the SIMP and the gateway software to provide a real-time command/control system, which has numerous interesting features. The objectives and structure of the experiments are presented and the results are reported in the context of a simple channel model that has been useful in establishing experimental parameters and predicting general system behavior.

4.2 SATNET CONFIGURATION

During the previous two years the SATNET configuration has been continuously developed and has changed significantly.

With increasing experience and sophistication of the experimenters, the experiments have become more intricate and comprehensive. This subsection describes the network topology, as well as the control, data collection, and monitoring paths available for experimentation.

The detailed organization of the SATNET space-segment and earth-station components has been described and will not be repeated. For reference, however, the structure as viewed by the experimenters will be outlined. The current configuration, shown in Figure 4-1, consists of the Atlantic Primary satellite in position off the west coast of Africa, together with the Etam, Goonhilly, Tanum, and Clarksburg earth stations. Etam, Goonhilly, and Tanum (the large stations) can receive and transmit a nominal 64-kbit/s data rate, while Clarksburg (the small station) can receive at a nominal 16-kbit/s data rate and transmit at either data rate. When the PSP Terminals are installed at the large stations, then all the stations can transmit at either data rate.

The satellite/earth station control computers, the SIMPs, are colocated with the radio equipment at the station and are connected through gateway computer to the ARPANET. All end-to-end user traffic, as well as experiment control and collection traffic flows through the gateways with certain exceptions. At Clarksburg the RCC line provides connectivity to the ARPANET for experiment control and collection traffic only. This line carriers no user traffic. At Etam, the SDAC line provides connectivity to the ARPANET if the BBN gateway fails. The "back door" lines at Goonhilly and Tanum are not normally available for this purpose.

In the latest SIMP-II control program, the experiment control and data collection functions are normally performed using the data path through the gateway (except at Clarksburg), which is also assumed in this report. In some cases, the gateways

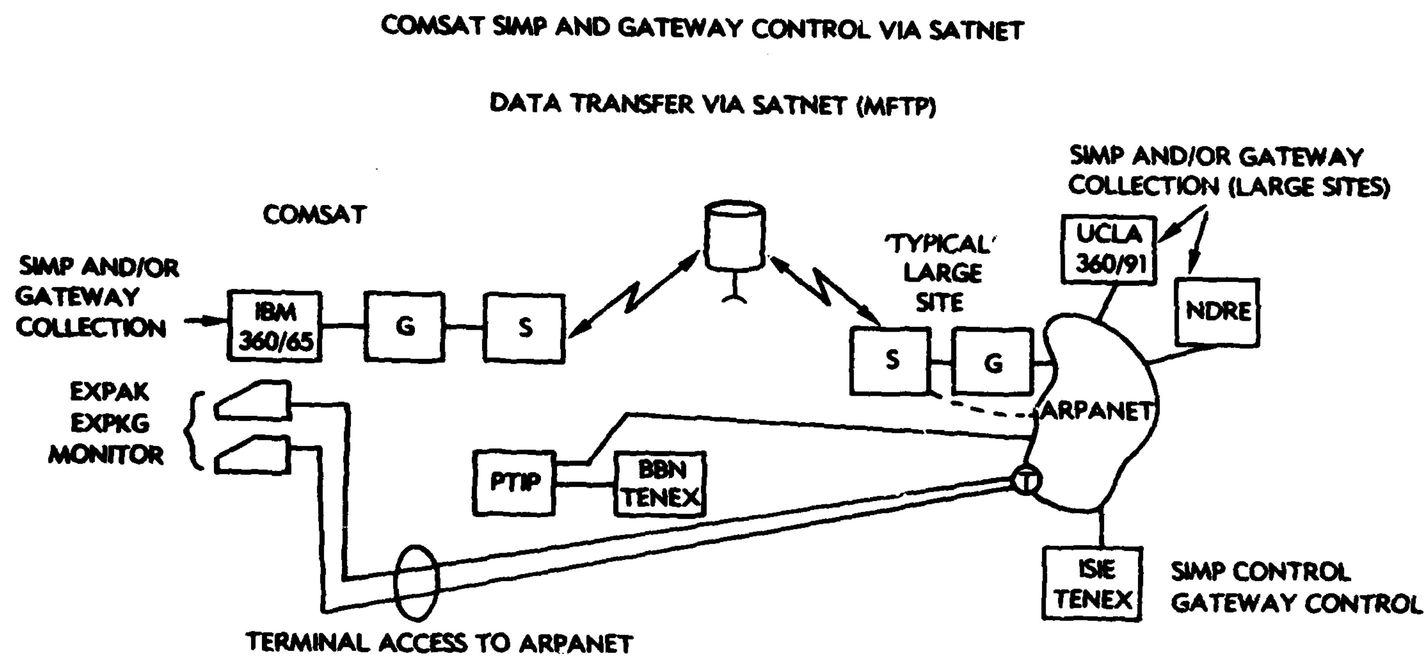

Figure 4-1. Experimental Configuration

are a considerable distance from their SIMPs, in one case not even in the same country. In this report, the gateways are identified as NDRE, which is connected to the Tanum SIMP; UCL, which is connected to Goonhilly; COMSAT, which is connected to Clarksburg; and BBN, which is connected to Etam.

The COMSAT gateway, which is not directly connected to the ARPANET, was originally intended as a data path between the ARPANET and the IBM System 360/65 computer located at COMSAT Laboratories in Clarksburg; however, this function has not been fully developed. Instead, it is planned to connect the gateway to a small minicomputer intended as a base for application-oriented experiments involving packet speech, facsimile and other specialized data. Presently, the 4.8-kbit/s data connection from the Clarksburg SIMP to the RCC in-house computer network at BBN enables the ARPANET-resident control programs to access the Clarksburg SIMP. A special feature of the code, which operates much like a miniature gateway, also provides access to the COMSAT gateway via the 4.8-kbit/s line.

Notwithstanding the connectivity via the gateways and the ARPANET, the control, collection, and monitoring traffic can be sent from one SIMP to another by the space setment. For example, if a connecting gateway or line is down, traffic to and from one SIMP can be routed via the space segment, another SIMP, its attached gateway and the ARPANET. Therefore, the possible configurations for experiment control, data collection and monitoring are numerous and an interesting feature for study in the experimental program.

4.3 EXPERIMENT CONTROL, COLLECTION, AND MONITORING FUNCTIONS

In the last Interim Report [4], the structure and operation of the SATNET experiment control, collection and monitoring software was described. In the latest reporting interval, significant new features were added to the COMSAT software to support non-steady-state experiments and to greatly reduce the effort required to perform them. In this subsection, the structure and operation of the complete package, including the new features will be described.

Figure 4-2 shows the configuration of the experiment control, collection, and monitoring system. In the design of the present experiment control software both the control and data-collection functions must occur on the same ARPANET host. Although two hosts, BBNE in Boston and ISIE in Los Angeles, are equipped for these functions, all experiments conducted to-date have used ISIE. The software package for experiment control, EXPAK, generates control messages, termed parameter-change (PC) messages, which are transmitted over the ARPANET paths to a selected SATNET gateway and then to the SIMP. The gateway selected is usually connected directly to the SIMP for efficiency and reliability; however, the PC message could be relayed via another gateway and the space segment if required by equipment or circuit outages. Data collected at the various SIMPs are returned to EXPAK by the same or different paths than those used by the PC messages. These data include configuration and status information pertaining to each SIMP, as well as accumulated statistics on equipment operation and data traffic, called cumulative statistics (cumstats).

The principal measurement tool in all experiments is a set of synthetic traffic generators, which can simulate various traffic rates, distributions, and network configurations. Each of the eight traffic generators in each SIMP can be specified to direct time-stamped messages to the same or another SIMP, which measures the transmission delays and constructs statistics such as means, maxima, minima, and histograms. These data are incorporated into the cumstat messages which are returned to the experiment control site.

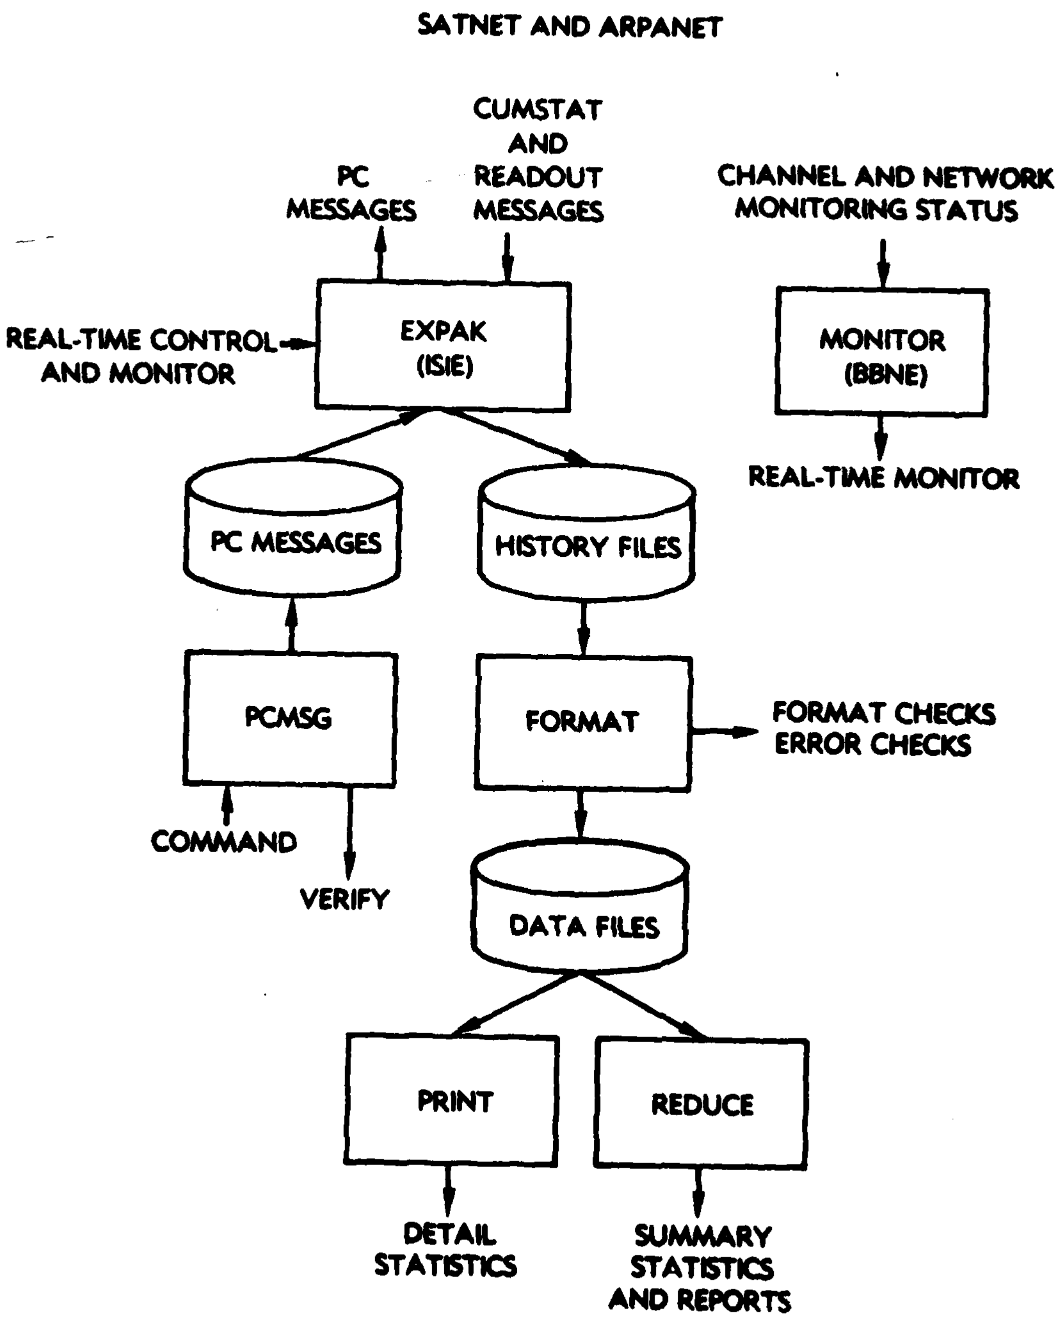

Figure 4-2. Experiment Control, Collection, and Monitoring

The traffic generators can be specified to generate traffic of several rates and types, including datagram and stream with and without acknowledgments and belonging to any of several priority-delay classes. Each of the eight generators can be separately enabled or disabled at specified times according to a schedule, stored in the SIMP, which can be specified by the experimenter.

A typical experiment involves the specification of the traffic-generator parameters for each SIMP, together with a specification of the schedule of times each generator is to be enabled or disabled. These data are sent as PC messages over the ARPANET and gateway paths to each SIMP as required. As part of this process, the destination and route (gateway) to be used by the cumstats is specified. The experiment itself is initiated by a special PC message which causes the schedule to be activiated in each SIMP at the same global time. As the binary-coded cumstats are received at the experiment collection host, they are converted to symbolic-text form suitable for printing and saved in the experiment data base. Following the experiment, the set of cumstats collected are processed and various reports are produced.

4.3.1 The Real-Time Network Monitoring Program (MONITOR)

In addition to the control and collection functions initiated by the experimenter, the SIMPs and the gateways automatically send periodic reports (as shown in Figure 4-3) to a designated ARPANET site, currently BBNE in Boston. The reports of the SIMPs are derived from channel monitoring data using the hello subframe, which occurs at about 1.3-s intervals. The hello subframe is divided into four subframes during each of which a single SIMP transmits a monitoring packet. All the SIMPs, including

```
SatNet Monitor, Internet Version: 350
Last Recorder(361) Restart: 12-Dec-78 07:17:55
Last GWMON(52) Restart: 10-Dec-78 14:14:09
Current mode is: Event monitoring
Type ? for help.

All..

E 1012 001777 (3.0:6) 0  0,0,0,0  0,0,0,0  0,0,0,0  16,16,16,16
E 1012 001777 E=0/0/192 S=0*0/0/256 C=4/4/4/.90 DS=0/0/4 DR=0/0/0
E 1012 001777 (3.0:6)    CPODA  R=64  BCK=3.21
T 1012 001777 (3.0:6)    CPODA  R=64  BCK=3.68
E 1012 001777 Host 1 Up PS=0/3 PR=0/4 HI=64/62 MS=0/1 MR=0/2
E 1012 001777 Host 2 Up PS=0/0 PR=0/4 HI=66/66 MS=0/0 MR=0/2
G 1012 001777 Host 1 Up PS=0/3 PR=0/4 HI=67/66 MS=0/1 MR=0/2
T 1012 001777 Host 1 Up PS=0/3 PR=0/4 HI=67/66 MS=0/1 MR=0/2
G 1012 001777 (3.0:6) 0  0,0,0,0  0,0,0,0  0,0,0,0  16,16,16,16
G 1012 001777 E=0/0/192 S=0*0/0/256 C=6/2/2/.90 DS=0/0/2 DR=0/0/0
T 1012 001777 (3.0:6) 0  0,0,0,0  0,0,0,0  0,0,0,0  16,16,16,16
T 1012 001777 E=0/0/192 S=0*0/0/256 C=6/2/2/.90 DS=0/0/2 DR=0/0/0
G 1012 001777 (3.0:6)    CPODA  R=64  BCK=3.51
- 1012 Gateways, B=0/2, U=0/4, N=0/3, C=Dn
- 1013 Gateways, B=0/4, U=0/4, N=0/4, C=Dn
G 1014 021777 (3.0:6) 0  0,0,0,0  0,1,0,0  0,0,0,0  16,16,16,16
G 1014 021777 E=1/0/191 S=0*0/0/256 C=6/1/1/.90 DS=0/0/1 DR=0/0/0
T 1014 021777 (3.0:6) 0  0,0,0,0  0,0,0,0  0,0,0,0  16,16,16,16
T 1014 021777 E=0/0/192 S=0*0/0/256 C=6/1/1/.90 DS=0/0/1 DR=0/0/0
G 1014 021777 (3.0:6)    CPODA  R=64  BCK=3.51
T 1014 021777 (3.0:6)    CPODA  R=64  BCK=3.68
G 1014 021777 Host 1 Up PS=0/3 PR=0/2 HI=67/65 MS=0/1 MR=0/1
T 1014 021777 Host 1 Up PS=0/3 PR=0/2 HI=66/65 MS=0/1 MR=0/1
- 1014 Gateways, B=0/2, U=0/4, N=0/4, C=Dn

 1014 12-Dec-78  ---  All Is OK.
```

Figure 4-3. Example Of MONITOR Output

the transmitter, monitor these packets and record the number of packets successfully received. The data derived from these packets are periodically sent via either the attached gateway and the ARPANET to BBNE or indirectly via the space segment and another SIMP to the same destination (lines beginning with "E", "G", or "T", and containing only numbers). Path selection depends upon numerous considerations, including the state of the line connecting the SIMP and its gateway, certain fixed tables assembled in the SIMP code, and explicit polling messages transmitted to the SIMPs from the monitoring program at BBNE. Other data reported by the SIMP include message traffic to and from the gateways and the status of the connected communications line (other lines beginning with "E", "G", or "T" in Figure 4-3).

The gateways also periodically report status to BBNE (lines beginning with "-" in Figure 4-3) using only the ARPANET paths. From these reports and from those of the SIMPs, the BBNE monitoring program (MONITOR) determines the components and circuits which are operable in the aggregate experimental network and makes available periodic state information which can be accessed by experimenters and network control personnel. During normal experiment operations, the experimenter has access to a terminal which displays this information in real time. Presently, this requires a terminal in addition to the one used for experiment control.

4.3.2 The Experiment Control Program (EXPAK)

Experiments involving the satellite channel and the SIMPs are normally controlled using a program called EXPAK, which was constructed by BBN and described in PSPWN 83. The program is designed to run on designated ARPANET hosts, currently ISIE and BBNE, and provides interactive control of several functions, including the following:

a. Authenticating user identity, enabling certain privileged system features, and setting up data collection files.

b. Translating the user-supplied PC messages into internal form and sending them as directed to the selected SIMPs using the selected ARPANET/SATNET paths.

c. Storing the incoming cumstat messages in a specified file termed the history file.

d. Accumulating the appropriate cumstat data into a specified file termed the summary file (also called the cumstat array).

e. Presenting the selected cumstat messages and portions of these messages in real time on the operator's terminal. (This is useful for debugging and monitoring certain data as the experiment proceeds.)

f. Establishing the ARPANET/SATNET data paths to be used for PC and cumstat messages automatically or as specified by the experimenter.

This software package has been revised as the SIMP control program and the SATNET monitoring and control procedures have matured. In the most useful mode of operation, PC messages are prepared in advance for an experiment using the appropriate program. The messages for a single experiment involving one or more SIMPs are stored in a file with routing information indicating the SIMP and access path for each message. Figure 4-4 is an example of such a message. A number of files are necessary, one for each data point in an experiment and a set for each experiment. Additional files are required for various functions including SIMP initialization, experiment activation, and panic stop. The PC message is represented by pairs of octal numbers, the first identifying the parameter and the second the value. Each message segment is preceeded by a header indicating the SIMP(s) and route (gateway). These messages can then be sent by EXPAK with a single command.

```
: E,* G,* T,*
205 000074   206 000074   207 005001   210 000064   211 000000
212 177600   213 000000   217 000002   361 000001   363 000224
364 000224   365 000224   366 000224   373 100200   374 100400
375 101000   376 102000   377 000100   400 000100   401 000100
402 000100   413 000224   414 000224   415 000224   416 000224
423 100200   424 100400   425 101000   426 102000   427 000100
430 000100   431 000100   432 000100   443 000074   444 000000
445 000074   446 101000   447 000074   450 042000   451 000074
452 146000   453 000074   454 021000   455 000074   456 125000
457 000074   460 063000   461 000074   462 167000   463 000074
464 010400   465 000074   466 114400   467 000074   470 052400
471 000074   472 156400   473 000074   474 031400   475 000074
476 135400   477 000074   500 073400   501 000074   502 177400
503 000000   504 000000   505 000000   506 000000   507 000000
510 000000   511 000000   512 000000   513 000000   514 000005
515 000006   516 000007   517 000010   520 000011   521 000012
522 000013   523 000014   524 000015   525 000016   526 000017
527 000020   530 000021   531 000022   532 000023   533 000024
534 177552   535 177634   536 177716   537 000000   540 000062
541 000144   542 000226   543 000310   544 000372   545 000454
546 000000   547 000000   550 000000   551 000000   552 000000
553 000000   554 000000   555 000000   556 000000   557 000000
560 000000   561 000000   562 000000   563 000000   564 000000
565 000000
:
: E,*
403 000011   404 000011   405 000011   406 000011   407 000005
410 000005   411 000005   412 000005   433 000015   434 000015
435 000015   436 000015   437 000005   440 000005   441 000005
442 000005
:
: G,*
403 000015   404 000015   405 000015   406 000015   407 000011
410 000011   411 000011   412 000011   433 000005   434 000005
435 000005   436 000005   437 000011   440 000011   441 000011
442 000011
:
: T,*
403 000005   404 000005   405 000005   406 000005   407 000015
410 000015   411 000015   412 000015   433 000011   434 000011
435 000011   436 000011   437 000015   440 000015   441 000015
442 000015
:
```

Figure 4-4. Example Of PC Message (Use File Format)

The cumstats collected during the experiment are recorded in both the history and the summary files. The history file contains a complete packet-by-packet description of all messages sent to and from the SIMPs, including PC messages, cumstat messages, and acknowledgments. In addition, the time-stamped data portions of all cumstat messages are recorded packet-by-packet. The summary file contains the aggregate cumstat data collected during the experiment. The individual entries consist of either the arithmetic sum or the maximum of corresponding entries in the individual cumstat messages as they arrive at the data collection site. With the use of the history file, the events during an experiment can be reconstructed in considerable detail, including the production of the summary file. However, the large size of the history file creates problems in file management and access.

The experiments reported previously have involved only steady-state measurements. Therefore, only the data of the summary file have been necessary and the only data used from the history file have been special cumstat messages, called readout messages, returned by the SIMPs participating in the experiment. The readout messages contain the current values of all storage areas that can be altered by PC messages and verify that the various parameters have been set properly. The experiments run during the last reporting interval included many that were designed to observe transient phenomena and other non-steady-state measurements. For this reason, the experiment data-reduction procedures now make use only of the history file and the summary file is no longer used. In order to reduce the bulk of these data, new reformatting and data-compaction procedures were developed. These will be described in the next section.

While the experiment is in progress, selected portions of the cumstat messages can be displayed upon arrival in real time at the controlling terminal. This feature has been useful in determining if the experiment is proceeding properly or if an external disruption such as equipment, line, or channel failure has occurred. In general, experiment plans have called for the return of a complete set of cumstats from each SIMP at intervals of between fifteen and thirty seconds and for the continuation of the experiment for several minutes. These values have been proven appropriate for the experiments conducted to-date.

4.3.3 Data Reduction And Editing Programs

The volume of the data collected during experiments has required construction of a sizable library of data reduction and editing programs. Although at least one program usable with EXPAK-collected data was known to exist (REDUCE, written at UCLA), this program was designed primarily to yield throughput/delay measurements and was not suited to the much wider range of measurements required in the COMSAT programs. In addition to throughput/delay measurements, the COMSAT programs were designed to yield performance data of the CPODA algorithm itself with respect to resource allocation and fairness principles, and to probe for potential problems in stability and robustness.

The COMSAT programs were all written in BASIC for the TOPS-20 host in the ARPANET. This language was chosen so that the data reduction and edition procedures could be easily constructed and modified and has, on hindsight, been a good choice. Although normally run on a TOPS-20 host, some of these programs have also been run on an off-line LSI-11 host. When operated in this mode, the experiment data files (history or summary) are transferred via an ARPANET TIP line to floppy-disk storage on the LSI-11 prior to processing.

As will be described in detail below, the manner in which data are collected during an experiment has changes since the last report. Where previously a separate and distinct run was performed for every data point, the new technique requires only a single run for all data points in a particular experiment. Thus, a run containing up to sixteen different sets of traffic-generator parameters can be performed without operator intervention. Besides providing the mechanism necessary for transient analysis, this technique has dramatically reduced the time and effort necessary to conduct the experiment. The additional software which implements these capabilities includes features for plotting throughput/delay curves and other similar data.

Of the programs incorporated into the COMSAT library, the following are representative:

- SEPAR - Extract selected cumstat messages from the history file for later processing.
- FORMAT - Reformat cumstat messages, check for errors and save in compacted form.
- PRINT - Print cumstat messages in detailed form (primarily used as a debugging tool).
- REDUCE - Print reduced data from cumstat messages in the form of histograms and tables.
- EXPLOT - Plot selected data points reduced from a collection of cumstat messages.
- PCMSG - Construct and edit PC messages to be used by EXPAK.

The structure and function of each of the programs described in abstract above will be described in this subsection.

SEPAR extracts selected cumstat messages from the history file. Normally only a single history file is recorded for a set of runs comprising one or more experiments conducted during a single session. This program provides a mechanism to sort and merge the cumstat data into separate files for further processing. Normally, these files, as well as the history file, are purged from the data base once the files have been processed by FORMAT (see below).

FORMAT extracts cumstat messages from the history file or the files prepared by SEPAR to produce a formatted file, in compact form, suitable for archival storage and for further processing by other programs. This program also checks for consistency in the readout messages returned by the SIMPs and summarizes the various types of errors that have occurred, including those due to channel errors, buffer overruns, and equipment faults. On the basis of this information, a data point can be discarded if necessary before more extensive processing is performed.

A specimen printout produced by FORMAT, which is shown in Figure 4-5, displays a set of PC messages sent to the SIMPs along with their transmission times. Occasionally, a single, long PC message is fragmented by EXPAK into several short packets before transmission; in these cases, separate acknowledgment packets are returned by the SIMP. Discrepancies between the readout messages returned by each SIMP are also recorded. Although most of the data of these messages is identical for each SIMP, some (e.g., that pertaining to the traffic generators) is not. These intentional discrepancies as well as the inadvertent ones, which result in an invalid experimental configuration, are shown. The printout also summarizes the errors reported during the experiment. Many of these, for example, down-link transmission errors, are a normal feature of contention systems like SATNET. Others, such as up-link errors and buffer overruns, are indicative of resource management problems or inadequacies in the SIMP hardware and software design.

PRINT reads the compact file produced by FORMAT and prints edited summaries of both readout and cumstat data. The operator can specify that selected SIMPs and subsets of the

```
Enter filename<cr> ?SUM35

EXPAK History File Started: Fri 21-Apr-78 11:31
By user: COMSAT-NETWIZ
11:32:05 Using data from file: EXP3.5A.2
11:32:07 ACK from Etam, PC msg: 0, Pkt: 1, Simp Ver: 2.5:2
11:32:10 ACK from Goonhilly, PC msg: 0, Pkt: 1, Simp Ver: 2.5:2
11:32:11 ACK from Tanum, PC msg: 0, Pkt: 1, Simp Ver: 2.5:2
11:32:13 ACK from Etam, PC msg: 0, Pkt: 2, Simp Ver: 2.5:2
11:32:13 ACK from Tanum, PC msg: 0, Pkt: 1, Simp Ver: 2.5:2
11:32:19 ACK from Goonhilly, PC msg: 0, Pkt: 3, Simp Ver: 2.5:2
11:32:24 ACK from Tanum, PC msg: 0, Pkt: 4, Simp Ver: 2.5:2
11:32:27 Using data from file: START..2
11:32:29 ACK from Etam, PC msg: 0, Pkt: 1, Simp Ver: 2.5:2
11:32:40 Readout from Etam, PC msg: 0, Pkt: 1, Simp Ver: 2.5:2
11:32:41 Readout from Tanum, PC msg: 0, Pkt: 4, Simp Ver: 2.5:2
11:32:43 Readout from Goonhilly, PC msg: 0, Pkt: 3, Simp Ver: 2.5:2
11:32:43 History file closed at: Fri 21-Apr-78 11:36

Readout discrepancies (all numbers are decimal)
Parm     Etam     Goony    Tanum    Clark
123      34048    34560    34560    0
124      34048    34560    34560    0
131      9        13       5        0
132      13       5        9        0
135      5        9        13       0
136      5        9        13       0

Error summary            Etam     Goony    Tanum    Clark
Uplink errors            0        0        0        0
Downlink errors          45       36       45       0
Frame and time synch     0        0        0        0
Res synch transitions    0        0        2        0
Info pkt retrans         9        18       11       0
Msg dropped              0        0        0        0
Buffer overrun           0        0        0        0
VDH error pkts           0        0        0        0
```

Figure 4-5. Example Of FORMAT Output

readout or cumstat data be included and that the output be meaningfully labeled. This program is most useful as a diagnostic tool to interactively probe for inconsistencies in the experiment design and conduct after anomalies in the readout or error reports have been discovered by the FORMAT program.

Figure 4-6 shows a part of a specimen printout of PRINT, which was produced from the portion of the summary file collected from the Etam SIMP during an experiment. The data are presented in raw form, and no processing other than formatting is performed.

REDUCE reads the compact file produced by FORMAT, computes specified performance data, and prints edited summaries in the form of tables and historgrams. As in the other programs, the operator can select the form and content of the output from one or more of the following categories:

a. Frame data, consisting of the mean lengths of the various frames and subframes with calculated utilizations.

b. Packet data, consisting of transmitted and received packet tallies for each station classified by subframe and type.

c. Message data, consisting of transmitted and received message tallies by packet type. (These include message arrival, reservation, transmission, reception, delivery and acknowledgment counts for each station.)

d. Header data, consisting of traffic and throughput of the channels formed by the available headers of each packet type (contention and piggyback).

e. Systems performance data, consisting of mean delays and throughputs for the various message source-sink combinations selected for the particular experiment. The displayed information includes means, upper and lower bounds and histograms, together with packet counts by SIMP, delay class and priority.

```
SUM20 E Cumulative Statistics

1. Channel I/O

Channel input: 474  Hardware error: 0  Checksum error: 0
 Receive-time error: 0  No header buffer: 0

Channel output: 157  Hardware error: 0  Transmit error: 0
 Send-time error: 0  Timeout error: 0  No output buffer: 0

Late SAT reset: 0  Max late (HT): 0  Total late (HT): 0

2. Received Packet Tallies

Accepted block data packets for me: 3

Downlink noise states     Accepted        Rejected
 Hello pkt                450             0
 Control pkt              16              0
 Block data header        8               0
 Block data pkt           8               0
 RTT pkt                  0               0

 Uplink noise pkts: 0  Dropped on downlink: 0

3. Buffers and Queues

Send freelist underflow: 0  Freelist underflow: 0  No res node: 0

Queue stats       Samples           Max             Total
 Freelist         1745                              5        162441
 QI               1745              1               3
 QH               1745              1               0
 QC               1749              1               3
```

Figure 4-6. Example Of PRINT Printout

An example printout of REDUCE is shown in Figure 4-7 and 4-8. Figure 4-7 illustrates the reduced data on frame and packet header utilization, packet tallies, mean queue lengths, and channel loading. For this particular experiment, the information subframe was almost completely full [relative rate in the information subframe (Section 1, line 3) is 0.818], and the control subframe was heavily loaded [traffic and throughput (Section 7, lines 2 and 3) are 0.158 and 0.123, respectively.] Occasionally, because of the particular algorithm used to transmit cumstats, the totals in the printouts produced by REDUCE may not be precise. This is evidenced by the small negative numbers in the specimen printout and by the small discrepancies in the packet tallies. In addition, well-known anomalies in the operation of the BASIC interpreter cause certain very small floating-point numbers to appear many orders of magnitude too large.

Figure 4-8 illustrates the format of the message delay and histogram data produced by REDUCE, which includes several statistics on delay means, maxima, and minima, as well as histograms of delay distributions. As explained previously, these histograms are valuable in the identification of the instabilities and anomalies in the channel-scheduling algorithm operation.

EXPLOT is a plotting package designed to provide a "quick look" feature for the various data collected. It scans a collection of cumstat messages and extracts successive values of a selected variable. It can also perform data reduction operations as in REDUCE to produce these values. It then scales and plots these values by cumstat number and labels the axes. An example of the plot produced can be found in the next subsection.

PCMSG is a special-purpose editing program used to produce PC messages for submission to EXPAK. It provides a number of editing and formatting features designed to simplify message construction and manipulation, especially for users who do not

SUM27 Combined Statistics

1. Frame Utilization	Hello frame	PODA frame	Info subfrm	Ctl subfrm	
Total frames	150	600	600	600	
Total length (VS)	19200	17400	15721	1679	
Relative rate	1	0.906	0.818	8.7000	
Max length (VS/frame)	128	29	27	29	
Mean length (VS/frame)	128	29	26.201	2.798	

2. Uplink Channel	Etam	Goony	Tanum	Clark	Total
Info pkt transmitted	1836	1973	2003	0	5812
Info pkt retrans	1	3	3	0	7
Ctl pkt successes	92	61	56	0	209
Ctl pkt failures	14	26	17	0	57
Hello pkt transmitted	150	150	150	0	450
Discarded	0	-1	-2	0	-3

3. Downlink Channel	Etam	Goony	Tanum	Clark	Total
Info pkt received	5813	5813	5812	0	17438
Ctl pkt received	208	208	208	0	624
Hello pkt received	450	450	450	0	1350
Discarded	2	2	3	0	7

4. Channel Protocol	Etam	Goony	Tanum	Clark	Total
Msg generated	1832	1965	2012	0	5809
Msg accepted	1832	1965	2012	0	5809
Res attempts	1863	2039	2059	0	5961
Res successes	1836	1991	2026	0	5853
Msg sent to channel	1837	1976	2006	0	5819
Msg rcvd from channel	2013	1930	1869	0	5812
Msg delivered to sink	2030	1978	1896	0	5904
Ack attempts	2018	1974	1901	0	5893
Ack successes	2008	1956	1884	0	5848
Msg dropped	0	0	0	0	0

5. Mean Queue lengths	Etam	Goony	Tanum	Clark	Total
Freelist	33.383	35.464	35.887	0	104.73
QI	18.455	19.073	18.7	0	56.229
QH	0.032	0.099	0.015	0	0.146
QC	15.351	15.38	15.35	0	46.083

6. Packet Utilization	Etam	Goony	Tanum	Clark	Mean
Mean res ctl subfrm	0.103	7.8000	0.061	0	0.08
Mean ack ctl subfrm	4.1000	0.028	0.027	0	0.032
Mean load ctl pkt	0.883	0.866	0.835	0	0.864
Mean load piggy pkt	0.574	0.574	0.569	0	0.572

7. Common Channel	Info	Ctl	Ctl hdr	Pig hdr
Capacity	15721	1679	266	5819
Traffic	0.37	0.158	0.864	0.572
Throughput	0.369	0.123	0.681	0.572

Figure 4-7. Example Of REDUCE Output

8. Sink Statistics

```
Class:  5  Total msg: 1853  Desired delay:  256  Mean delay:  304.2
Mean early:  14.8  Max early:  69  Mean late:  55.4  Max late: 621
      0 -9999 !
      0  -150 !
      5  -100 !
    185   -50 !*****
   1090     0 !***********************************
    361    50 !***********
     99   100 !***
     48   150 !*
     44   200 !*
     18   250 !
      3   300 !

Class:  9  Total msg: 1996  Desired delay:  256  Mean delay:  299.5
Mean early:  12.4  Max early:  64  Mean late:  49.5  Max late: 634
      0 -9999 !
      0  -150 !
      1  -100 !
    190   -50 !*****
   1171     0 !***********************************
    435    50 !*************
    178   100 !*****
     13   150 !
      5   200 !
      0   250 !
      3   300 !

Class: 13  Total msg: 2055  Desired delay:  256  Mean delay:  296.5
Mean early:  13.2  Max early:  45  Mean late:  46.8  Max late: 627
      0 -9999 !
      0  -150 !
      0  -100 !
    214   -50 !******
   1209     0 !***********************************
    481    50 !**************
    147   100 !****
      1   150 !
      0   200 !
      0   250 !
      3   300 !
```

Figure 4-8. Example Of REDUCE Printout
(Sink Statistics)

need to understand in detail the construction or use of PC messages by EXPAK and the SIMPs. The features currently implemented include the following:

a. Create PC message file or merge with existing file.

b. Insert and delete individual unformatted (octal) parameters in the message.

c. Change the destination or routing of one or more PC messages in the file.

d. Edit individual formatted parameters in the message using an interactive script in which each member of a set of selected parameters is displayed with a short explanation of its function. The experimenter is then invited to provide a new value, change the existing one, or delete it.

e. Print the contents of the PC message file in either formatted or unformatted (octal) form. The unformatted form is in the EXPAK format (see Figure 4-4). An example of the formatted form, which is keyed to the editing features, is shown in Figure 4-9. This form is also produced by the PRINT program from a readout message, allowing quick comparisons for verification and error detection.

This program is most useful in the construction of new PC message files or in systematic changes to the existing ones. Features for the construction of parameters whose values are encoded into portions of the same 16-bit SIMP word are especially valuable, since this ordinarily would involve tedious computations by hand in the octal radix.

```
SUM20 E Readout Message

1. Experiment Control Parameters

Destination ID: 2561 52 0  Path code: 0  Cumstat groups: 65408 0

Switches
 Stop all: 7  Stop msg generator: 0  Stop cumstats: 0  Send readout: 0

Start times (HF)
 Parameter turnon: 0  Message generators turnon: 0
 Cumstat turnon: 65529  Cumstat interval: 30  Cumstat turnoff: 65377

2. Special Parameters (enabled at parameter turnon)

Protocol code (0-TPODA 1-CPODA 2-FTDMA): 1  TTG/priority mask: 65535
 Msg buffer chunk size: 10  My rec rate: 1  Lowest rec rate: 1  Mode: 0

Slot/frame sizes
 FTDMA slot: 3  PODA frame: 29  Min control subframe: 2  Ack waiting window: 512

3. Delay Class/Priority Parameters

Active delay classes: 8  Active priority classes: 4

Delay class definitions (4-7 map onto 0-3)
 Delay class      0      1      2      3      4      5      6      7
 Desired delay    0      64     128    256    512    1024   2048   4096
 Max hold time    0      128    256    1024   1024   2048   4096   8192

4. Miscellaneous Protocol Parameters

Max reservations per pkt: 2  Queue length sampling interval: 10

Congestion controls      0      1      2      3
 Thresholds              5      10     15     20
 Holding-time offsets    50     100    150

Station synchronization thresholds
 Out-of-sync: 5  In-sync: 5  Return-to-initial-acquisition: 5
 Reservation-exist: 20  Channel-empty: 5  Channel-noisy: 5
```

Figure 4-9. Example of PCMSG Output

4.4 EXPERIMENT RESULTS

4.4.1 Introduction

The COMSAT experimental program has emphasized the algorithmic verification and performance evaluation of SATNET under a wide variety of operating conditions at the expense, if necessary, of parametric optimization and fine tuning under well-behaved operating conditions. This report does not attempt to present a complete and detailed history of the experiments. In many cases, the experiments were designed to reveal inadequacies in the channel-scheduling algorithm and to aid in its fine tuning. In most cases, after an inadequacy was found, its cause was identified and the problem was solved. Accordingly, the experimental results presented in this subsection demonstrate that the operation of the hardware and software conforms to design expectations. Of particular interest is system behavior when confronted with traffic intensities near the extrema of the capacity envelope where the resource management and fairness principles of the scheduling algorithm are most vulnerable.

There were three classes of experiments identified in the COMSAT experiment plan: those involving steady-state measurements, reported in the last Interim Report; those involving nonsteady-state measurements; and those involving the small station. The results of the latter two will be the principal topic in this subsection; however, the results of the former will be summarized for completeness.

All of these experiments with the exception of the Clarksburg experiments involved three SIMPs: Etam, Goonhilly, and Tanum. The Clarksburg experiments involved only the Clarksburg SIMP in loopback mode to itself. These experiments will be discussed in a later section. Figure 4-10, which shows the SATNET configuration used, evidences that each SIMP had two active message classes, the α class and the β class, with each individual generator sending along

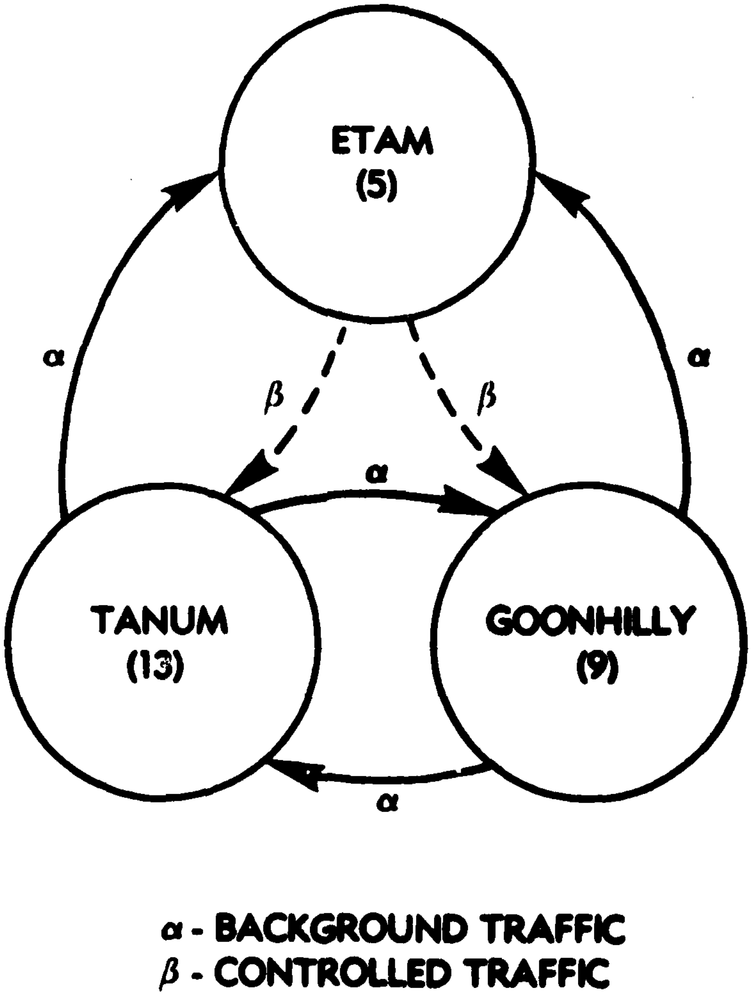

Figure 4-10. Three-Station Experimental Configuration

the source/sink path shown. This configuration was chosen, both for its flexibility in adapting to different experimental designs and for the detection of asymmetries due to anomalies and breaches of fairness principles in the scheduling algorithm.

The experiments are distinguished only by the parameters assigned to the two generator classes α and β, and in the activation schedule for the message generators. For both classes, a geometric message interarrival distribution with a mean of $1/\lambda$ message per second (depending on the activation schedule) and a uniform message length distribution with a mean of 64 words (1024 bits) was chosen. These distributions were selected for the following reasons:

a. to approximate the characteristics of the ARPANET traffic, which will probably be carried on this system in the near future;

b. to provide reasonably close agreement with assumptions made in the simple channel model (described in the last Interim Report); and

c. to avoid the spurious behavior which occurred with fixed interarrival intervals and/or message lengths when the operation of the scheduling algorithm and the message generation algorithm developed undesirable synchronization tendencies.

The remaining parameters in the α and β message classes are the mean generation rate, priority, delay class and holding time specifications. In all experiments, the priority was fixed and the holding time was as long as practicable, so that the behavior could be compared with that predicted by the model. Thus, the operative parameters that distinguish the various experiments were the generation rate and the delay-class specifications.

4.4.2 Channel Calibration and Proof-of-Performance

In order to assure that the channel was performing satisfactorily, at least under operating conditions likely to prevail in practice, a set of calibration measurements were performed and documented in the last Interim Report. These experiments, which parallel those made by UCLA and NDRE, verify that the scheduling algorithm performs as expected and that the channel can be loaded under steady-state conditions to approximately full utilization of the information frame. These experiments are briefly summarized below.

In the first experiment, which was designed as a baseline for the other two, the generation-rate and delay-class parameters were identical for both the α and β classes, so that the configuration was symmetric. This experiment involved measurements over a wide range of generation rates from a level corresponding to essentially a single message in the system at any given time to a level corresponding to complete saturation of network resources. Figure 4-11 shows the system behavior for traffic rates over this range. In this and the following figures, the traffic rates are relative to the 128-vs CPODA frame (65 kbit), so that the maximum relative rate is $(128-16-2)/128 \simeq 0.84$. Both mean message delay (in vs) and control-channel traffic (in vs/vs) are evident in the Figure. The dip in control-channel traffic at about 0.82 relative rate is believed to correspond to a quasi-stable state in which substantially all the reservations and acknowledgments are transmitted in piggyback mode (see below). Note that, even at relative rates close to channel saturation, control-channel traffic is less than the slotted-ALOHA saturation limit of $2/e (\simeq 0.36)$.

The throughput-delay performance of the channel under these steady-state uniform traffic conditions can be summarized as follows: At a relative rate of 0.8, the mean message delay is about three seconds, corresponding to an incident traffic level of about 17 kbit per station, or 17 64-word (1 kbit) packets per second. At low relative rates, the delay is approximately constant at a little over a second and includes components due to CPODA frame synchronization, reservation transmission, and message transmission. On the average, a packet arrives at the station in the middle of a CPODA frame and the reservation is transmitted as a control packet in the next frame. New reservations can be entered in the central queue only after the end of this frame plus the round-trip time. In the experiments, the

CPODA frame is 29 vs long and the round-trip time is about 27 vs. Thus the delay is about 1.5 x 29 + 2 x 27 = 98 vs. The delay time plus the transmission time of slightly less than 2 vs plus minor synchronizing delays accounts for the minimum value mean delays shown in Figure 4-11.

The second experiment was designed to exercise the algorithm that allocates reservations and acknowledgments to packet headers. This algorithm either initiates a control (contention) packet or waits for a scheduled information (message) packet to send reservations or acknowledgments according to the algorithm described in PSPWN 109. In the experiment, the algorithm was applied by varying the "foreground" traffic in class α from low to high rates while maintaining the "background" traffic in class β constant at about one-half the rated channel capacity.

Figure 4-12 shows the mean message delay, control-channel traffic, and piggyback-channel traffic under these conditions. The two stations in the β class transmit at about 0.24 relative rate (15 kbit). The third station, which is in the α class, transmits at relative rates ranging from approximately zero to about 0.36. At the lower rates, the third station can seldom piggyback reservations and acknowledgments and must use the control channel. At higher rates piggyback opportunities increase, as is evident from the dip in the control-header traffic as the α-class traffic level increases from zero. The piggyback/control header selection algorithm loads piggyback headers up to 90 percent of the full capacity, and loads the control headers to about half this value. Obviously, this is a compromise. Loading control headers more heavily would result in longer delays, while loading them more lightly would result in increased control-channel congestion.

The third experiment was designed to test the operation of the scheduling algorithm when two delay classes were present. As in the second experiment, the foreground traffic was varied and the background traffic was held constant. However, the desired delay of the background traffic was set at twice that of the foreground traffic;

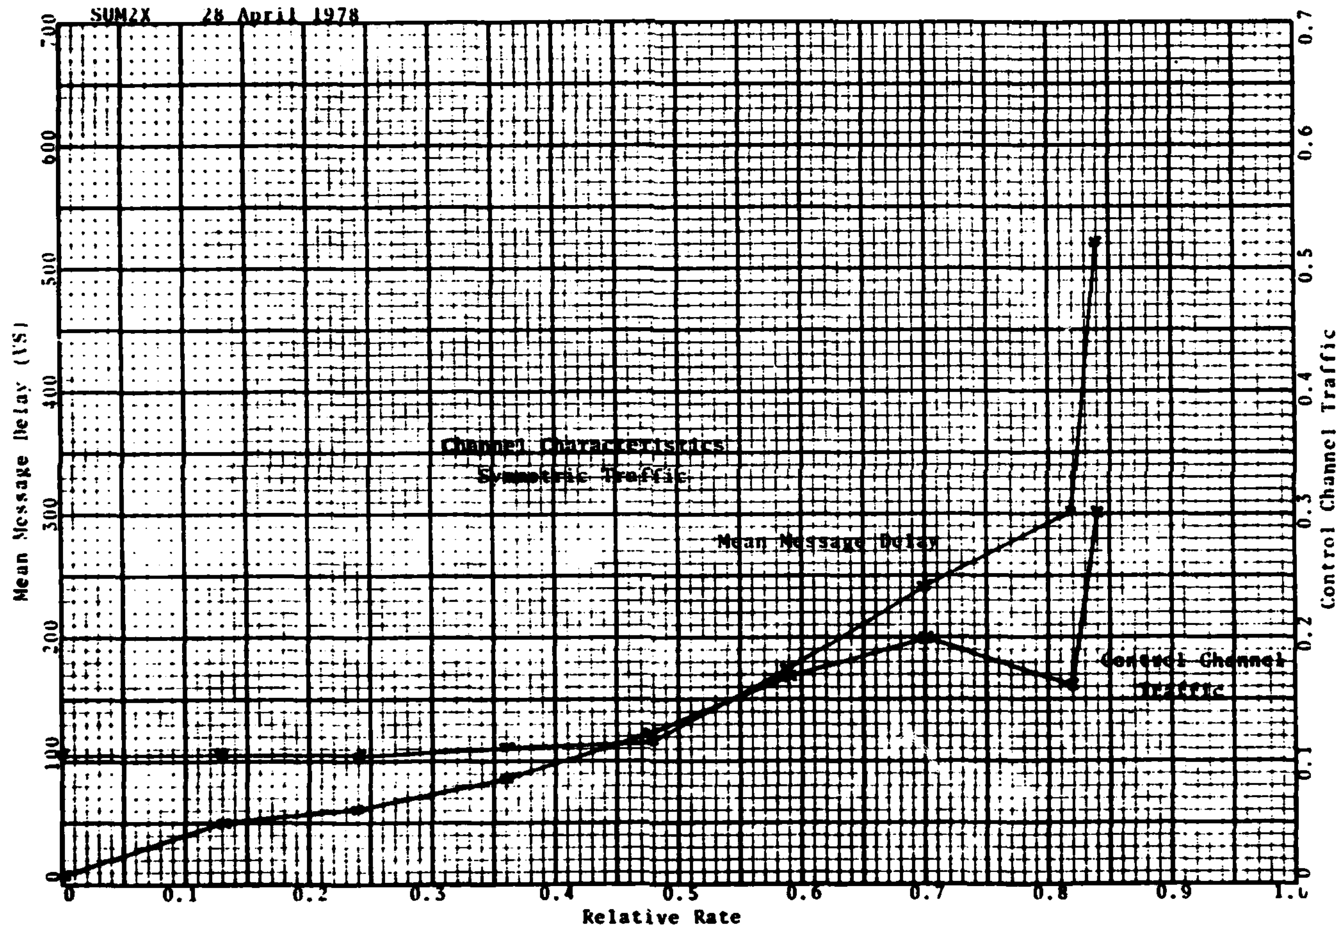

Figure 4-11. Baseline Experimental Results

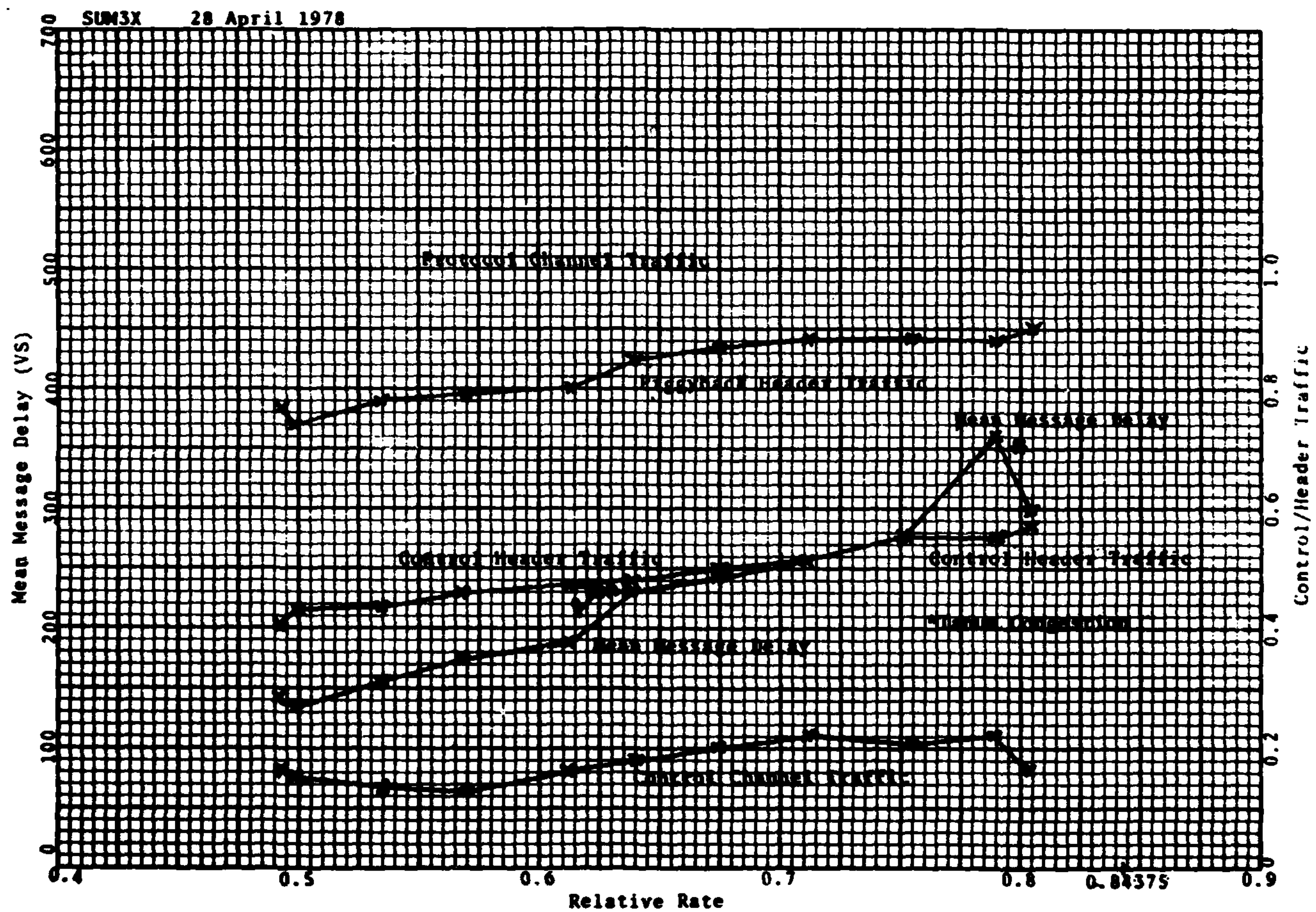

Figure 4-12. Experimental Results; Background Traffic at One-half Rated Channel Capacity

therefore, it would be expected that an increase in the foreground traffic would preempt the background traffic, which is clearly evident in Figure 4-13. As the relative rate increases for the 256-vs traffic, (α) the mean delay for the 512-vs traffic (β) increases and the 256-vs traffic remains at or near the nominal desired value. As in the previous experiment, the decrease in control-channel traffic is due to increased piggybacking opportunities.

4.4.3 Problems in Channel Scheduling and Stability Controls

The experiments performed during the last reporting interval and documented in the last Interim Report identified a number of problems in channel scheduling and stability controls. These problems fell into three main areas: (a) apparent processing inadequacies, manifested as excessive reservation-synchronization dropouts and other transient phenomena; (b) breach of fairness principles, manifested as capture of the channel by one or a group of stations causing denial of service to the remainder; and (c) large variations in delay for end-to-end traffic in excess of what could be expected according to the control models developed by UCLA and COMSAT. These problems will be discussed in this section.

Reservation synchronization refers to the state in which all stations have verified that their copy of the global scheduling information is consistent with each other and that normal reservation, transmission and acknowledgment operations can proceed. If reservations are not received by a station or if a station cannot transmit at a previously-scheduled time, then this consistency is lost. If a station perceives, by listening to the packets transmitted by it and the other stations, that this consistency is lost for an extended time, then it declares itself out of synchronization and proceeds to try to rebuild a consistent copy of the global data base. During the time this is going on, it is unable to transmit data messages and the message delays through the station increase accordingly, leading in extreme cases to messages being dropped due to expiration of holding time, etc.

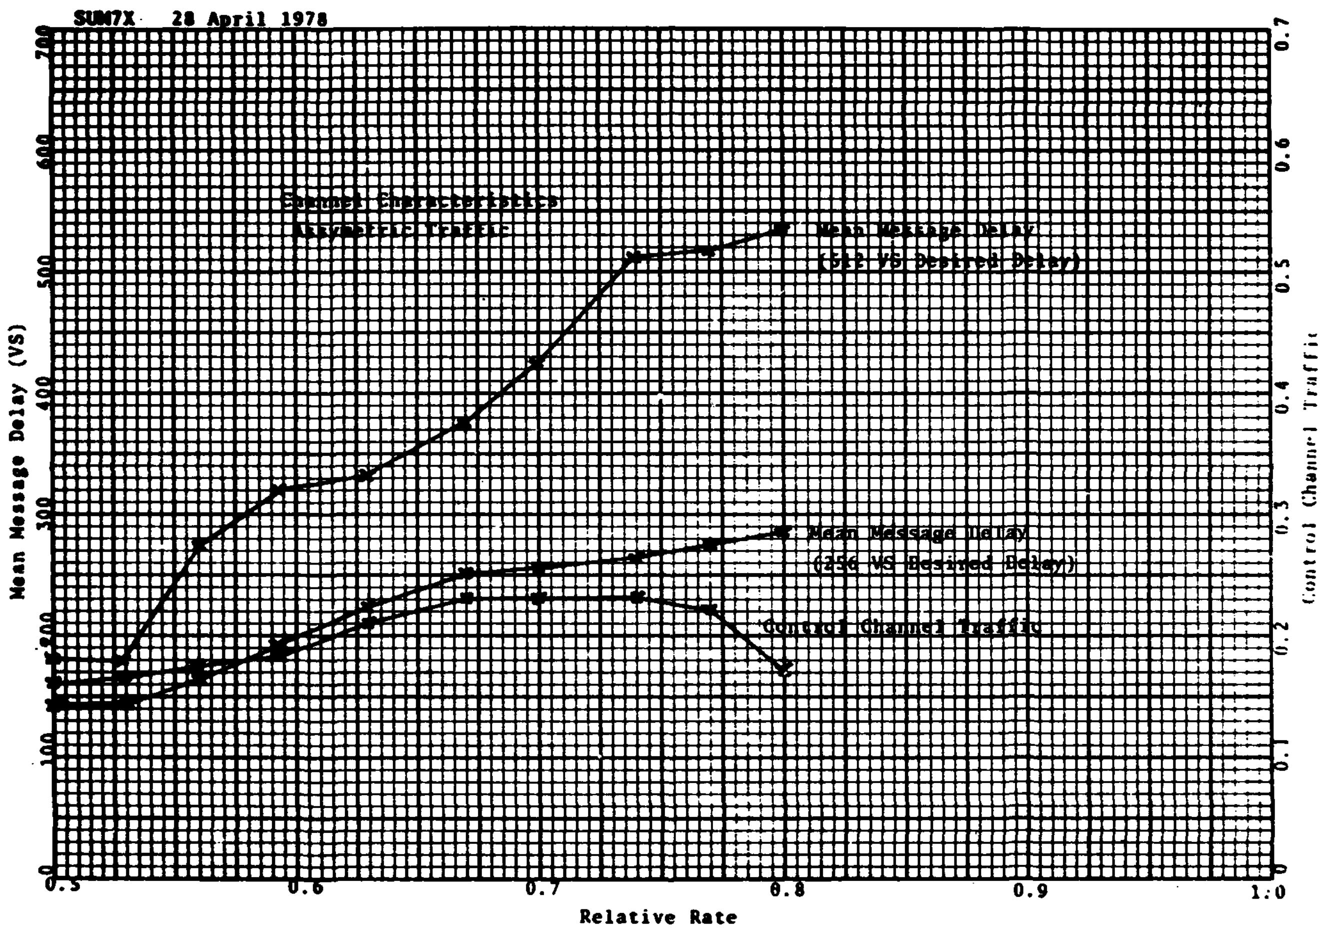

Figure 4-13. Experimental Results With Two Delay Classes

A station can lose reservation synchronization when it is unable to hear packets or transmit them at the scheduled time. This can be due to local noise at the station, packet collision due to lack of global data-base consistency, or inability of the SIMP program itself to keep up with the processing required. On a number of occasions during the refinement of the CPODA implementation in the Honeywell processor, it was observed that processing delays, rather than channel noise, were causing reservation-synchronization loss. This was most evident at the higher traffic levels and with either very short or very long messages. Subsequently, it was found that the overhead for sorting the central queue used for global scheduling was impacting the short-message case and the overhead to copy and reformat buffers used internally in the SIMP was impacting the long-message case. Both of these problems have since been solved, in one case by fixing a bug in the software and in another by restricting the maximum message length. However, during an appreciable period in the experiment and measurement program, it was not possible due to these causes to obtain consistent three-station measurement data and many data points had to be discarded.

A considerable number of measurements were performed to evaluate the robustness of the system when subject to traffic levels near and exceeding saturation. At these levels, the congestion controls and fairness principles come into play as they affect the acceptance and reservation-making behavior at the individual stations. In order to examine the behavior in detail, a number of experiments were performed using traffic levels near to and slightly below and above information-subframe saturation. Figures 4-14 and 4-15 illustrate the results obtained. Figure 4-14 is a detailed summary of the steady-state conditions for the three-station homogeneous-traffic ($\alpha = \beta$) system operating at an aggregate relative rate in the information subframe of about 0.82. During the approximately two-minute interval (150 vs) of the experiment, each station transmitted about 2000 information packets and 50-100 control packets, of which 70-80 percent were received correctly. (The small numbers

SUM27 Combined Statistics

1. Frame Utilization	Hello frame	TDMA frame	Info subfrm	Ctl subfrm	
Total frames	150	600	600	600	
Total length (VS)	19200	17400	15721	1679	
Relative rate	1	0.906	0.818	8.7000	
Max length (VS/frame)	128	29	27	29	
Mean length (VS/frame)	128	29	26.201	2.798	
2. Uplink Channel	**Etam**	**Goony**	**Tanum**	**Clark**	**Total**
Info pkt transmitted	1836	1973	2003	0	5812
Info pkt retrans	1	3	3	0	7
Ctl pkt successes	92	61	56	0	209
Ctl pkt failures	14	26	17	0	57
Hello pkt transmitted	150	150	150	0	450
Discarded	0	-1	-2	0	-3
3. Downlink Channel	**Etam**	**Goony**	**Tanum**	**Clark**	**Total**
Info pkt received	5813	5813	5812	0	17438
Ctl pkt received	208	208	208	0	624
Hello pkt received	450	450	450	0	1350
Discarded	2	2	3	0	7
4. Channel Protocol	**Etam**	**Goony**	**Tanum**	**Clark**	**Total**
Msg generated	1832	1965	2012	0	5809
Msg accepted	1832	1965	2012	0	5809
Res attempts	1863	2039	2059	0	5961
Res successes	1836	1991	2026	0	5853
Msg sent to channel	1837	1976	2006	0	5819
Msg rcvd from channel	2013	1930	1869	0	5812
Msg delivered to sink	2030	1978	1896	0	5904
Ack attempts	2018	1974	1901	0	5893
Ack successes	2008	1956	1884	0	5848
Msg dropped	0	0	0	0	0
5. Mean Queue lengths	**Etam**	**Goony**	**Tanum**	**Clark**	**Total**
Freelist	33.383	35.464	35.887	0	104.73
QI	18.455	19.073	18.7	0	56.229
QH	0.032	0.099	0.015	0	0.146
QC	15.351	15.38	15.35	0	46.083
6. Packet Utilization	**Etam**	**Goony**	**Tanum**	**Clark**	**Mean**
Mean res ctl subfrm	0.103	7.8000	0.061	0	0.08
Mean ack ctl subfrm	4.1000	0.028	0.027	0	0.032
Mean load ctl pkt	0.883	0.866	0.835	0	0.864
Mean load piggy pkt	0.574	0.574	0.569	0	0.572
7. Common Channel	**Info**	**Ctl**	**Ctl hdr**	**Pig hdr**	
Capacity	15721	1679	266	5819	
Traffic	0.37	0.158	0.864	0.572	
Throughput	0.369	0.123	0.681	0.572	

Program to format and print summary statistics

Figure 4-14. Summary of Steady-State Conditions

```
SUM27 Combined Statistics

8. Sink Statistics

Class:  5  Total msg: 1853  Desired delay:  256  Mean delay:  304.2
Mean early:  14.8  Max early:  69  Mean late:  55.4  Max late: 621
    0 -9999 :
    0  -150 :
    5  -100 :
  185   -50 :*****
 1090     0 :*************************************
  361    50 :************
   99   100 :***
   48   150 :*
   44   200 :*
   18   250 :
    3   300 :

Class:  9  Total msg: 1996  Desired delay:  256  Mean delay:  299.5
Mean early:  12.4  Max early:  64  Mean late:  49.5  Max late: 634
    0 -9999 :
    0  -150 :
    1  -100 :
  190   -50 :*****
 1171     0 :*************************************
  435    50 :**************
  178   100 :*****
   13   150 :
    5   200 :
    0   250 :
    3   300 :

Class: 13  Total msg: 2055  Desired delay:  256  Mean delay:  296.5
Mean early:  13.2  Max early:  45  Mean late:  46.8  Max late: 627
    0 -9999 :
    0  -150 :
    0  -100 :
  214   -50 :******
 1209     0 :*************************************
  481    50 :***************
  147   100 :****
    1   150 :
    0   200 :
    0   250 :
    3   300 :
Program to format and print summary statistics
Enter filename<cr> ?
```

Figure 4-15. Experimental Throughput-Delay Characteristics

in the Discard rows are not significant and are due to limitations in the cumstat recording mechanisms in the SIMP. This minor problem also results in some data not adding precisely to the totals indicated and some minor inconsistencies such as slightly more messages delivered to the destination than were received from the channel.)

At the time these measurements were made, it was possible to track, on the average, where messages were spending time in the system. Thus it was known from the length of the freelist, that is, the number of buffer "chunks" available for allocation to a message, and the length of the local input queues (QI), host output queues (QH) and central queue (QC) that about 105 chunks were allocated to messages (one chunk per message), 56 of which were waiting on the QI queues, 46 on the QC queue and the remaining three scattered on other queues in the system. During the 196-second run, about 5800 messages were generated, or about 30 messages/second. From Little's result [ ] the mean delay T can be deduced from the number in system $N = 105$ and arrival rate $\lambda = 30$,

$$T = N/\lambda = 105/30 = 3.5 \text{ seconds}$$

The fact that this figure is somewhat high, compared with the approximately 3.1-sec value actually measured (see Figure 4-14), can be explained by the fact that the queue-length statistics are derived from a sampling procedure which is known to be rather imprecise. Nevertheless, this kind of calculation serves as a check on other calculations, as well as an indicator of where the principal delays are in the system.

Continuing in the assessment of the data summarized in Figure 4-14, it can be noted that, at the traffic level indicated, the bulk of all reservations and acknowledgments were transmitted using information-subchannel headers with the remainder (8 percent reservations, 3 percent acknowledgments) transmitted using control-subframe headers (contention packets). In both cases, the efficiency of the headers, that is, the fraction of the available bits actually used, is comfortably high between about 0.57 and 0.68.

Figure 4-15 summarizes the throughput-delay characteristics of the system under the conditions already described. All data in this Figure are in virtual-slot units, where a virtual slot is 10.24 ms. Traffic in class 5 is going to Etam, that in class 9 to Goonhilly, and that in class 13 to Tanum. Note the consistency in the mean delay for traffic in each class and, especially, the high degree of similarity between the shapes of the delay histogram of each class. The conclusion is that the scheduling operations are stable and that perturbations are minor.

Increasing the generation rate slightly above the 30 message/sec level can cause the scheduling operations to break down, as is evident from Figure 4-16. Here the relative rate is 0.84, but the mean delay has increased to over 4.3 seconds for Etam and Tanum and over twice that for Goonhilly. It is the latter observation which is of most concern, since it is evident from the delay histograms that control-subchannel congestion and retransmissions are affecting Goonhilly much more than the other stations. What is happening here is an instability in the algorithm which assigns reservation/acknowledgment requests to packet headers and initiates control-packet transmissions. In the case shown, Etam and Tanum are piggybacking substantially all their reservation/acknowledgment data on messages to each other, and squeezing the control subchannel size down to only about 2.2 vs. Meanwhile, Goonhilly has been unable to capture enough information-subchannel bandwidth to piggyback its own requests and must compete with the other station's requests (primarily to Goonhilly itself) for use of the control subchannel. Although relatively few control packets are lost due to contention, the high utilization of the control-subframe header (greater than 0.9) indicates that traffic to Goonhilly is limited by the available bandwidth in the control-subframe itself. In this case, increase in the minimum control-subchannel size above 2 vs would not help unless a station could transmit more than once in the control subframe.

```
SUM29 Combined Statistics

1. Frame Utilization      Hello    MODA     Info     Ctl
                          frame    frame    subfrm   subfrm
 Total frames             150      600      600      600
 Total length (VS)        19200    17400    16059    1341
 Relative rate            1        0.906    0.836    0.069
 Max length (VS/frame)    128      29       27       29
 Mean length (VS/frame)   128      29       26.765   2.235

7. Common Channel         Info     Ctl      Ctl hdr  Fis hdr
 Capacity                 16059    1341     499      5947
 Traffic                  0.37     0.372    0.919    0.552
 Throughput               0.369    0.324    0.811    0.551

8. Sink Statistics

Class:  5  Total msg: 2131  Desired delay:  256  Mean delay:  417.4
Mean early:  10.5  Max early:  31  Mean late: 163.9  Max late: 790
    0 -9999 :
    0  -150 :
    0  -100 :
   31   -50 :
  315     0 :*********
  673    50 :******************
  313   100 :*********
  249   150 :********
  118   200 :***
  120   250 :***
  314   300 :*********

Class:  9  Total msg: 1715  Desired delay:  256  Mean delay:  926.8
Mean early:  21.9  Max early:  42  Mean late: 679.3  Max late:1633
    0 -9999 :
    0  -150 :
    0  -100 :
   21   -50 :
   61     0 :**
  138    50 :****
  131   100 :****
   93   150 :***
   25   200 :
    5   250 :
 1260   300 :*********************************************

Class: 13  Total msg: 2221  Desired delay:  256  Mean delay:  437.0
Mean early:  11.4  Max early:  37  Mean late: 185.7  Max late:1112
    0 -9999 :
    0  -150 :
    0  -100 :
   53   -50 :*
  401     0 :***********
  690    50 :******************
  289   100 :********
  183   150 :****
  141   200 :***
  103   250 :**
  361   300 :**********
```

Figure 4-16. Throughput-Delay Characteristics With Increased Generation Rate

The phenomenon of channel capture graphically illustrated in the preceding example has been known and studied by other SATNET experimenters and can happen at traffic levels less than saturation. At the levels considered here, however, the limitation in control-subchannel bandwidth makes it even less likely that a station needing its fair share of information-subchannel bandwidth can break the deadlock created by the two other stations. Clearly, the congestion controls and throttling mechanisms are inadequate, as perhaps is the decision procedure used for reservation/acknowledgment header allocation and control-packet initiation. These problems are understood to be now under study at BBN.

A different kind of fairness problem is illustrated in Figure 4-17. This experiment involved traffic from two delay classes, one of 256-vs desired delay to Etam and the other of 512-vs desired delay to Goonhilly and Tanum. As in previous cases, traffic for these sites is split with two sites contributing about one-half the total traffic to the third (see Figure 4-6). One might expect from this description that the 512-vs traffic to Goonhilly and Tanum would exhibit about the same mean delay and delay histogram shape, but from Figure 4-17 this clearly is not the case. Subsequent to these measurements, a bug was found in the SIMP code which explained this anomaly, as well as another described presently. A point worth mentioning here is that this kind of problem would probably not have been discovered during the developmental period except for three-station experiments such as these, and might have caused considerable disruption after transition to operational status.

4.4.4 Experimental Results with Time-Varying Traffic

The experimental results reported previously have pointed up a need to investigate the channel response in message throughput and delay, to time-varying traffic excitations. Previous results have been reported in this project by UCLA for the slotted-ALOHA

```
.UN73 Combined Statistics

1. Frame Utilization       Hello   PODA    Info     Ctl
                           frame   frame   subfrm   subfrm
 Total frames              150     600     600      600
 Total length (VS)         19200   17400   11392.   6007.6
 Relative rate             1       0.906   0.593    0.312
 Max length (VS/frame)     128     29      27       29
 Mean length (VS/frame)    128     29      18.987   10.012

7. Common Channel          Info    Ctl     Ctl hdr  Pig hdr
 Capacity                  11392.  6007.6  1113     4182
 Traffic                   0.367   0.185   0.853    0.461
 Throughput                0.363   0.12    0.551    0.455

8. Sink Statistics

Class: 5  Total msg: 802  Desired delay: 256  Mean delay: 193.0
Mean early: 95.7  Max early: 190  Mean late: 81.5  Max late: 568
  106 -9999 !########
  219  -150 !#################
  185  -100 !##############
  144   -50 !###########
   86     0 !#######
   39    50 !###
    9   100 !
    2   150 !
    0   200 !
    0   250 !
   12   300 !

Class: 9  Total msg: 1707  Desired delay: 512  Mean delay: 154.4
Mean early: 364.0  Max early: 447  Mean late: 182.3  Max late: 259
 1687 -9999 !###############################################################:
    0  -150 !
    0  -100 !
    0   -50 !
    0     0 !
    0    50 !
    7   100 !
    5   150 !
    7   200 !
    1   250 !
    0   300 !

Class: 13  Total msg: 1722  Desired delay: 512  Mean delay: 483.3
Mean early: 68.6  Max early: 381  Mean late: 48.8  Max late: 629
   99 -9999 !###
  165  -150 !#####
  335  -100 !############
  538   -50 !###################
  444     0 !################
  107    50 !###
   16   100 !
    3   150 !
    0   200 !
    0   250 !
   15   300 !
```

Figure 4-17. Experimental Results

rated

case, but not for the CPODA case. Unfortunately, it has not been possible to conduct CPODA experiments similar to the ALOHA experiments due to the limited time available. The EXPAK software and the extensive set of supporting software described previously in this report do not provide sufficiently fine granularity to study channel behavior phenomena in the order of a roundtrip delay. Such experiments would have to be conducted on an end-to-end basis using time-stamped messages such as might be done using the Gnome controllers in the gateways, for example.

On the other hand, the EXPAK-related software can be used to detect and study unstable regions where unsuspected interactions between the channel-scheduling and fairness algorithms could occur, for example. The existence of such regions could be inferred by discontinuities or hysteresis effects in the throughput-delay characteristic, for example. A number of experiments of this type were performed using EXPAK features and supporting software constructed for this purpose. A useful byproduct was a technique that greatly increased the efficiency of conducting the experiments and allowed the collecting of the data for an entire plot of channel characteristics versus time in a single run, rather than the multiple runs previously required.

In order to perform time-varying experiments, a method to control the SIMP message generators in a time-varying manner is required. This was accomplished using a set of four message generators in each SIMP for each of the two classes of traffic, α and β, mentioned previously in this report. The composite of eight message generators was controlled using a script which specified the time (in units derived from the basic virtual-slot interval) and duration for each activation of each generator. In most experiments the traffic intensity of each of the four generators was prespecified in the ratios 1:2:4:8, so that sixteen traffic levels could be established using appropriate combinations of the generators and the assumption of superposition. At all except the highest traffic levels, where the processing required by the CPODA algorithm significantly detracts from the processing required by the message generators themselves, the superposition assumption is reasonable.

In one set of early experiments a number of runs were performed with three stations and uniform traffic levels varying from zero to saturation. In some of these runs a distinct tendency toward instability could be observed at various discrete traffic levels. After considerable effort had been expended toward identification of the cause they simply disappeared one day. Subsequently it was learned that a programming bug in the channel-scheduling algorithm had been found and fixed at about that time. This was the same bug that resulted in the instability noted in the previous section with respect to the fairness principles.

The channel throughput-delay characteristic in the present system, with the bug fixed, is shown in Figures 4-18 and 4-19. In these and subsequent Figures the time-interval number in the range 1-16 is shown along the left margin. Each interval corresponds to 60 Hello Frames (about 78 sec) during which three sets of cumstats were collected from each of the three stations. The three sets of cumstats were compared among themselves to verify their validity and combined to yield the data shown for each interval. In all Figures the first point (interval 1) and the last (interval 15) represent possibly inconsistent data, due to various synchronizing and housekeeping functions inherent in the data collection process. The results illustrated in Figures 4-18 and 4-19, which were collected at Etam, clearly show the effect of the congestion controls at the higher traffic intensities (intervals 12-15) where the SIMP is refusing some generated messages.

The operation of that portion of the channel-scheduling algorithm that initiates control packets or waits for piggyback opportunities is clearly shown in Figure 4-20. This Figure shows the fraction of all reservations sent via control-packet headers, demonstrating the shift from control-packet to information-packet headers as the traffic level increases. In spite of this shift the total offered control-packet rate increases with traffic to a maximum just short of saturation (interval 12) and then decreases above that. The explanation for the decrease above this level can be found in Figure 4-19, which shows the mean delay characteristic. At the traffic levels in question, the mean delay for information packets rises significantly, but the throughput does not.

5. PSP TERMINAL DEVELOPMENT AND IMPLEMENTATION

COMSAT's main activity during the period covered by this final report has been the design, fabrication and installation of the new PSP terminal. The terminal includes the Phase 2 SSI, the burst modem, the Command and Monitoring Module (CMM) and the data test set (DTS) which were developed for the Phase 2 system. The terminal also includes a complete Phase 1 system that contains all of the capabilities of the present SPADE system. These subsystems modules will be described in this section in the context of overall PSP terminal operation.

The PSP terminal is a self-contained baseband-to-IF package which interfaces the SIMP at baseband and the earth station common IF subsystem at 70 MHz. This terminal will replace the SPADE modem and Phase 1 SSI which have been used for packet transmission in the DARPA Packet Satellite Experiment for the past three years.

The PSP terminal contains redundant QPSK modems, frequency synthesizers, and an SSI data module, as well as a non-redundant transmission mode controller/encoder, a reception mode controller/decoder, a switch matrix, a DTS, a CMM, a time-frequency unit (TFU), and an IF subsystem. Hence, redundancy is not provided for transmission to small stations, although transmission to large stations is fully redundant. In fact, the PSP terminal can support simultaneous operation on two different SCPC channels, with only one channel supporting a 64-kbit/s transmission rate. Figure 5-1 shows the physical characteristics of the PSP terminal. It is 24 in. wide by 25.5 in. deep by 70 in. high. A photograph of the terminal is shown in Figure 5-2.

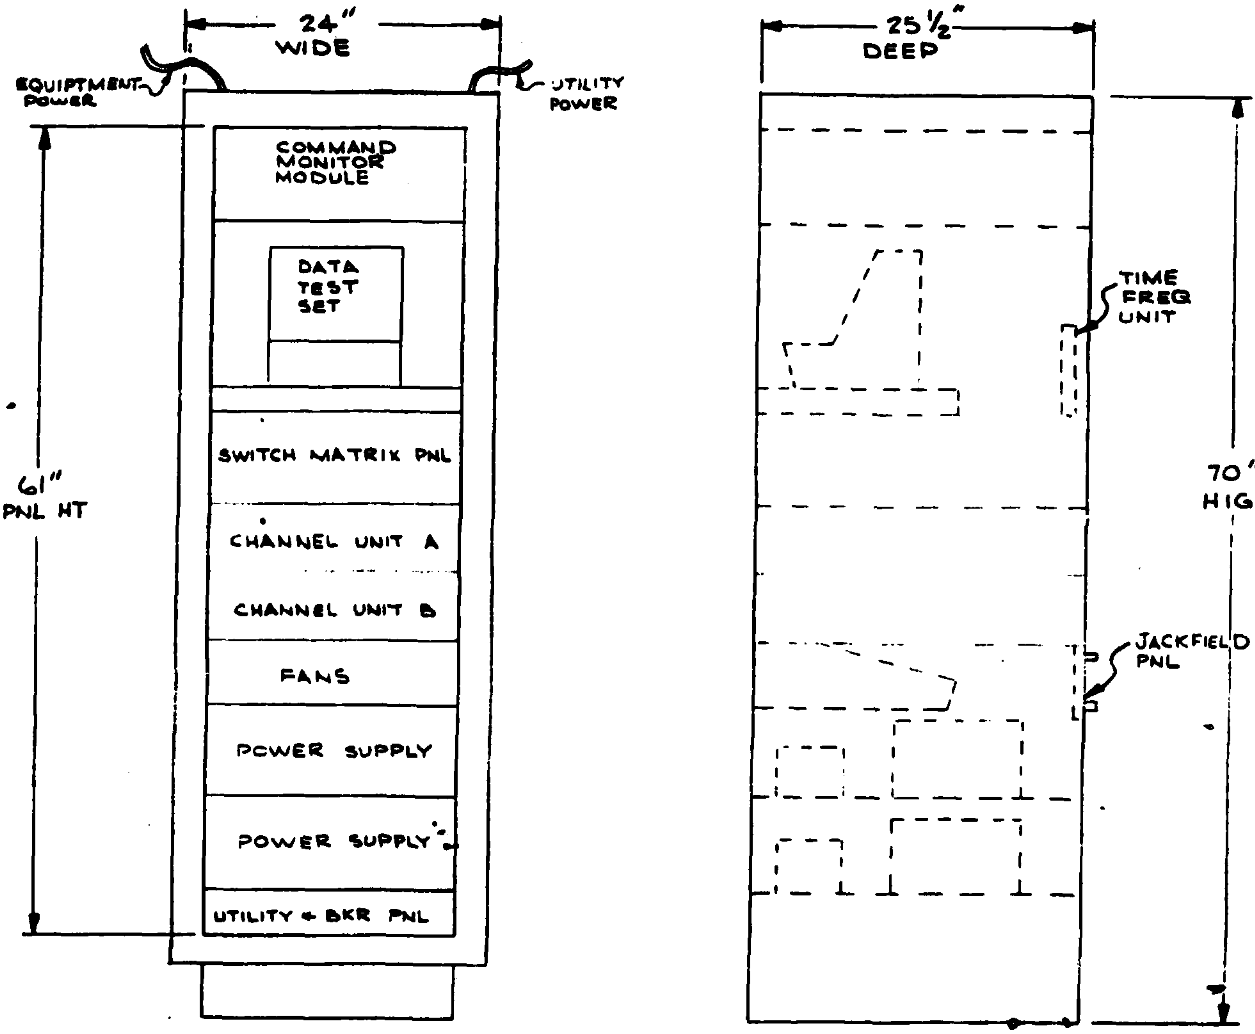

Figure 5-1. Physical Characteristics of the Packet Satellite Program Terminal

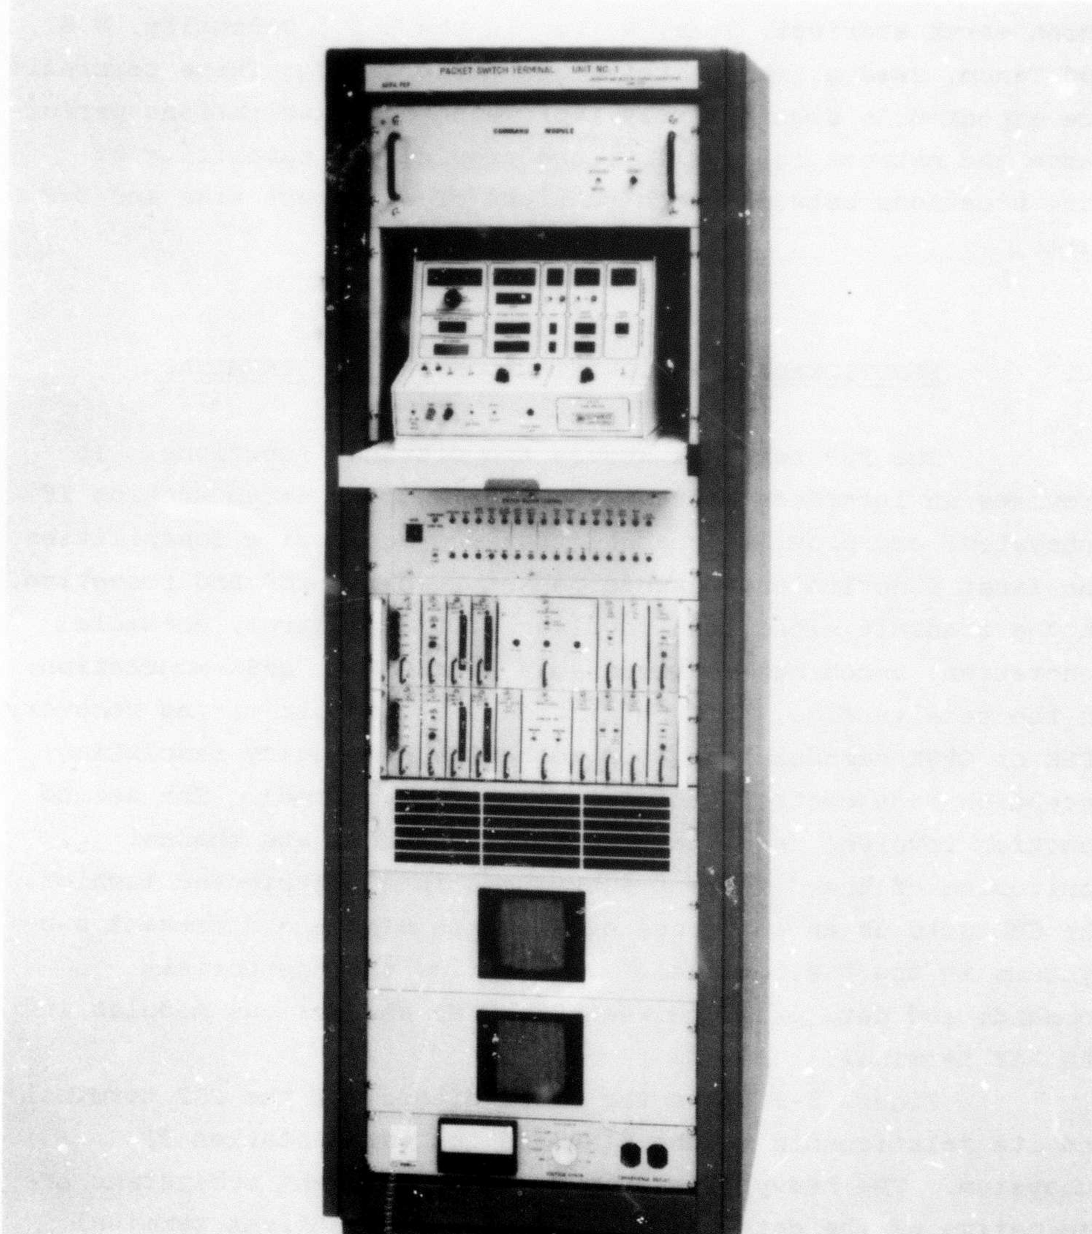

Figure 5-2. The PSP Terminal

The PSP terminals are scheduled for installation at the three earth stations, Etam, W. Va. in the U.S., Goonhilly, U.K., and Tanum, Sweden, in the third quarter of 1979. These terminals are expected to significantly improve both communications performance and network reliability, and provide the capability of communications between earth stations of different size and G/T ratio.

5.1 FUNCTIONAL CHARACTERISTICS OF THE PSP TERMINAL

The PSP terminal serves two distinct functions. It provides an interface between the SIMP and the earth station IF subsystem, and provides channel test and monitoring capabilities. The first function encompasses packet transmission and reception. On the transmit side, this involves carrier control, preamble generation, encoding (if required), and BPSK or QPSK modulation; on the receive side, this entails carrier and bit timing recovery, BPSK or QPSK demodulation, SOM detection, ambiguity resolution, reception mode control, and decoding (if required). The second function involves parameter and state control, and channel monitoring of the different subsystems in a distributed fashion. The CMM acts as an interface between the SIMP and different subsystems in the PSP terminal. It supplies the appropriate commands and data paths between the SIMP and various modules in the PSP terminal.

Figure 5-3 shows the block diagram of the PSP terminal and its relationship to the SIMP and the earth station IF subsystem. The heavy lines between the different subsystems are indicative of the data paths and emphasize the first terminal function. The light lines between the CMM and other subsystems indicate the explicit command and monitoring paths and are related to the second terminal function.

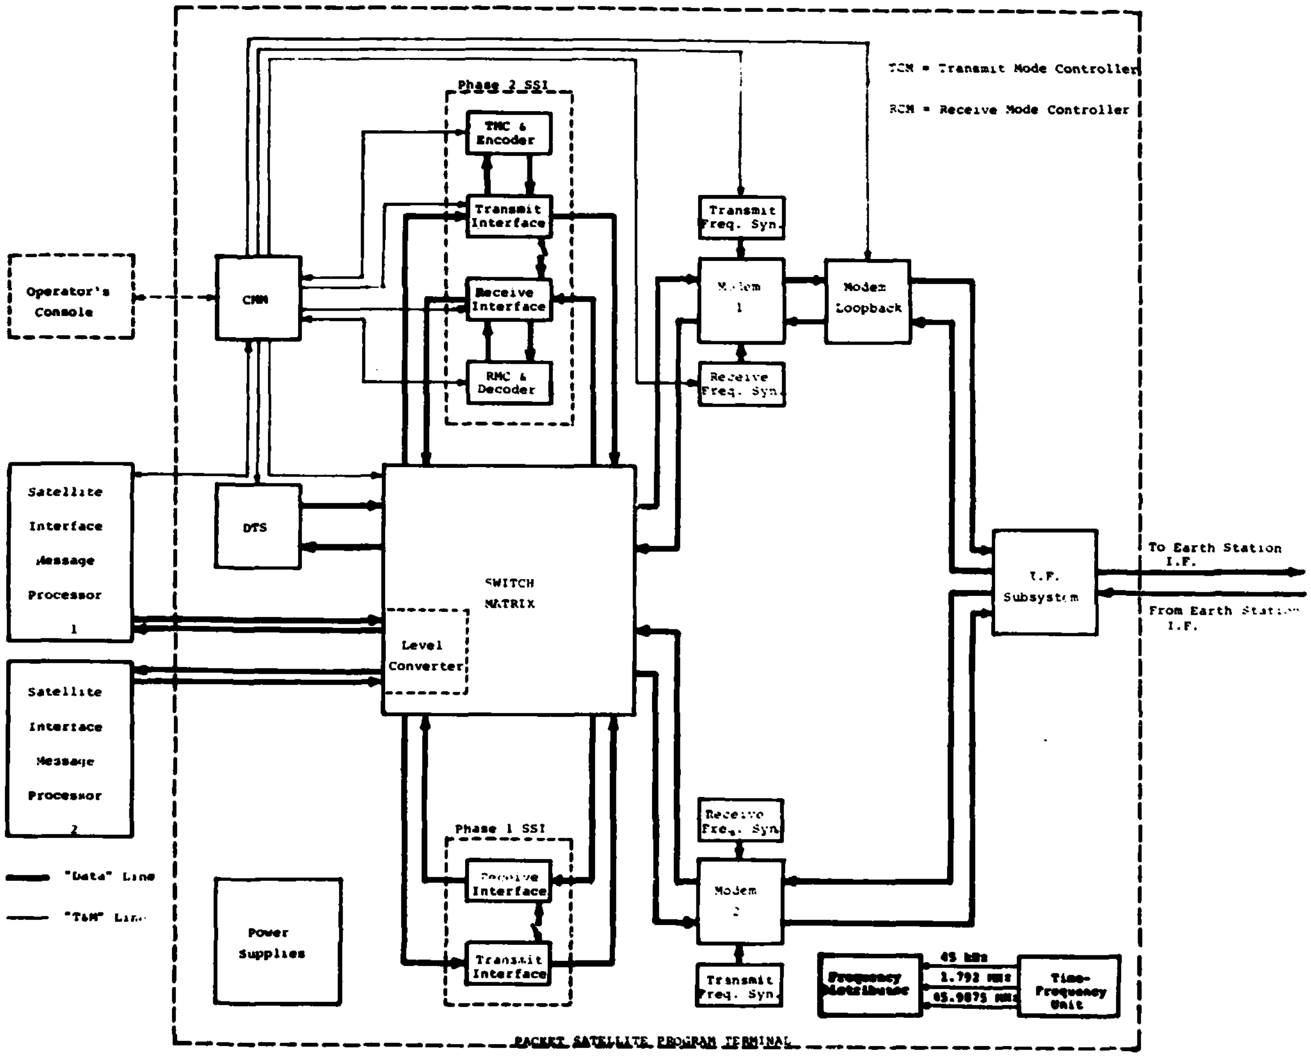

Figure 5-3. Functional Block Diagram of the Packet Satellite Program Terminal

The TFU generates the three reference clock signals (i.e., 45 kHz, 1.792 MHz, and 45.9875 MHz) used through-out the terminal. The terminal also has self contained power supplies.

The following subsections discuss the operation of the various subsystems of the PSP terminal.

5.2 OPERATION OF THE MODEM-SIMP INTERFACE

The interface between the SIMP and the new packet modem is denoted by SSI (SPADE-SIMP interface) in Figure 5-3. Although the SPADE modem will no longer be used, the historical acronym SSI, which is a familiar term to the DARPA PSP community, will be used for the modem-SIMP interface.

Both the Phase 1 and Phase 2 SSIs provide the data interface between the SIMP and the modem. The Phase 1 SSI supports only one packet transmission mode, i.e., a 64-kbit/s data transmission rate, and until now has provided the interface between the SIMP and SPADE modem in the three large stations. The Phase 2 SSI supports the packet transmission modes specified in Table 5-1. Definition of some of the parameters in this table and the basic format of a transmitted packet are shown in Table 5-2 and Figure 5-4, respectively. The data may be encoded using a rate 1/2 convolutional encoder. The decoding function, if required, is performed in the receive Phase 2 SSI, which also plays a major role in providing an overall test and monitoring capability for the satellite communications system. In particular, it can automatically transmit certain test and monitoring information to the SIMP at the end of each packet reception. This information includes all the measured modem performance parameters. All the available modem parameters are supplied by the modem to the SSI after each packet is received. Transmission of this data back to the SIMP is performed only on command from the SIMP.

In the following sub-paragraphs a brief general description of each module will be presented. A more detailed explaination of the operation of this individual modules and the system as a whole will be presented in paragraph 5.3.

Table 5-1. Phase 2 SSI Mode Chart

Mode	F/S	Coding	CL, PL	Pre	1st SOM	2nd SOM	1st Info Rate[a]	2nd Info Rate[a]	1st Info Rate[a]	2nd Info Rate[a]
1A	F	None	CL = 0	1	1	X	64	X	64Q	X
1B	S	None	CL = 0	2	3	X	64	X	64Q	X
2	S	All	CL = PL	2	2	X	16	X	32B	X
3A	F	All	CL = PL	1	1	X	32	X	64Q	X
3B	S	All	CL = PL	2	3	X	32	X	64Q	X
4	S	Part	CL < PL	2	2	1	16	64	32B	64Q
5A	F	Part	CL < PL	1	1	X	32	64	64Q	64Q
5B	S	Part	CL < PL	2	3	X	32	64	64Q	64Q
6A	S	All	CL + CL^1 =PL	2	2	1	16	32	32B	64Q
6B	S	All	CL + CL^1 =PL	2	2	3	16	32	33B	64Q
7	S	None	CL = 0	2	2	1[b]	64	X	64Q	X

[a]Rates are in kbit/s, Q = QPSK, B = BPSK.
[b]Second SOM is contiguous to first SOM.

Table 5-2. Parameter Definitions

PL ≡ Packet Length

CL ≡ Coded Length

CL^1 ≡ 2nd Coded Length for Mode 6, $CL^1 = PL - CL$

F/S ≡ Signal to Designate Preamble (sometimes referred to as L/S)
(F → transmission to large station)
(S → transmission to small station)

Pre ≡ Preamble
(1 ≡ large station preamble - nominal 60 baud times at 32 kilosymbol/s)
(2 ≡ small station preamble - nominal 96 baud times at 32 kilosymbol/s)

SOM Designation

Nominal Value

(1 ≡ large station QPSK SOM, 32 bits at 64 kbit/s)
(2 ≡ small station BPSK SOM, 32 bits at 32 kbit/s)
(3 ≡ small station QPSK SOM, 64 bits at 64 kbit/s)

Info Rate
Information rate in kbit/s

Chan Rate
Channel rate in kbit/s

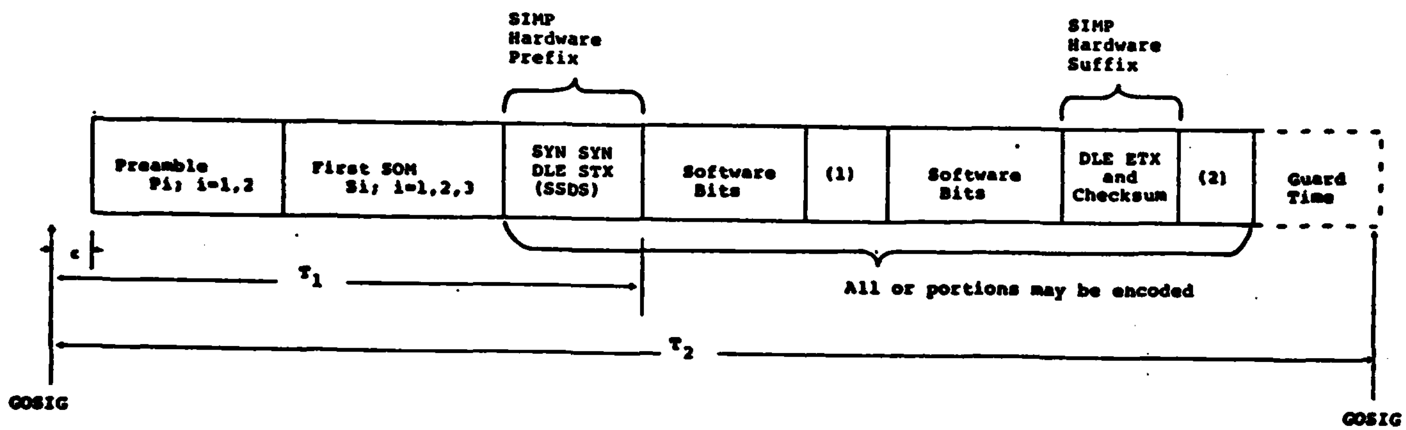

(1) Encoder tailing bits and second SOM for Mode 4
Encoder tailing bits for Mode 5
Second SOM for Mode 6

(2) Encoder tailing bits for Modes 2, 3, and 6

Note: Both fields (1) and (2) are generated by the transmit interface and will be removed by the receive interface. Thus the SIMP operation will be transparent to these fields as it is to the preamble and first SOM.

Figure 5-4. Basic Channel Packet Format

5.2.1 TRANSMIT SIDES OF THE PHASE 2 SSI

Figure 5-5 is a block diagram of the transmit side of the Phase 2 SSI. Data are clocked out into the SSI after a packet transmission is initiated by a GOSIG from the SIMP. The F/S (fast/slow) signal specifies the preamble to be used in the packet transmission. In general, transmission to large earth stations can utilize a shorter preamble and is designated by F; transmission to lower performance index (small) earth stations is designated by S. The transmit interface module provides carrier control; preamble, SOM, and clock generation; and data handling between the SIMP and the modem. The SOM and packet transmission mode selection are implemented by the transmission mode controller (TMC)/encoder module using the information supplied by the SIMP in the first two words of the packet header. This module also performs a 1/2 rate convolutional encoding, if required. The encoding is specified by information contained in the first two words of the packet header.

The transmit side of the Phase 2 SSI can perform numerous test and monitoring functions; some are prompted by the SIMP via the CMM, and others are performed routinely. Parameters such as preamble length, SOM length, and SOM pattern can be changed by commands from the CMM. Anomalies such as inconsistencies in packet header information and excessive duty cycles are detected automatically and reported to the SIMP by the CMM. In addition, the transmit side of the Phase 2 SSI is able to abort packet transmission if an error in the SIMP's unique word or anomalies in the first two words of packet header are detected and to transmit a dummy packet to all SIMPs to initiate rescheduling.

The Phase 2 transmit SSI also provides a data loopback and a codec bypass switch for testing. With the data loopback

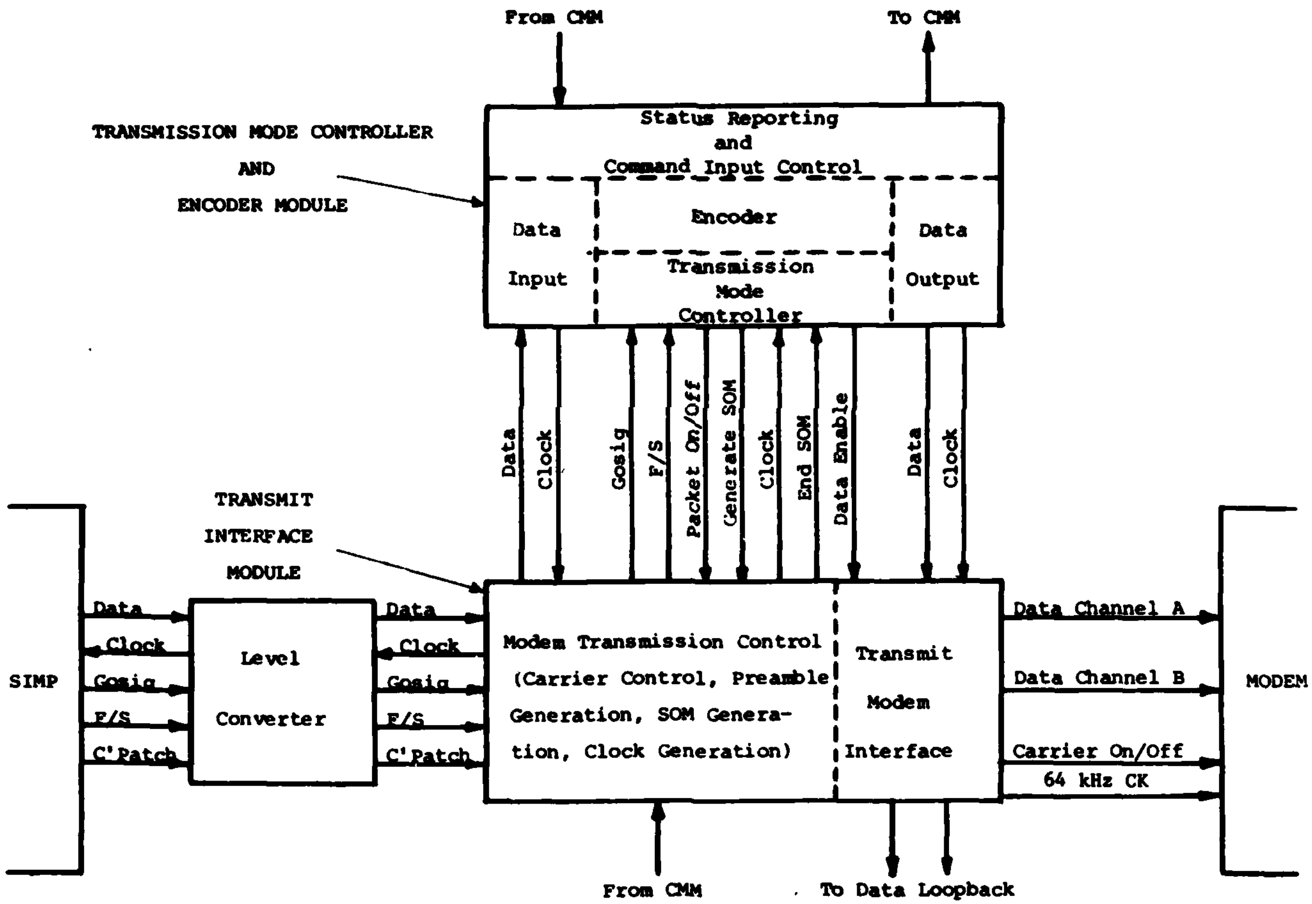

Figure 5-5. Functional Block Diagram of the Transmit Side of the Phase 2 SSI

switch, the transmit SSI can be connected directly to the receive SSI, hence bypassing the modem. The codec bypass switch allows the TMC/encoder and RMC (reception mode controller) decoder modules to be removed from the loop in the event of a failure of either unit, thus relying only on Phase 1 SSI capabilities.

5.2.2 TRANSMIT SIDE OF THE PHASE 1 SSI

The transmit side of the Phase 1 SSI does not provide any test and monitoring capabilities or the option of automatically modifying the preamble and SOM. It can be used only for packet transmission to large earth stations (INTELSAT Standard A) at 64 kbit/s. This module and the corresponding Phase 1 receive SSI have been operational for the past one and one-half years and will be used in the PSP terminal to provide redundancy for mode 1 (64 kbit/s) operation.

5.2.3 RECEIVE SIDE OF THE PHASE 2 SSI

A block diagram of the receive side of the Phase 2 SSI is shown in Figure 5-6. The modem supplies the receive interface module with bursts of data and clock from the demodulator. Data are clocked out simultaneously on the BPSK modem output lines (3 lines) and the QPSK modem output lines (3 lines for channel A and 3 lines for channel B). QPSK SOM detection and phase ambiguity resolution are performed in the receive interface module. The receive interface also receives a TTL preamble detect signal, which goes high and remains high until the modem receives an end of packet (EOP) signal. The SSI receive interface supplies the EOP signal to the modem to terminate packet reception or to

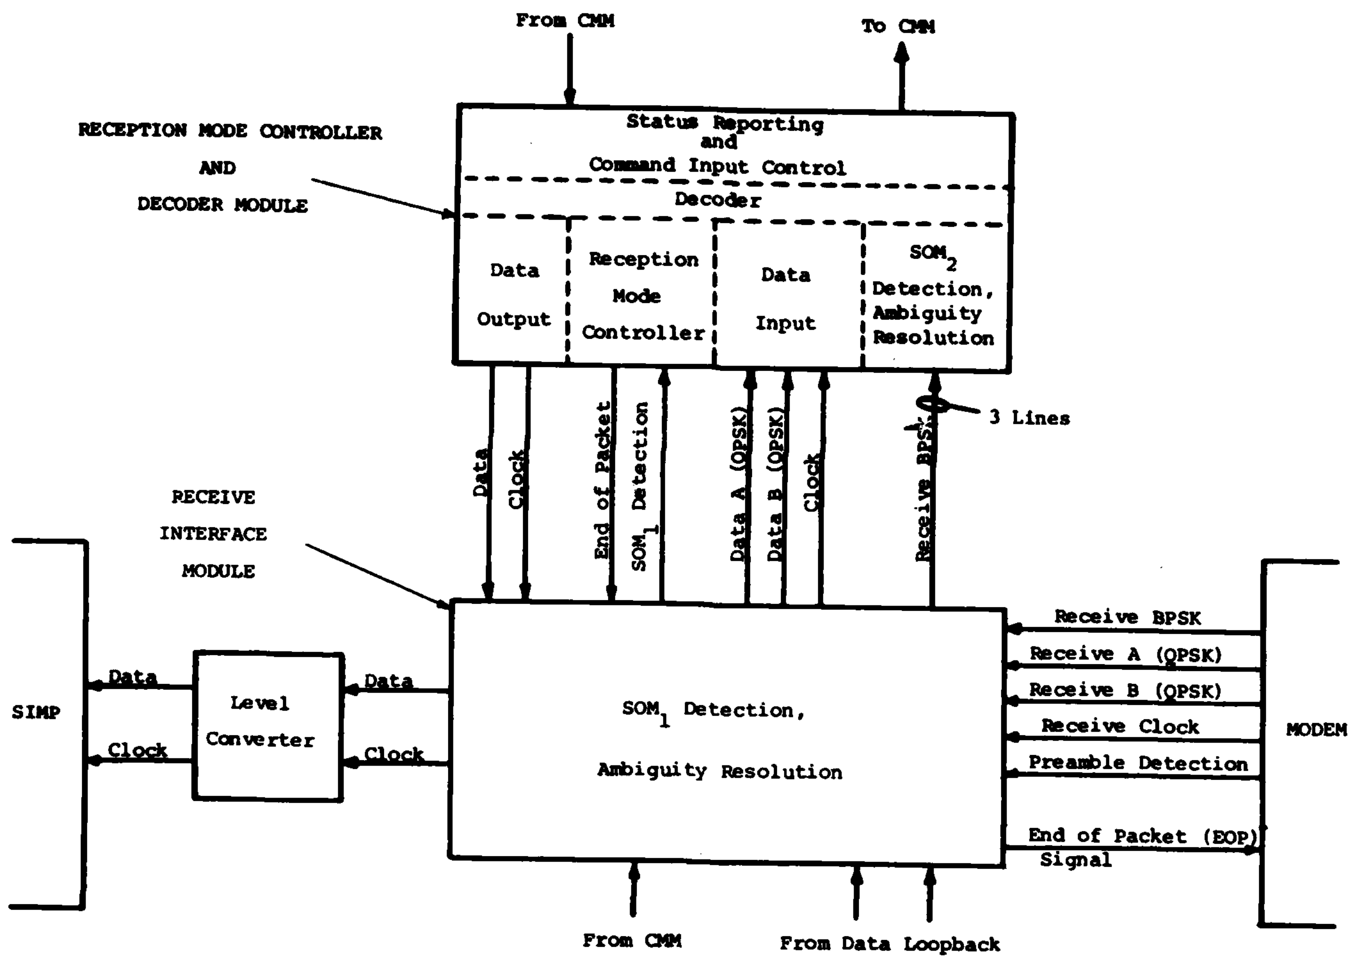

Figure 5-6. Functional Block Diagram of the Receive Side of the Phase 2 SSI

reinitialize the modem if a preamble but no SOM is detected. Each transmission of the EOP signal to the modem both initializes the modem and initiates the transfer of the test and monitoring parameters.

The reception mode controller (RMC)/decoder module selects the mode of operation of the receive SSI, decodes the data, if required, and clocks the data into the SIMP via the receive interface module. The output clock rate is 128 kbit/s. The output clock is intermittent (i.e., the data are outputted in spurts) during packet reception with the exact pattern depending on the packet transmission mode. Decoding is performed by a Module LV 7015C Viterbi decoder using soft decisions (i.e., 3-bit quanitiztion of the demodulator output). In addition, the RMC/decoder module performs BPSK SOM detection and ambiguity resolution.

The receive side of the Phase 2 SSI performs certain test and monitoring functions and stores the test and monitoring parameters passed to it from the modem. Based on commands received from the CMM, the RMC/decoder module transmits to the SIMP over the data path selected test and monitoring parameters, which are inserted into a 4-word slot at the end of the packet. Some of the parameters such as automatic gain control (AGC) on noise can be transmitted to the CMM upon command. The RMC/decoder module also tests the received packet; if inconsistencies appear in the header, the module will abort the packet and output to the SIMP a dummy packet, which will contain test and monitoring data.

5.2.4 RECEIVE SIDE OF THE PHASE 1 SSI

The receive side of the Phase 1 SSI, which has been in operation for the past one and one-half years, consists of the receive interface module without the CMM capabilities. This module performs QPSK SOM detection, ambiguity resolution, and parallel-to-serial conversion of the received data.

5.2.5 OPERATION OF THE MODEM

The modem used in the PSP terminal is a full-duplex phase-shift-keyed (PSK) modulator-demodulator for packet (burst) satellite communications. It has been designed to minimize the time required for burst mode acquisition in a low signal-to-noise ratio environment by utilizing the same preamble symbols to acquire both carrier and bit timing.

The modulator is a 4-phase 32 kilosymbol per second PSK modulator. Modulation of the in-phase and quadrature channels is controlled independently through two modulator inputs. Figure 5-7 shows a block diagram of the QPSK modulator. The modulator RF output center frequency is selected by an external frequency source, an SCPC frequency synthesizer that provides any one of 800 channels separated by 45 kHz over a 36-MHz band. The modulating data signals are filtered using 4-pole Butterworth low-pass filters with a 3-dB bandwidth of 18.2 kHz.

The demodulator is a 32 kilosymbol per second 2-phase (BPSK) and 4-phase (QPSK) demodulator. The BPSK capability is provided for reception by small stations. The demodulator frequency and bit timing are acquired independently for each new packet received. Each packet must start with an alternating zero-one sequence (48-96 symbols) and an SOM sequence for packet

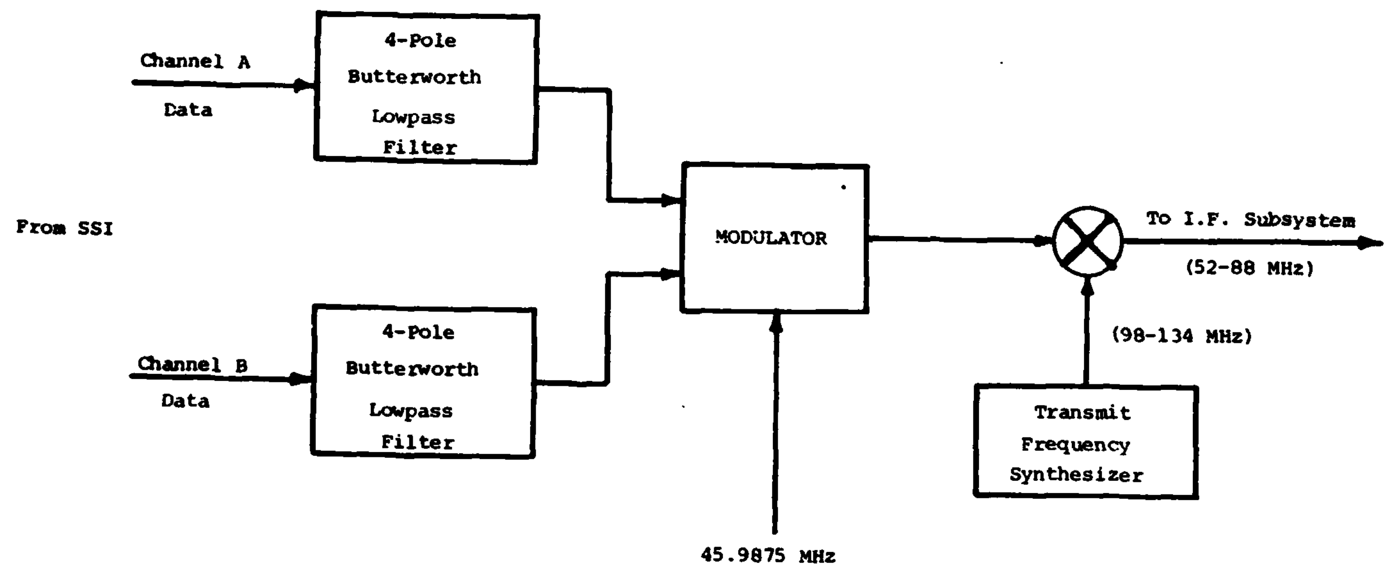

Figure 5-7. Block Diagram of the Modulator

framing. When the demodulator is not actively demodulating a packet, it is continuously in the acquisition mode, searching for a preamble in the received data.

The demodulator, which is shown in Figure 5-8, consists of an analog front end and a digital microprocessor. Most of the demodulator functions are performed in the microprocessor, which has been specially designed for digital signal processing. The analog portion has a dual conversion front end. The first conversion is accomplished by mixing the incoming signal with an external frequency source, which (similar to the modulator) may be an SCPC synthesizer to select any one of 800 channels in a 36-MHz band. The second conversion lowers the signal to baseband (±2 kHz) yielding both in-phase and quadrature signals. The ±2 kHz residual frequency offset results from the uncertainty in the received carrier frequency. At this point, both the in-phase and quadrature channels are filtered using 7-pole Butterworth filters with delay equalization. These filters have a 3-dB bandwidth of 18.2 kHz. Each channel is analog-to-digital (A/D) converted at a 64-kilosample/s rate (two samples per received symbol per channel), and each of these samples is quantized to 4 bits. These samples are fed to a microprocessor which performs the remaining demodulator functions; namely, frequency and phase recovery, bit timing recovery, and data demodulation. The demodulated packet data are fed into the receive SSI until the modem receives an EOP signal, which ends the demodulation process and returns the modem to the preamble search mode.

The modem can continuously monitor certain critical parameters of the communications link and make these results available to the receive SSI at the end of each packet. Data available in the modem include the following: AGC voltage with noise alone, AGC voltage with signal and noise in the receiver, frequency offset of the incoming signal during preamble detection,

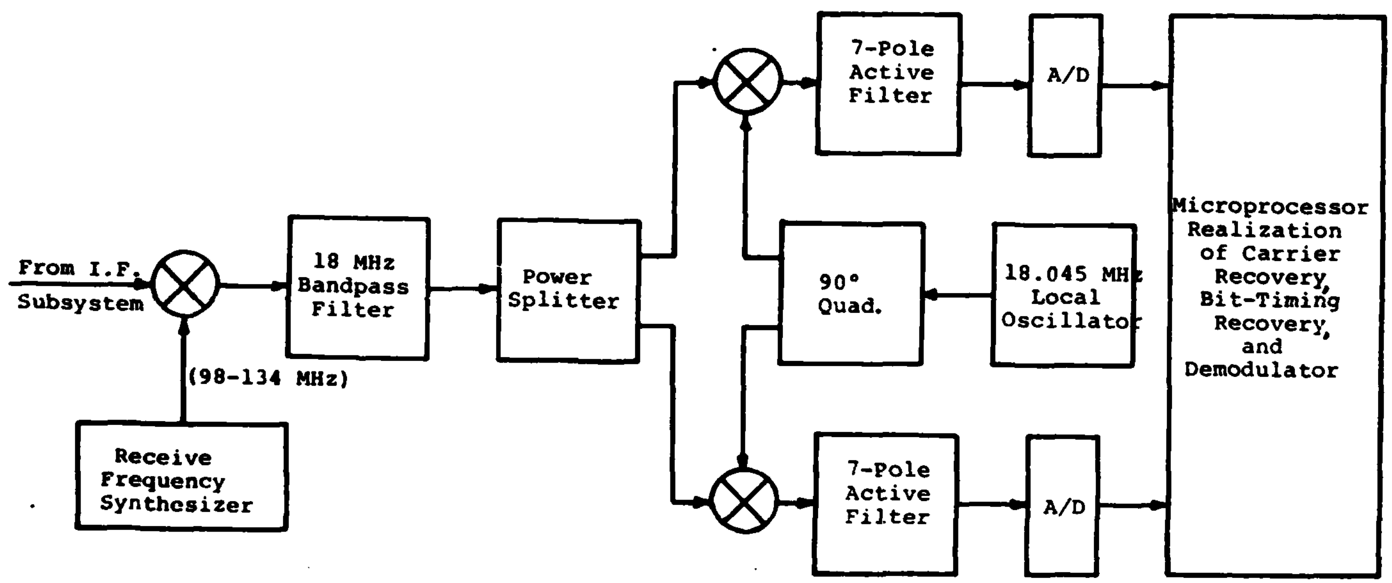

Figure 5-8. Block Diagram of the Demodulator

QPSK frequency offset at the end of the packet, and information for the calculation of the effective value of E_b/N_o both at the beginning and at the middle of the packet. Upon command from the CMM to the receive SSI, these parameters are appended to the packet and transmitted to the SIMP. The AGC voltage with noise alone is always available and can be sent to the SIMP via the CMM path. Furthermore, one of 800 possible frequencies of the transmit/receive frequency synthesizers can be selected by a command from the CMM. This command has been implemented in the CMM software, but disabled in the transmit and receive synthesizers.

5.2.6 MODEM LOOPBACK AND IF SUBSYSTEM

The modem loopback and the IF subsystem are implemented in a single module (board) which consists of a switch, a frequency translator, a power combiner, and a power divider. With this module, the modem output can be directly connected to its input (via a frequency translator), and the two channels (SSIs and modems) can operate simultaneously, hence using the full-duplex capability of the PSP terminal without reconfiguring the earth station interface. The first capability, i.e., the modem loop-back, facilitates SSI and modem testing and troubleshooting.

5.2.7 DATA TEST SET

The DTS has been implemented to simulate the SIMP as a packet source and receiver. It has been designed to assist earth station personnel in channel troubleshooting and has been proven essential in the testing of the Phase 2 SSI. The DTS packet format is similar to that generated by the SIMP. Figure 5-9 shows the front panel of the test set. The DTS transmission parameters which are F/S, packet length, coded length, mode of transmission, packet rate, and interpacket interval as well as the data pattern (all zeros, or all ones, or pseudorandom data) can be set manually. The on/off and ready/clear switches can be operated either manually or remotely via the CMM. The receive side can indicate the number of packets with bit errors, total bit errors, the number of SYN DLE STX and DLE ETX sequences detected, received packet mode, packet length, and coded length. Values of most of these parameters can be readily compared with the corresponding transmission parameters.

5.2.8 SWITCH MATRIX

With the switch matrix, the PSP terminal can be operated in different modes. A DTS, a SIMP, a DTS and a SIMP, or two SIMPs can be chosen as packet source/sink devices. The full duplex capability of the PSP terminal can be exercised with either two independent SIMPs or a SIMP and a DTS connected to the 2-channel equipment in the terminal. Furthermore, the switch matrix enables either Phase 1 or Phase 2 SSI to be used with all four transmit/ receive combinations of the two modems (i.e., modem 1 transmit and modem 1 receive, modem 2 transmit and modem 1 receive, modem 1 transmit and modem 2 receiver, or modem 2 transmit and modem 2 receive). The different switches on the switch matrix can be set either manually or via command from the CMM.

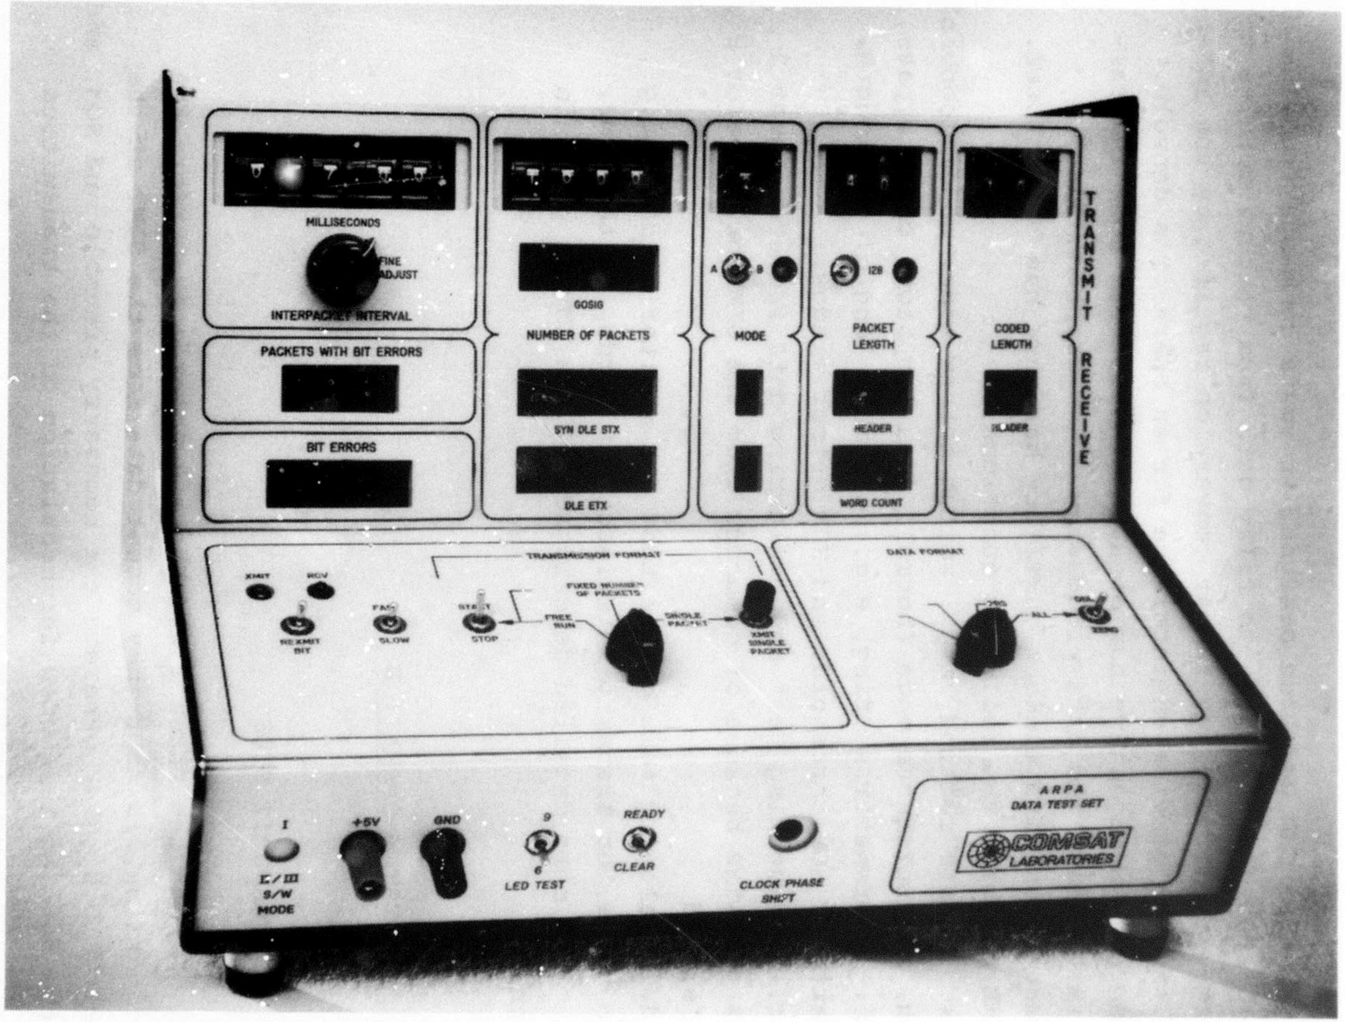

Figure 5-9. Front Panel of the Data Test Set

5.2.9 COMMAND AND MONITORING MODULE

The CMM has been designed to route commands from the SIMP to specified devices, to relay data from various modules to the SIMP, to provide any format conversion required by the interfaces, and to provide an interface for an operator's console. The CMM design objective has been to minimize its role in real-time involvement and data processing. A microprocessor implementation with minimum auxiliary hardware has been chosen. The CMM performs the following functions:

a. state setting, e.g., switch matrix mode, modem loopback switch, and interface mode (normal, data loopback, codec bypass);

b. parameter setting, e.g., preamble length, SOM length, SOM pattern, SOM threshold, and SOM detect window;

c. requesting test and monitoring data from the Phase 2 SSI or the SIMP because of commands from the SIMP or operator's console; and

d. relaying some of the test and monitoring data, e.g., AGC on noise, or transmit interface trap words, to the SIMP.

A more detailed description of the CMM is given in Section 5.4.

5.3 SYSTEM OPERATION

5.3.1 TRANSMIT SIDE

Both the Phase 1 and Phase 2 interfaces perform the same functions, in that they accept data from the SIMP at the transmit clock rate, generate the necessary preamble and SOM, and pass the data to the modem for modulation of the transmitted carrier. Carrier on/off control is also performed by the transmit interface. The major difference between the two interface units is that the phase 1 system will generate only QPSK - uncoded data, whereas the phase 2 system can generate QPSK-uncoded and/or coded data, and BPSK-coded data. The particular form of modulation in the phase 2 interface is determined by the information contained in the first two software header words generated by the SIMP and the F/S level, also generated by the SIMP.

The following sequence of events takes place once the GOSIG is activated in the SIMP.

a. GOSIG goes high - this signal to the interface indicates that the SIMP is ready to transmit data.

b. Packet ON/OFF is generated from GOSIG and passed to the TMC/encoder. The Packet ON/OFF causes the interface to generate the BPSK or QPSK preamble and SOM (based on level of F/S signal). The modem carrier is turned on by Packet ON/OFF signal.

c. TMC/encoder clocks out the first 64 bits of data from the SIMP, at a bit rate of 128 kbit/s, and from the header data determines if encoding is required. The data are stored in buffer registers in the TMC/encoder board.

d. After the SOM has been generated, an END SOM pulse is sent to the TMC/encoder requesting the data for transmission.

e. Data at 64 kbit/s are transfered from the TMC/encoder to the Transmit Interface, where it is transformed from serial to parallel data at 32 k symbols/sec.

f. When all data have been transmitted to the modem, the carrier is turned off by the Packet ON/OFF signal, which is controlled by a signal from the TMC/encoder.

For transmission to a Standard station, values of carrier recovery and bit timing recovery can be varied from 2 bits/channel to 255 bits/channel each, with typical values of 2 and 56 bits/channel, respectively. The length of the SOM is 16 bits in both channel A and channel B data lines to the Modem. Typical timing diagrams and packet format are shown in Figure 5-10.

For transmission of a Mode 2 type packet, i.e., transmission to a non-standard station, the same sequence of events takes place, except that the long preamble of 96 bits/channel of only bit timing recovery is generated (no carrier recovery), followed by the BPSK SOM. Data coming from the TMC/encoder has encoded the SIMP data at a rate of 1/2 before it is sent to the transmit interface. The TMC/encoder transmits the data to the transmit interface at a 32 kbit/s rate (i.e., each bit is transmitted twice at 64 kbit/s) which results in BPSK modulation of the transmitted carrier.

The general format for all the modes is shown in Figures 5-11(a) through 5-11(f). These figures include the type of preamble used, the SOM's if more than one is required, what portion of the packet is encoded (if required), and the location of the encoder trailing bits. The second SOM is generated for those packets where both BPSK and QPSK modulation is performed by the modem. The second SOM is required by the receive interface to resolve ambiguities of the demodulated QPSK data. The first SOM is used to resolve the phase ambiguity in the BPSK demodulated data. It should be noted that in those modes where both BPSK and QPSK modulation are required, BPSK is transmitted first (see Modes 4 and 6).

5.3.2 RECEIVE OPERATION

The data received from the modem can consist of either demodulated QPSK or BPSK data. The receive interface and RMC/decoder must determine which types of data has been received and

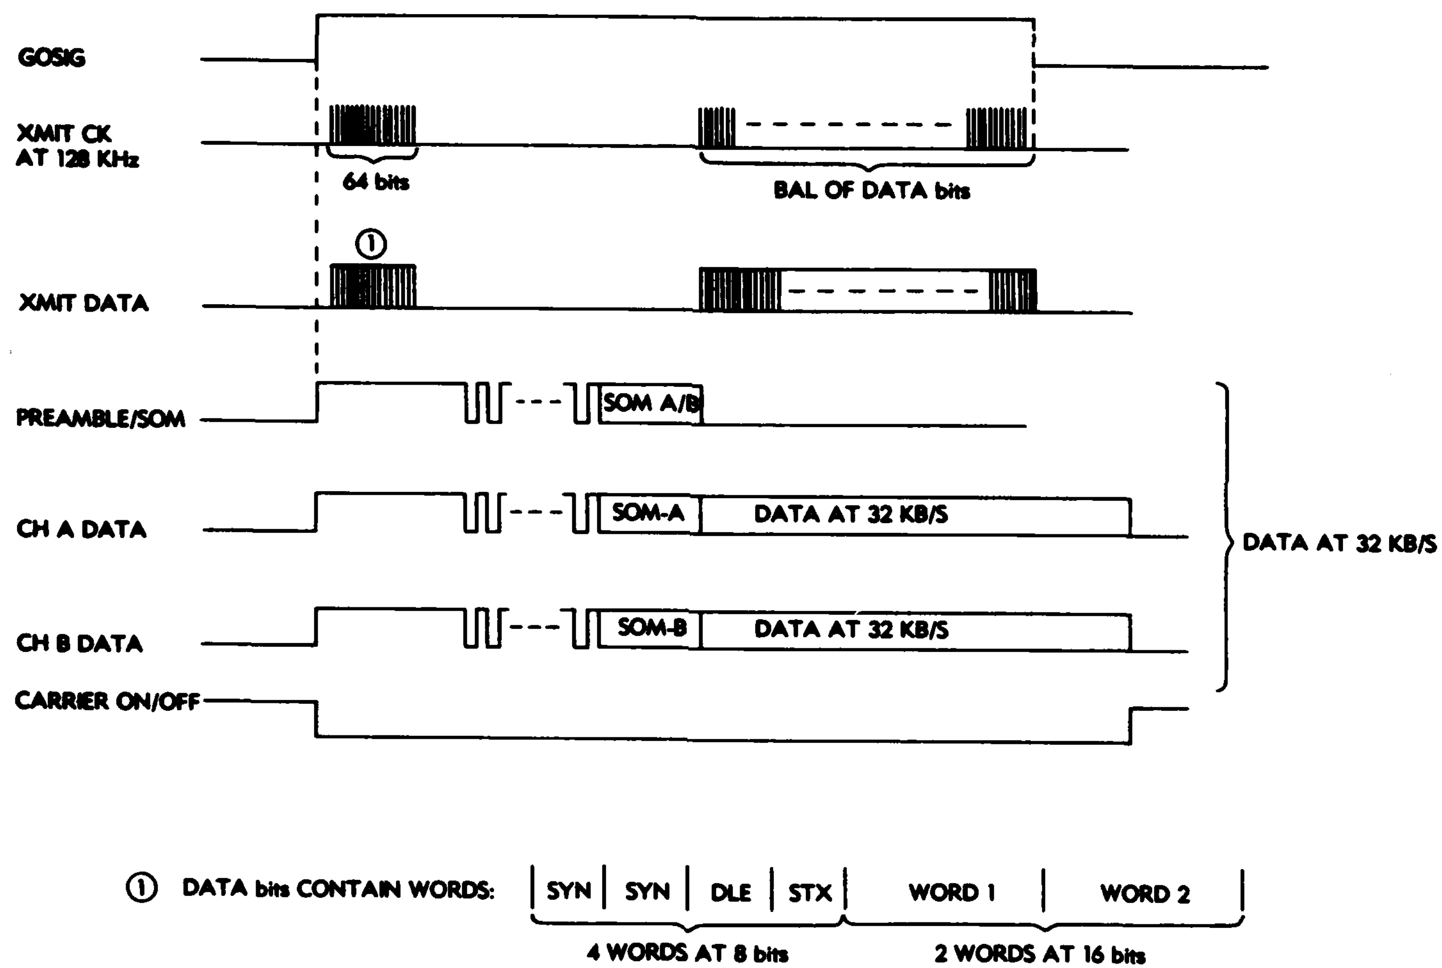

Figure 5-10. Timing Diagram and Packet Format

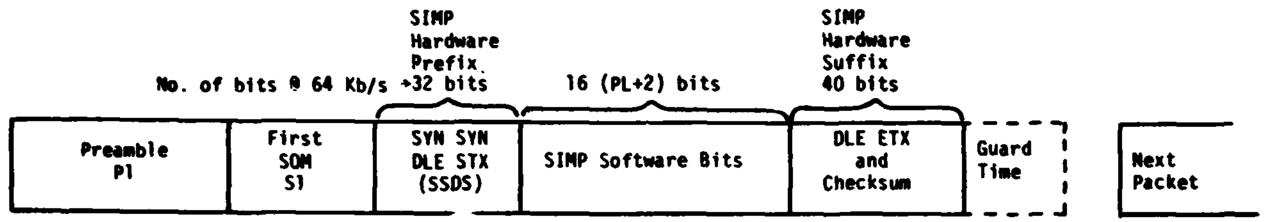

Figure 5-11a. Mode 1 Information Rate 64 kbit/s, Channel Rate 64 kbit/s

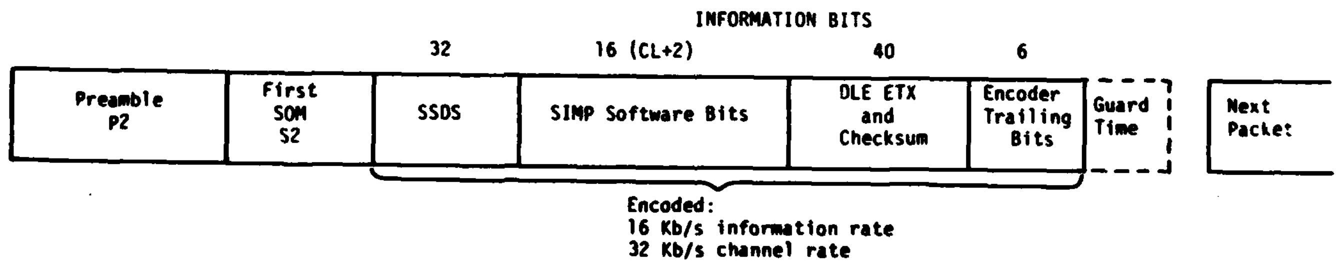

Figure 5-11b. Mode 2 Information Rate 16 kbit/s, Channel Rate 32 kbit/s

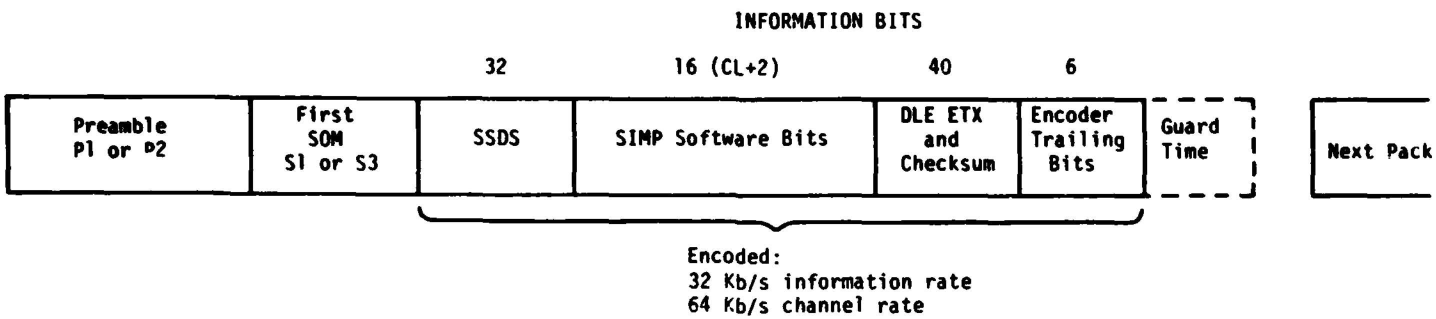

Figure 5-11c. Mode 3 Information Rate 32 kbit/s, Channel Rate 64 kbit/s

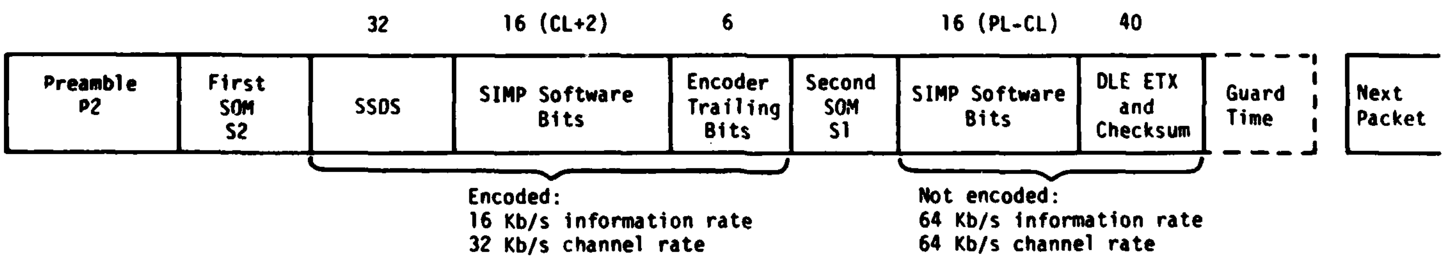

Figure 5-11d. Mode 4 Information Rate 16 kbit/s, Channel Rate 32/64 kbit/s

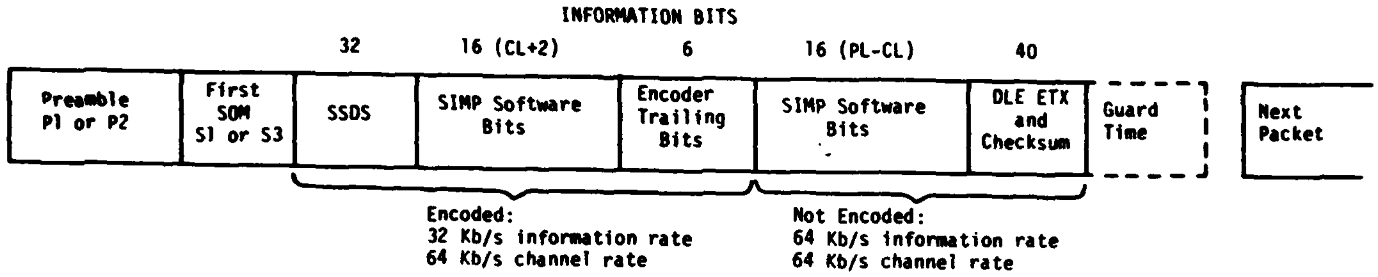

Figure 5-11e. Mode 5 Information Rate 32/64 kbit/s, Channel Rate 64/64 kbit/s

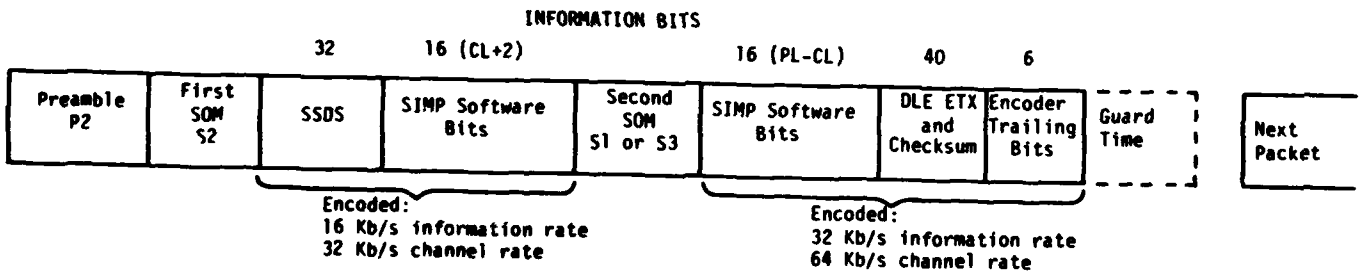

Figure 5-11f. Mode 6 Information Rate 16/32 kbit/s, Channel Rate 32/64 kbit/s

perform the proper functions in order to restore the data to its original bit pattern before it is sent to the SIMP. Each packet upon arrival can be any of the ten different modes shown in the mode chart of Table 5-1. In order to determine which mode is received, the following sequence of events takes place.

a. Upon detection of the preamble (carrier recovery and bit timing recovery) by the modem, a Preamble Detect signal is sent to the receive interface and RMC/decoder modules. This commands the modules to look for both the QPSK and BPSK SOM.

b. Demodulated data is sent to the receive interface and RMC/decoder. The receive interface searches for the QPSK SOM, the RMC/decoder searches for the BPSK SOM.

For a Mode 1 type packet:

c. The receive interface searches for the QPSK SOM. The data received by the interface will be in one of eight possible states, listed in Table 5-3.

d. The receive interface determines which of the eight states exists, resolves the ambiguities, and then pass the data to the RMC/Decoder for further processing. The unambiguous state is defined as State 1.

e. A "QSOM Detect" signal is passed to the RMC/Decoder indicating that it is now receiving valid data.

f. The RMC/Decoder looks for a valid bit sequence of SYN DLE STX (ASCII Code) and then inspects the next 32 data bits to determine packet length, packet mode, and coded length. The latter should be zero for a Mode 1 packet.

Table 5-3. Ambiguity States*

State	I-Ch.	Q-Ch.
1	SOM-A	SOM-B
2	SOM-A	$\overline{\text{SOM-B}}$
3	$\overline{\text{SOM-A}}$	SOM-B
4	$\overline{\text{SOM-A}}$	$\overline{\text{SOM-B}}$
5	SOM-B	SOM-A
6	SOM-B	$\overline{\text{SOM-A}}$
7	$\overline{\text{SOM-B}}$	SOM-A
8	$\overline{\text{SOM-B}}$	$\overline{\text{SOM-A}}$

*Note "$\overline{}$" indicates inversion

g. If the inspection of packet length, coded length, and mode number show valid information, i.e., coded length equals zero, then the RMC/decoder transforms the data from parallel to serial form and passes it on to the SIMP. The data and clock coming from the RMC/decoder is at a "bursty" 128 KHz rate, i.e., data and clock are transmitted three or four bits at a time.

For a Mode 2 packet.

a. BPSK data is read into the RMC/Decoder which scans the data for the BPSK SOM.

b. Upon detection of the SOM, a determination is made as to which of two states the data arrived, viz, inverted or not inverted, and inverts the data if required.

c. After detection of the BPSK SOM, the Decoder starts decoding the data, using the soft decision Viterbi decoding algorithm.

d. The RMC/decoder looks for a valid SYN DLE STX and then inspects the next 32 data bits to determine packet length, coded length (= to packet length) and mode.

e. Decoded data is then sent to the SIMP at the "bursty" 128 KHz clock rate.

For mixed tranmission modes, e.g., Mode 4, the above sequences are both employed as follows:

a. The RMC/decoder reads in the BPSK data and searches for the BPSK-SOM, while the receive interface searches for the QPSK-SOM.

b. Upon detection of the BPSK-SOM by the RMC/decoder, it resolves any phase ambiguity in the BPSK data, sends it through the decoder, and from this data, determines the packet length, mode (in this case = 4), and coded length. Data are then sent to the SIMP at the "bursty" 128 KHz rate.

c. The receive interface continues looking for a QPSK SOM as long as the Packet Detect Signal from the modem is high. Upon detection of the QPSK SOM, the QSOM signal is sent to the RMC/Decoder along with the resolved QPSK data.

d. The RMC/decoder converts the parallel data from the receive interface to serial data, and sends it to the SIMP at the "bursty" 128 KHz clock rate. The RMC/decoder must insure the fact that the last BPSK datum bit is followed immediately by the first QPSK datum bit, i.e., no bits are dropped nor added between the coded and uncoded portions of the packet.

At the end of each data packet the RMC/decoder generates an "END OF PACKET" pulse. This pulse is sent to the modem to reset the microprocessor in the modem. This pulse causes the modem to start searching for a new packet, which it recognizes by the detection of the Bit Timing Recovery sequence at the beginning of the new packet.

5.3.3 SWITCH MATRIX OPERATION

The switch maxtrix controls the channel unit data I/O by determining which channel unit, phase 1 or phase 2, the SIMP will be transmitting data to, and receiving data from, and which of the two modems will be used for transmission and reception. Figure 5-12 shows functionally how this is done. As shown in the figure, the switches are connected for a Switch Matrix Mode 1 operation. In this mode, the SIMP is sending and receiving data from the Phase 2 interface. The interfaces are transmitting and receiving data to/from the "B" channel unit. The Data Test Set will operate through the Phase 1 interface and the "A" channel modem.

Similarly for S/M Mode 2, the SIMP transmits and receives through the Phase 2 interface, but transmits through the "A" channel modem and receives through the "B" channel modem. The remaining six operational modes are also shown in the table in Figure 5-12.

Two test modes are included. These are modes "zero" and "nine". In these two modes, either the SIMP or the DTS can be used as a data source, with the same receive data going to both the SIMP and DTS. In mode zero, the SIMP transmits through the Phase 1 interface and modem "A". Or, in the same mode, the DTS could be the data source transmitting through the Phase 2 interfaces. In mode nine, the opposite would be true for data source, whereas both SIMP and DTS receive the data regardless of which unit transmits.

Switch "S5" is used to connect a second SIMP input to the PSP Terminal. In this case, the DTS could not operate with the terminal.

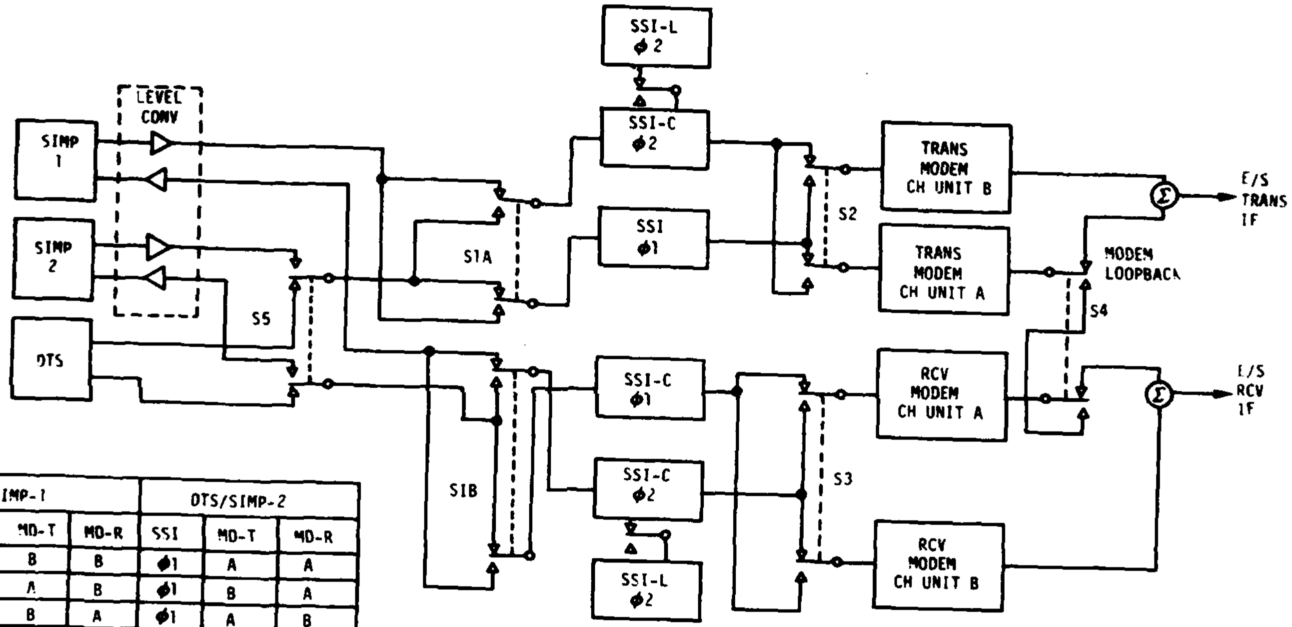

S/M MODE	SIMP-1			DTS/SIMP-2		
	SSI	MD-T	MD-R	SSI	MD-T	MD-R
1	ϕ2	B	B	ϕ1	A	A
2	ϕ2	A	B	ϕ1	B	A
3	ϕ2	B	A	ϕ1	A	B
4	ϕ2	A	A	ϕ1	B	B
5	ϕ1	B	B	ϕ2	A	A
6	ϕ1	A	B	ϕ2	B	A
7	ϕ1	B	A	ϕ2	A	B
8	ϕ1	A	A	ϕ2	B	B

	SIMP XMIT	DTS XMIT
0	CH UNIT A	CH UNIT B
9	CH UNIT B	CH UNIT A

NOT SHOWN: PATH LOOPBACK FROM XMIT SSI-ϕX TO RCVE SSI-ϕX

Figure 5-12. Switch Matrix Configuration - Mode 1

5.3.4 IF SUBSYSTEM

The IF subsystem in the PSP Terminal is used only to combine the 52-88 MHz transmit IF signals from the two modems, and divide the 98-134 MHz receive IF into the terminal. There is no AFC or AGC capability in the IF subsystem. These functions are controlled by the earth station GCE.

The IF subsystem does have the capability to do modem IF loopback testing with the Channel "A" modem, when transmitting in the 52-70 MHz band (LO Band). This test capability can be used to test and/or repair a faulty modem. The IF loopback test does not interrupt the operation of the Channel "B" modem.

5.3.5 TIMING/FREQUENCY UNIT (TFU)

The PSP terminal has a TFU which generates all the necessary timing signals used throughout the terminal. They include oscillators at;

a. 45.000 KHz
b. 1.792 MHz
c. 45.9875 MHz

The 45.000 KHz oscillator is a temperature controlled unit. Its output is used in the transmit and receive synthesizers. The 1.792 MHz oscillator is used to generate the transmit clocks (64 KHz and 32 KHz) in the transmit interface. It is also used in the receive interface for SOM detection. The 45.9875 MHz oscillator is used in the IF subsystem for frequency translation in the modem loopback test.

5.3.6 POWER SUPPLIES

The PSP terminal has self contained power supplies. All the necessary voltages used in the terminal are supplied by individual, commercially purchased units.

5.4 COMMAND AND MONITORING MODULE

The Command and Monitoring Module (CMM) is a relatively recent addition to the PSP terminal which provides two basic functions:

1. the CMM receives and interprets commands from the SIMP and then relays these commands to the various modules of the PSP terminal to modify transmission formats or change the configurations of the PSP terminal hardware.
2. the CMM will, again in response to a command from the SIMP, collect certain test and monitoring parameters from the modules in the PSP terminal, and then transfer these parameters to the SIMP.

The various command and monitoring data paths between the CMM and the PSP terminal subassemblies are shown in Figure 5-13. A more detailed description of these interfaces can be found in Appendix A.

There are several SIMP initiated commands that the CMM can carry out which change the configuration of the PSP terminal. The CMM has two initialization schemes which cause either a full or partial initialization of the PSP terminal. RESTART of the PSP terminal is a full initialization which occurs following a hardware reset of the CMM due to power-up or operator control, or

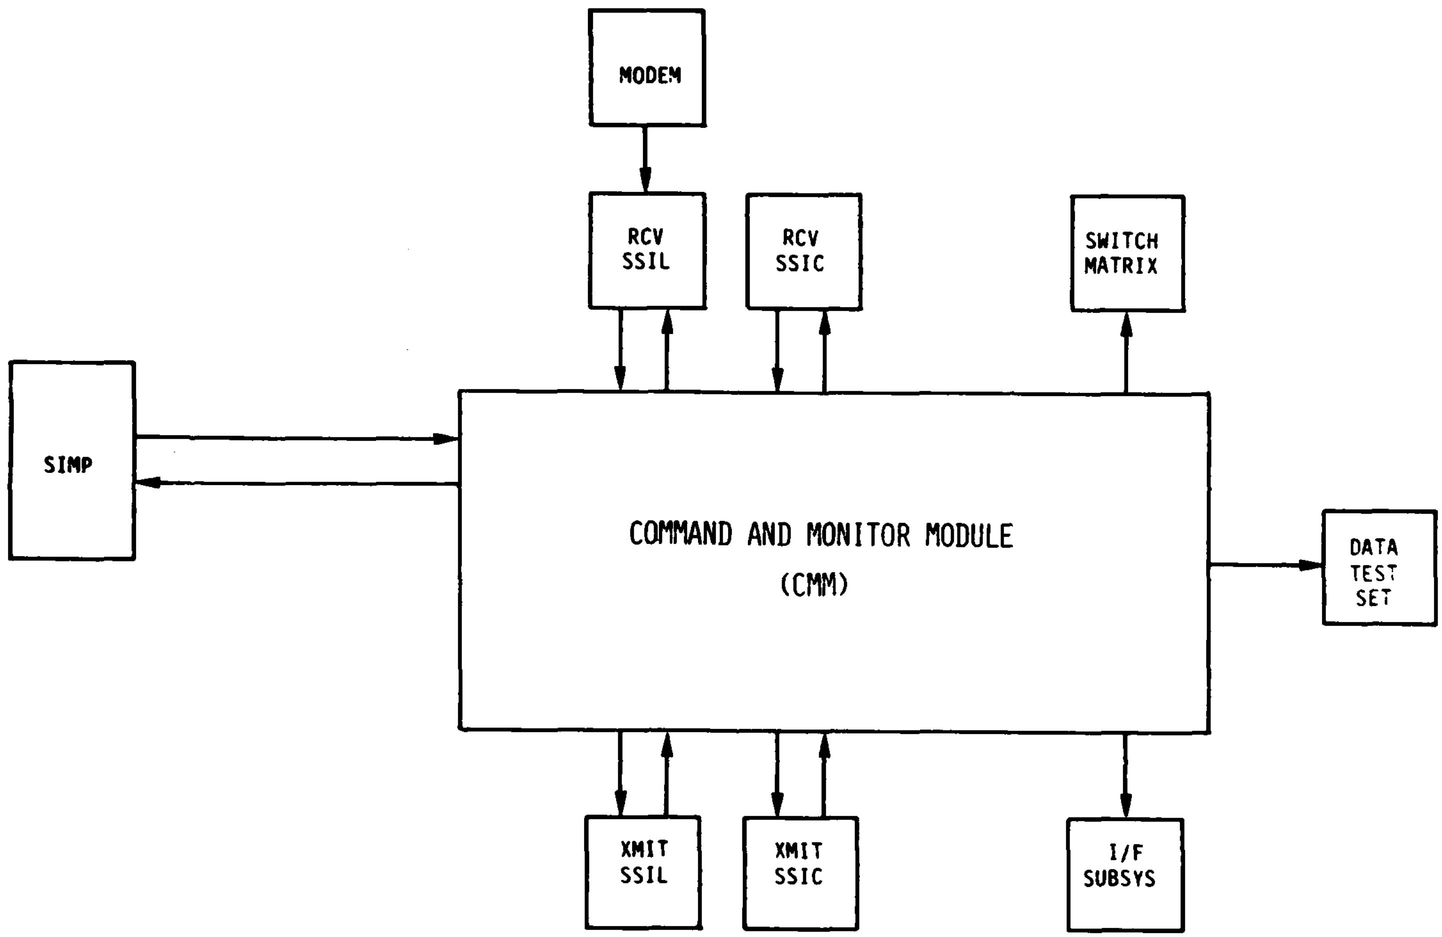

Figure 5-13. Basic Information Between CMM and Other Subassemblies in the PSP Terminal

by SIMP command. This initialization causes the CMM to initialize all of its internal registers, set the PSP terminal to default parameters and mode, and reset the other PSP modules. The partial initialization is refered to as RESET, and is a SIMP initiated command. RESET sets a default mode, SOM pattern and length, and resets the PSP terminal modems and interfaces. CMM internal registers remain unchanged following a RESET command from the SIMP.

The CMM also allows the SIMP to change the hardware configuration of the PSP terminal through the CMM interfaces to the Switch Matrix, Data Test Set, and other PSP modules. By commanding the CMM to change the Switch Matrix, all of the PSP terminal modes can be software controlled and changed by the SIMP. The CMM allows control of four DTS functions, specifically: setting DTS to FAST or SLOW transmission rate, setting to single packet transmission mode, fixed number of packets mode, or free running mode. When the fixed number of packets mode is selected, the actual number of packets is set manually on the DTS front panel.

Another important function of the CMM is to allow the SIMP to change various packet parameters. The SIMP can change the carrier recovery length, the small and large station bit timing recovery length, the SOM patterns and lengths, the allowable SOM correlation error threshold (2 to 5 errors) and control of test and monitoring data and dummy packets from the receive interfaces. In all cases the parameters changed by the SIMP commands to the CMM remain until a new change is specified or the CMM is initialized.

An important function of the CMM is the monitoring and collection of various PSP terminal parameters. The CMM constantly monitors (packet by packet) the number of SOM correlation errors, and stores the most recent value in an internal register. By a simple command to the CMM, the SIMP can obtain this value which can be stored or monitored by SIMP software to help indicate channel and PSP terminal conditions. Other system parameters are not continually monitored, but are gathered and stored by the CMM following a SIMP command to do so. These parameters include: preamble parameters, SOM patterns and lengths, and AGC values from the receive modems. Finally, the SIMP can monitor all of the 48 internal registers of the CMM, and can initiate a CMM routine used to check the status of the SIMP-CMM interface.

More detail of the actual CMM functions and the interfaces shown in Figure 5-14 are given in Appendix A.

5.5 MODEM TEST DATA

Tests have been performed on the six microprocessor-based modems that were developed by the Linkabit Corporation for inclusion in the delivered PSP terminals. These data give bit error rate versus E_b/N_o for operation with mode 1 packets (uncoded QPSK). Nominal power levels (-30 dBm) were supplied to the modem for all tests. These results are summarized in Figures 5-15 through 5-17. The test results can be summarized as follows:

Figure	Modem S/N#	PSP Terminal	Departure From Theory at 10^{-6} BER
5-15	4	#1 (ETAM)	2.8 dB
5-15	6	#1 (ETAM)	3.2 dB
5-16	3	#2 (Goonhilly)	2.1 dB
5-16	5	#2 (Goonhilly)	2.8 dB
5-17	1	#3 (Tanum)	2.2 dB
5-17	2	#3 (Tanum)	2.3 dB

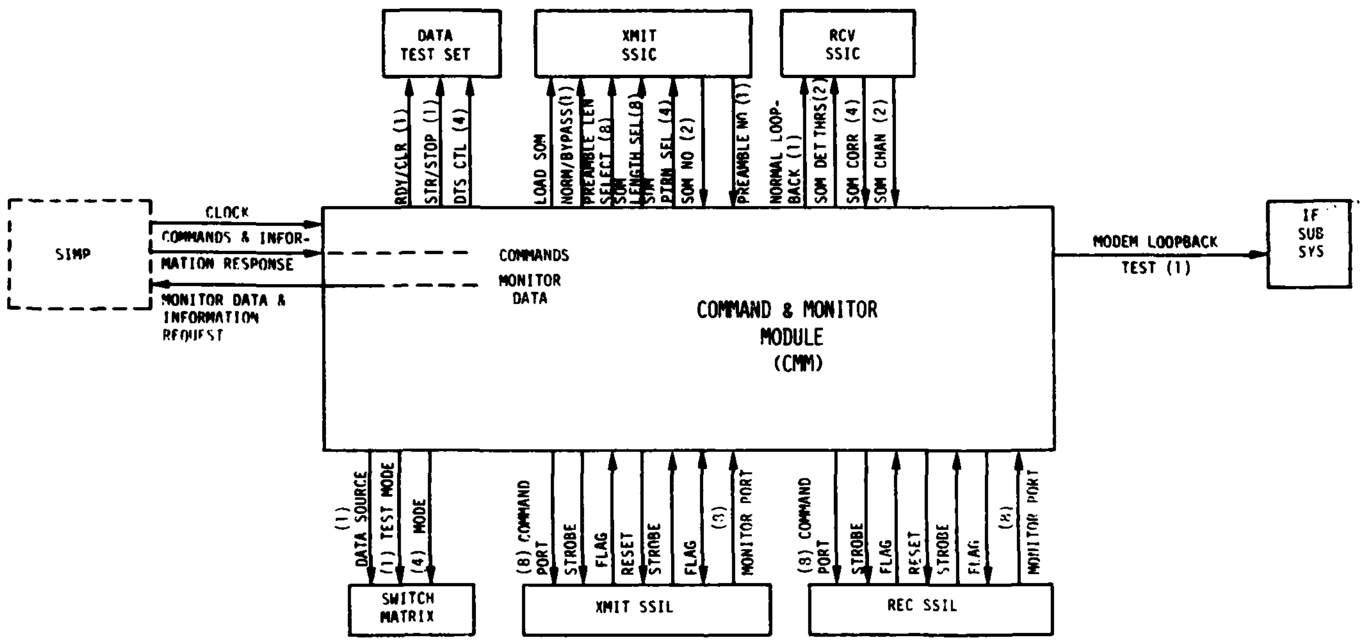

Figure 5-14. CMM Functional Connections

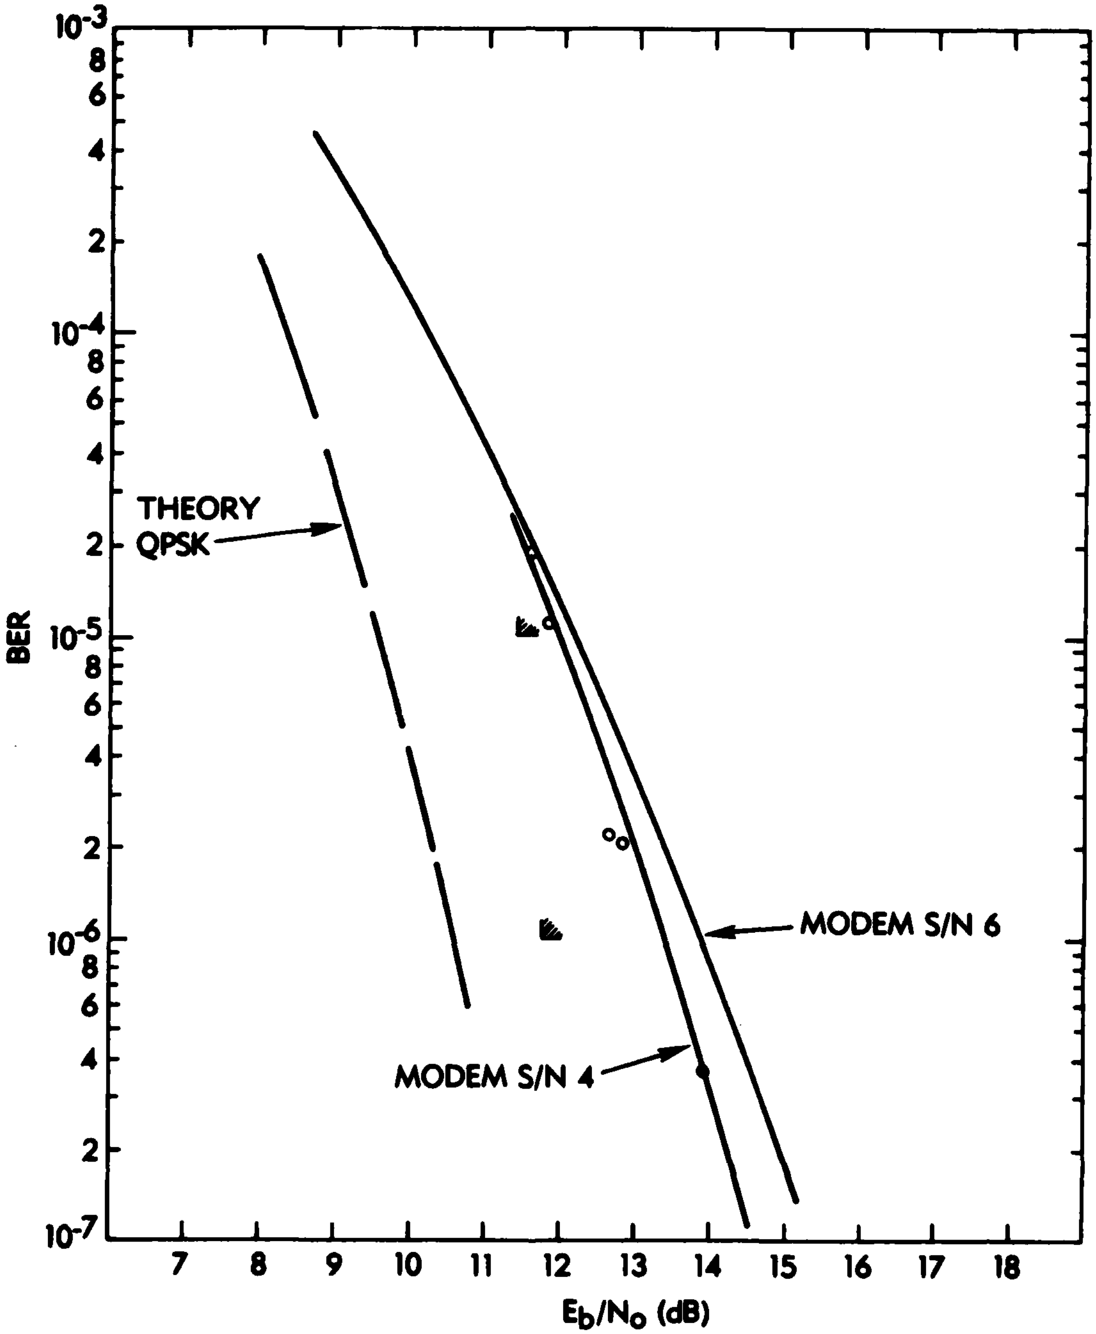

Figure 5-15. Modem Test Data, Etam PSP

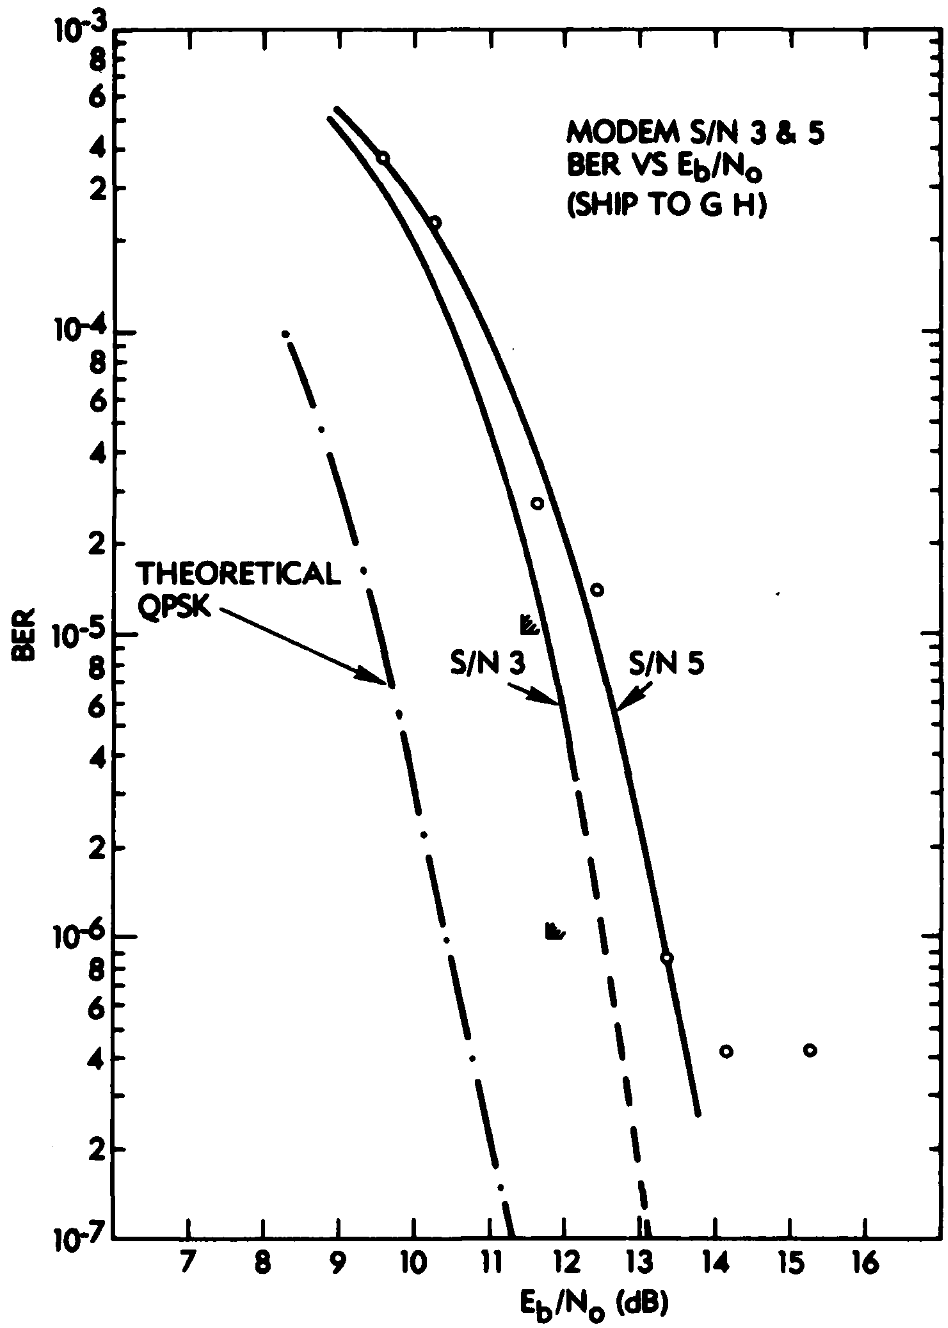

Figure 5-16. Modem Test Data, Goonhilly PSP

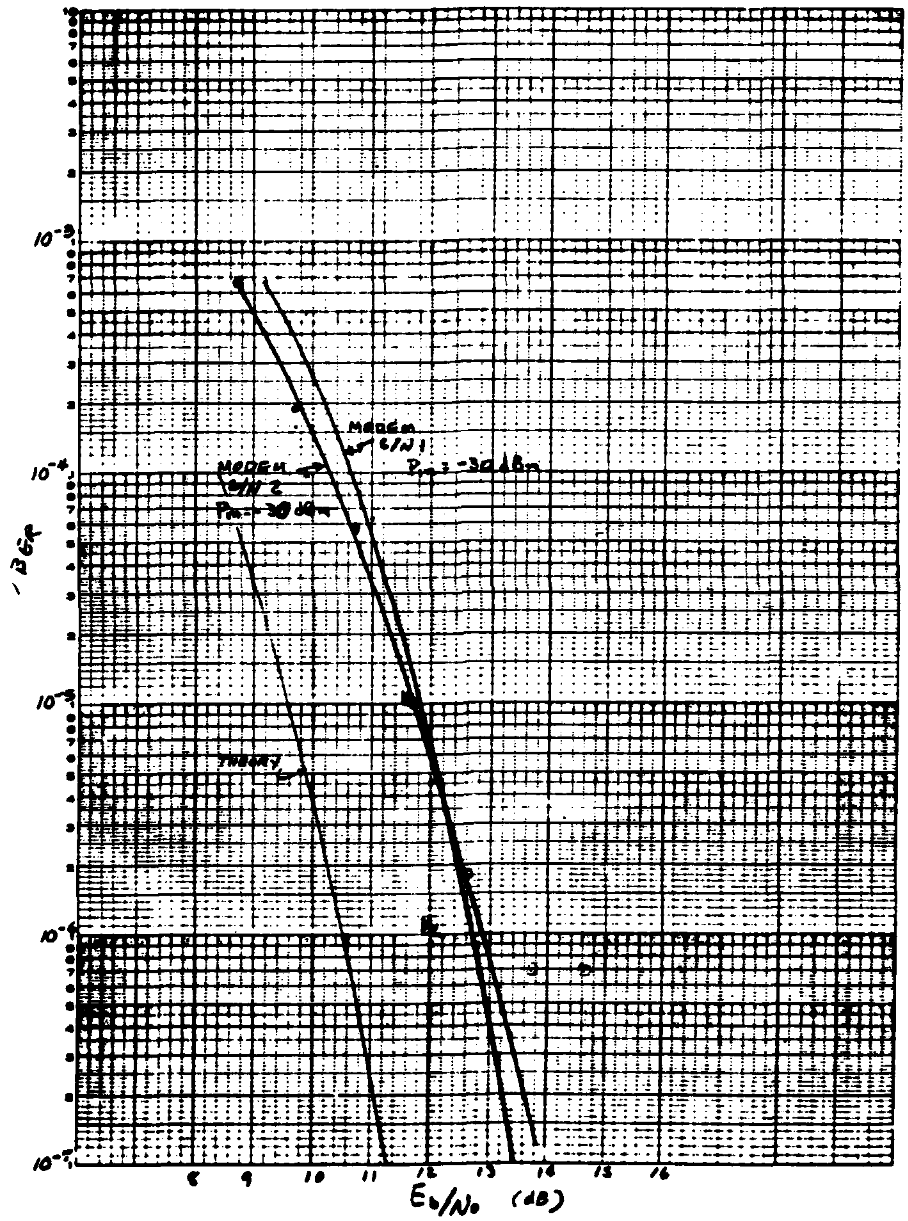

Figure 5-17. Modem Test Data, Tanum PSP

6. CHANNEL PERFORMANCE MONITORING

6.1 INTRODUCTION

COMSAT's recent involvement with SATNET has included the monitoring and interpretation of the channel summary data collected by the network control center at BB&N in Boston. These data are collected and distributed daily and provide a valuable summary of channel quality. On certain occassions, these data have indicated problems with the satellite links which required corrective action at the participating earth stations. The purpose of this section of the final report is to document the experience in this area during the period December 1978 to July 1979. The data presented apply only to the hello packets which are sent approximately once per second from each station and received by all stations in the network. The degree of success achieved in receiving these hello packets at each receiving station is a useful gross indicator of channel quality.

This section is divided into several subsections. In subsection 6.2, the SATNET frame format is reviewed, and the exact structure of the hello packet is summarized. From this description, a heuristic model can be developed for "failure" and "success" in receiving the hello packets at a particular receiving station. This model is then related to the count of missed hello packets which is one of the important items provided in the channel summary data from BB&N. In subsection 6.3, the SATNET link budgets are reviewed which provide an estimate of the performance to be expected on the SATNET links. The compilation of the channel summary data itself is provided in subsection 6.4. These data are presented as they were originally compiled in the form of a running summary. Certain improvements were made in this process as it evolved but no attempt has been made to correct the earlier results according to this experience. A final subsection 6.5 contains the conclusions reached as a result of the data analysis.

6.2 REVIEW OF SATNET FRAME AND HELLO-PACKET FORMAT

The SATNET frame format is shown at the top of Figure 6-1. The basic timing for the frame is provided by the SIMP where a hardware count of the 10 μs computer clock is made for a total of 1024 counts. This time epoch defines a virtual slot (VS) of duration 10.24 ms. The total frame consists of 128 of these virtual slots.

A 12 VS segment of the frame is reserved for the hello packets.* These packets have a nominal duration of 30 ms as shown in the lower part of Figure 6-1. The four earth stations in the network (Etam, Goonhilly, Tanum, Clarksburg; E,G,T,C) cyclically rotate their hello-burst transmissions through these slots with each station transmitting for an equal time in each of the four slots in the hello subframe. Since Clarksburg did not participate in the network during the data-collection period, its position was vacant.

*The format of both the SATNET frame and the hello bursts was changed on 9 April 1979. The hello subframe was reduced to 4 VS and each hello burst was reduced to a duration of approximately 9 ms by shortening the "data" segment to 72 symbols.

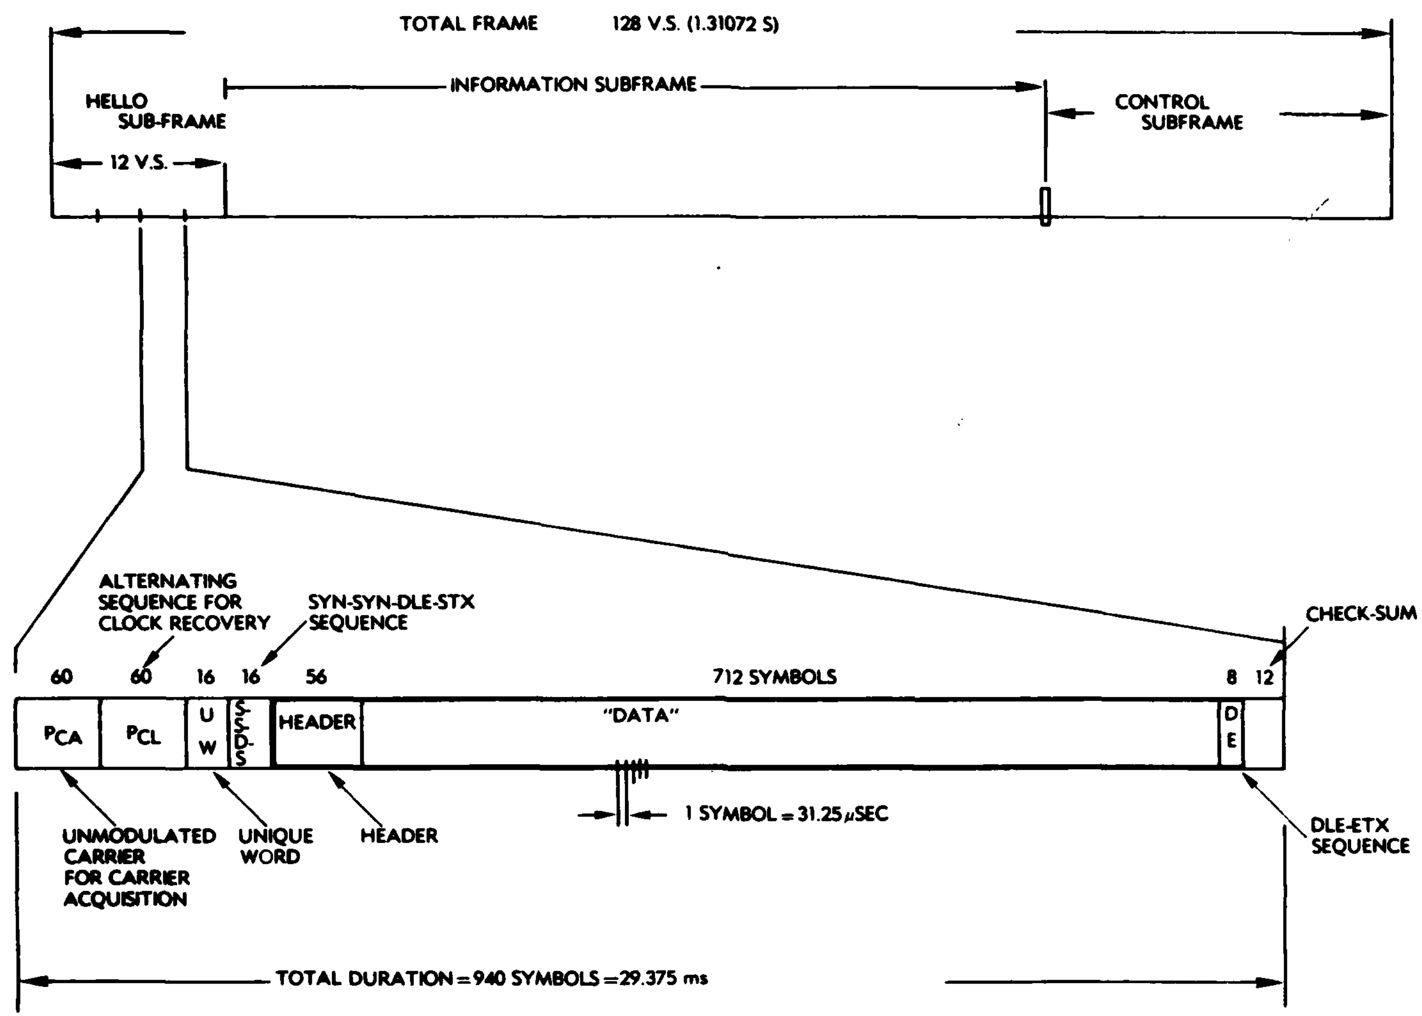

Figure 6-1. SATNET Frame and Hello-Packet Format
(prior to 9 April 1979)

The sequence was therefore as follows;

Position →	1st	2nd	3nd	4th
	E	G	T	_
	_	E	G	T
	T	_	E	G
	G	T	_	E

	E	G	T	_
	etc,			

The hello burst itself is identical at all stations and has the format shown at the bottom of Figure 6-1. The burst has eight distinct segments:

a. 60 symbols of unmodulated carrier
b. 60 symbols of an alternating 1, 1; 0, 0; 1, 1; 0, 0; . . ., QPSK sequence for clock recovery,
c. a 16 symbol unique word which signifies start-of-data and allows for ambiguity resolution,
d. 16 symbols (32 bits) for the SYN-SYN-DLE-STX sequence,
e. 56 symbols (7 16-bit words) of header,
f. 712 symbols of "data" (in the hello burst, this data is actually "filler"),
g. 8 symbols (16 bits) for the DLE-ETX sequence,
h. 12 symbols (24 bits) for a check sum (which is formed on all bits including and following the header).

The exact format of the transmitted hello burst is important because it impacts the interpretation of the SATNET channel data. The data provide a count of "missed" hello packets according to the logic diagram shown in Figure 6-2. The count of packets

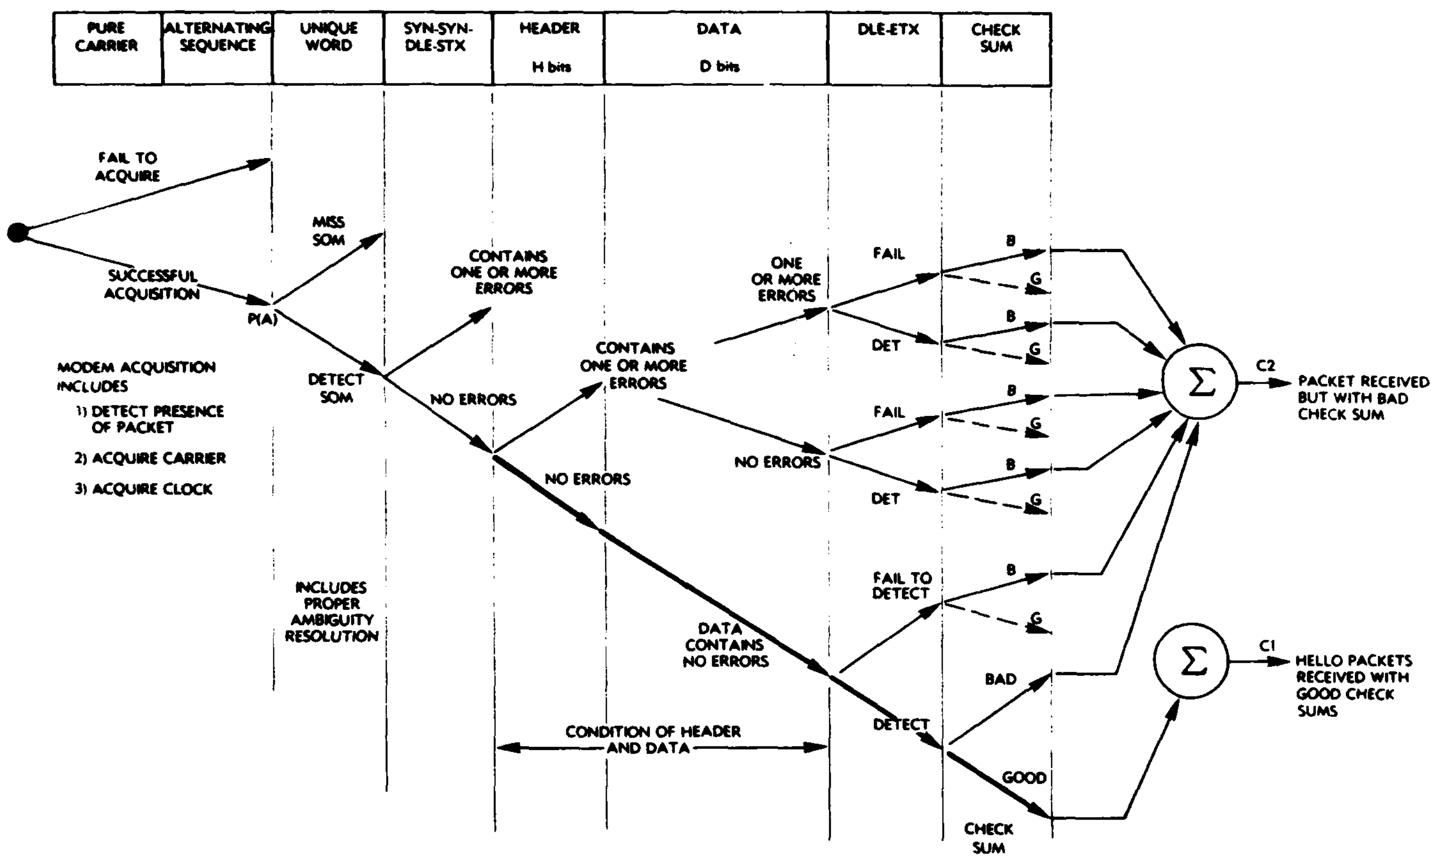

Figure 6-2. Missed Hello Packet Counting Procedure

received correctly takes place over a time interval when exactly 64 hello packets are sent by each station. By counting the number actually received with valid check sums, the number missed is inferred as the difference (i.e., #MISSED = #EXPECTED - #RECEIVED). The processing also keeps track of packets actually received, but with bad check sums. The events that must succeed for a hello packet to be received with a valid check sum are:

E_1 - the modems must acquire carrier phase and clock phase from the packet preamble. This event is dependent on link carrier-to-noise ratio (C/N), frequency offsets, and possibly interpacket intervals. The dependence of the probability of successful acquisition, P(A), on these factors is not known. In the following, it will be assumed that the preamble used in SATNET is long enough that P(A) approaches 1.0.

E_2 - the SOM unique word must be detected correctly (a few bit errors are allowed). Unique word detection takes place by correlating two 16-bit unique-word sequences, one on the I-channel and one on the Q-channel, against the recovered binary data sequences produced by the modem. If it is assumed that the false detection probability is negligibly low, then the probability of missing the unique word is dominated by the probability that more than two disagreements (errors) occur in one or both of the 16-bit words. Denoting the channel bit error probability as p, and letting $q \doteq 1 - p$, then the probability that the SOM is missed is;

$$\Pr\left\{\begin{matrix}\text{SOM}\\ \text{missed}\end{matrix}\right\} = \Pr\left\{\begin{matrix}\text{two or more errors occur in either,}\\ \text{or both, of the 16-bit unique words}\\ \text{on the I and Q channels}\end{matrix}\right\}$$

$$= 1 - \left\{\left[\binom{16}{0} p^0q^{16}\right]^2 + 2\binom{16}{1} p^1q^{15}\binom{16}{0} p^0q^{16}\right.$$

$$+ \left[\binom{16}{1} pq^{15}\right]^2 + 2\binom{16}{2} p^2q^{14}\binom{16}{0} p^0q^{16}$$

$$\left.+ 2\binom{16}{2} p^2q^{14}\binom{16}{1} pq^{15} + \left[\binom{16}{2} p^2q^{14}\right]^2\right\}$$

This relationship can be approximated for p reasonably small as;

$$\Pr\left\{\begin{matrix}\text{SOM}\\ \text{missed}\end{matrix}\right\} \cong 496p^2 + 11040p^3 + 96960p^4$$

Conversely,

$$\Pr\left\{\begin{matrix}\text{SOM}\\ \text{detected}\end{matrix}\right\} \doteq \Pr\{S\} \cong (1 - 496p^2)$$

E_3 - the SYN-DLE-STX sequence must be detected with no errors), only the last 24 bits of this 32-bit sequence must be error free. Again assuming a bit error probability p, the probability that this sequence is missed (contains one or more errors) is

$$\Pr\left\{\begin{matrix}\text{SYN DLE STX}\\ \text{missed}\end{matrix}\right\} = 1 - (1 - p)^{24} \cong 24p$$

Conversely, the probability that the sequence is detected correctly is $\Pr\{SDS\} \cong (1 - p)^{24}$

E_4 - Both the header sequence and the following data contain no bit errors. For "H" header bits and "D" data bits, the probability that neither contains bit errors is

$$\Pr\left\{\begin{matrix}\text{no errors in}\\ \text{header or data}\end{matrix}\right\} = (1 - p)^{H+D} \doteq \Pr(H + D)$$

In this analysis, no distinction is made between the header and the data bits. Both are evidently processed to produce the check sum. In practice, errors in the header may be more serious than errors in the data.

E_5 - DLE-ETX sequence following the data is detected. The probability that the DLE-ETX sequence is error-free is simply $\Pr(DE) = (1 - p)^{16}$.

E_6 - The check sum is examined. The 24-bit check sum will give a reliable indication of errors in the header, the data, and the DLE-ETX sequence provided that the check-sum itself does not experience bit errors. The probability of a good check-sum is thus $\Pr(C) = \Pr(H + D) \cdot \Pr(DE)(1 - p)^{24}$. Note that there are several other paths in Figure 6-2 that can yield a good check sum when, in fact, bit errors exist in that portion of the packet that is protected by the check sum. These cases can be combined by considering the use of the 24-bit check sum to detect errors in a sequence of length $H + D + 16 + 24$ bits. For the extreme case where $H + D$ is approximately 1000 bits, an approximation to the error-detecting capability of the CRC check can be made by considering it as analogous to a (1023, 1003, t) BCH code where t, the number of errors that can be corrected is 2. The minimum distance of such a code, $d = 2t + 1 = 5$, will allow the detection of four errors. Using this analogy for the CRC check, the CRC check will fail (yield a "good" check sum where in fact data errors are present) when five or more errors occur in the data sequence. The probability of this event is approximately

$$\Pr\left\{\begin{matrix}\text{good check sum}\\ \text{erroneously}\end{matrix}\right\} = \binom{H + D + 40}{5} p^5 (1 - p)^{H+D+40}$$

The probability that a hello packet is received with a good check sum can now be obtained as

$$\Pr\left\{\begin{matrix}\text{hello packet}\\ \text{received with}\\ \text{a good check sum}\end{matrix}\right\} = \Pr(A)\Pr(S)\Big\{\Pr(SDS)\Pr(H + D)\Pr(DE)(1 - p)^{24}$$

$$+ [1 - \Pr(H + D)\Pr(DE)] \cdot \binom{H + D + 40}{5} p^5 (1 - p)^{H+D+40}\Big\}$$

where the first term in brackets applies to the event where there are no errors in the header, data, DLE-ETX, <u>or</u> the check-sum bits, and the second term to the case where there are errors, but the check sum fails to detect these errors. This second term is negligible for the channel error probabilities of interest here, therefore;

$$\Pr\left\{\begin{matrix}\text{Hello packet}\\ \text{received with}\\ \text{a good check sum}\end{matrix}\right\} = \Pr(A)(1 - 496p^2)(1 - p)^{24}(1 - p)^{H+D}$$

$$\cdot (1 - p)^{16} (1 - p)^{24}$$

The probability of a missed hello packet is

$$\Pr\left\{\begin{matrix}\text{missed Hello}\\ \text{packet}\end{matrix}\right\} \cong 1 - \Pr(A)(1 - 496p^2)(1 - p)^{H+D+64}$$

where

$$p = \text{channel bit error probability}$$
$$H{+}D = \text{number of bits of header + data}$$

This expression is plotted in Figure 6-3 where the unknown acquisition probability, Pr(A), is retained as a parameter. In the region of high error probabilities, $p = 10^{-4}$ to 10^{-3}, the acquisition probability is likely to decrease somewhat, and this decrease tends to increase the missed-hello packet probability slightly. In the region of bit error probability around 10^{-6}, missed-packet percentages are approximately 0.2 percent for the long hello packet ($H + D = 1424$ bits), and approximately 0.05 percent with the short hello packet. In the latter case, the exact miss probability expected is fairly sensitive to the modem acquisition probability.

As bit error rate increases to 10^{-5}, the missed packet percentages would be expected to increase to slightly over one percent, and to approximately 0.3 percent for the long and short hello packets, respectively. A further increase in bit error rate to 10^{-4} results in missed packet percentages of 12.5 percent and 2.5 percent, respectively, for the long and short hello packets. The expected missed-packet percentages are approximately in proportion to packet length which is 1424/256 = 5.5 for the two cases in Figure 6-3.

The data presented in subsection 6.4 applies to missed hello packets in a 24 hr period. During this total period, the maximum number of hello packets transmitted by each station is

$$\#\text{Hello Packets} = \frac{24\ \text{hr/day} \times 3600\ \text{s/hr}}{10\ \mu\text{s} \times 1024 \times 128\ \text{s/frame} \cdot 1\ \text{frame/Hello Packet}}$$

$$= 65{,}918\ \frac{\text{Hello packets}}{\text{day}}$$

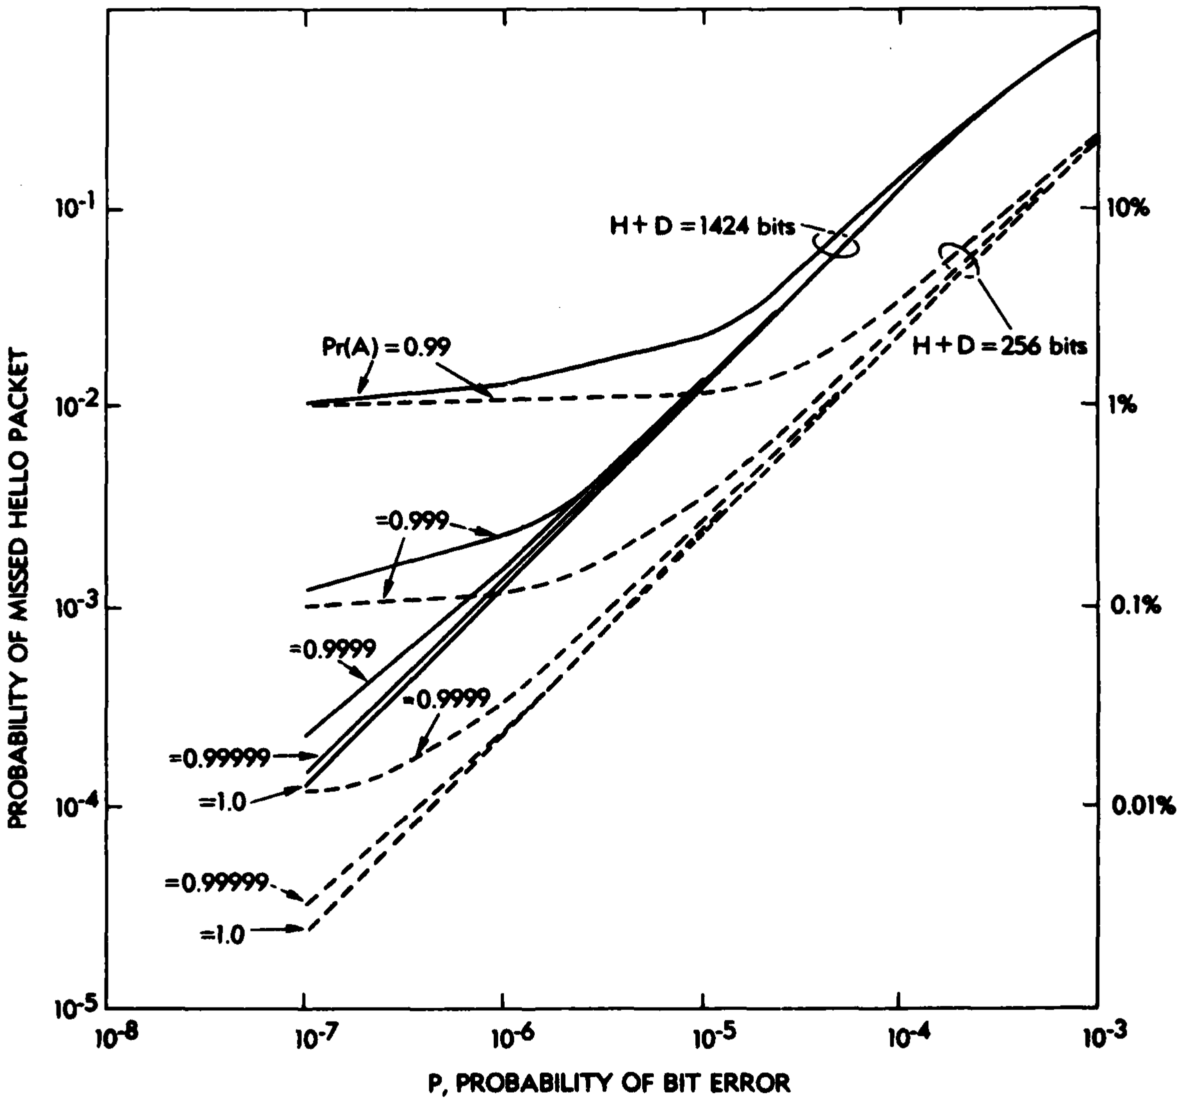

Figure 6-3. Probability of Missed Hello Packet vs Bit Error Probability

For a miss probability of 0.1 percent, for example, the expected number of hello packets missed at a particular receiving station, from one other transmitting station, would be 66.

A second element of data that is collected routinely on the SATNET links is the number of packets received correctly, but with a bad check sum. This second "count", C_2, (see Figure 6-2) is less useful than count C_1 since the current data is not sorted by link but is sorted only by receiving earth station. The probability that a packet is received with a bad check sum is approximately

$$\Pr\left\{\begin{matrix}\text{Hello packet received}\\ \text{with bad check sum}\end{matrix}\right\} = \Pr(A)\ \Pr(S)\ \Pr(SDS)\ \Pr(H + D)\ \Pr(DE)$$

$$\cdot\ [1 - (1 - p)^{24}] + \Pr(A)\ \Pr(S)\ \Pr(SDS)$$

$$\cdot\ \left\{1 - \Pr(H + D)\ \Pr(DE)\right\}$$

where the first term indicates a packet received with no data or header errors but with errors in the 24 bits of the check sum, and the second term applies to the reception of a packet with one or more errors in the header or data segment. For the latter situation, failure of the check sum (to detect these errors) has been assumed to be negligible. Substituting for the probabilities and simplifying gives

$$\Pr\left\{\begin{matrix}\text{Hello packet received}\\ \text{with bad check sum}\end{matrix}\right\} = \Pr(A)\ (1 - 496p^2)\ (1 - p)^{24})$$

$$\cdot\ \left[1 - (1 - P)^{H+D+40}\right]$$

Note that the probability of a missed hello packet is greater than the probability that a packet is received correctly but with a bad check sum; i.e.,

$$\Pr\left\{\begin{matrix}\text{missed Hello}\\ \text{packet}\end{matrix}\right\} < \Pr\left\{\begin{matrix}\text{packet received}\\ \text{with check sum}\\ \text{bad}\end{matrix}\right\}$$

This follows from simply subtracting the probabilities

$$\Pr\left\{\begin{matrix}\text{missed Hello}\\ \text{packet}\end{matrix}\right\} - \Pr\left\{\begin{matrix}\text{bad check}\\ \text{sum}\end{matrix}\right\} = 1 - \Pr(A)\ (1 - 496p^2)(1 - p)^{24}$$

When Pr(A) is close to 1.0, and p is small, this difference is approximately 24p. This difference is very small when "p", the channel bit error probability, is small. The ratio of packets with bad check sums to those missed completely is given approximately by

$$\frac{\Pr\{\text{bad check sum}\}}{\Pr\{\text{Miss}\}} = \frac{H + D + 40}{H + D + 64}$$

Note that this approximation applies only for very small values of p.

The assumption is now made that the SATNET channel experiences the same link conditions as any other SPADE or SCPC carrier in the INTELSAT IV-A transponder. In other words, it is assumed that from a link budget standpoint, the SATNET transmission is the same as any other carrier in the transponder. Most of the other carriers are carrying 56 kbit/s speech via 7-bit PCM, which, with overhead, are 64 kbit/s QPSK transmissions. The entries in Table 6-1 compare packet transmission quality to speech quality for various values of p. Also shown are the tentative quality standards for PCM speech transmission.

Table 6-1. Comparison of Quality for Various Values of p

Channel Bit Error Probability	Quality Standard (actually applies to 8-bit PCM)	Quality of PCM Speech*	% Missed Hello Packets in SATNET†	
			"Long" %	"Short" %
10^{-6}	Shall be achieved 80% of time	Excellent	0.16	0.050
10^{-5}		Very good	1.2	0.3
10^{-4}	0.3% of worst month	Noise barely perceptible	12.0	3.0
2×10^{-4}		Noise definitely noticeable	25	5.0
5×10^{-4}		Noise bad but tolerable to some users	50	12
10^{-3}	0.01% of the year	PCM speech is severely degraded	76	20

*approximate classifications
†"long" header + data 1424 bits; "short" 256 bits.

If error probability is 10^{-6} or better, the goal for 80 percent of the hours per year, speech quality would be excellent and the miss rate for the long SATNET hello burst should be approximately 0.1 percent. Of the approximately 66,000 hello packets sent each day, the expected number missed is 66. On some of the links in SATNET, on some days, this performance is achieved. On a few of the links, it is achieved routinely as noted in the data presented in subsection 6.4.

At an error rate of 10^{-5}, speech quality is still very good for those users who are sending speech in the transponder. Note, however, that the packet miss rate has increased to one percent. On the SATNET links, one would now expect 660 missed hello packets. The difference between 66 missed packets (0.1 percent) and 660 (1 percent) has a significant effect on the observed data and the latter value is usually cause for concern. It is likely that such variations occur routinely and they are probably not noticed by users who are transmitting PCM speech in the transponder.

A further degradation in error rate to 10^{-4}, which should happen only rarely (less than 20 hrs out of the "worst" month) will cause a significant degradation to the long SATNET hello packets. Likewise, it would probably produce a "noisy channel" to "discriminating listeners" of the speech traffic. Whereas a 10^{-4} error rate is certainly tolerable for PCM speech, it produces unacceptable conditions for the SATNET hello bursts. Evidently, from the data presented in Section 6.4, this error rate is not quite as rare as the quality standards would lead one to believe, particularly on certain SATNET links.

It can be concluded that the difference between 10^{-6}, 10^{-5}, or 10^{-4} error rates is very significant to the SATNET hello burst transmission. Such variations in error rate could result from variations of only a few dB in link C/N; such variations are likely. Other speech users in the transponder would probably be very tolerant of such variations and only notice the extremes.

6.3 REVIEW OF SATNET LINK BUDGETS

This subsection presents a summary of the link budgets that apply to the nine satellite links which interconnect the Etam, Goonhilly, and Tanum earth stations. These values follow the general format of Reference [1]. Certain elements in the link budget are commented upon and any possible variability is noted.

Table 6-2 gives the up-link relationships.

$$\left(\frac{C}{N_o}\right)_{UP} = \Omega_{SAT}\ bO_i\ \frac{\lambda_u^2}{4\pi} \cdot \left(\frac{G}{T}\right)_{SAT} \cdot \frac{1}{k} \cdot \frac{1}{N} \cdot L_u$$

where

Ω_{SAT} = saturation flux density
bO_i = transponder input backoff
λ_u = up-link wavelength
(G/T) = satellite receive G/T ratio
k = Boltzmann's constant
N = number of carriers in the transponder
L_u = up-link loss

The down-link C/N_o ratio can be obtained from the relationship;

$$\left(\frac{C}{N_o}\right)_{DN} = \frac{e.i.r.p._{SAT}}{4\pi R^2} \, bO_o \, \frac{\lambda_D^2}{4\pi} \cdot \left(\frac{G}{T}\right)_{ES} \cdot \frac{1}{k} \cdot L_D \cdot L_u \cdot \frac{1}{N}$$

where

$e.i.r.p._{SAT}$ = satellite e.i.r.p. at beam edge
R = range to the satellite
bO_o = transponder output backoff
$(G/T)_{ES}$ = earth station G/T ratio
L_D = down-link loss
L_u = any up-link loss

The appropriate values for INTELSAT IV-A are given in Table 6-3.

Table 6-2. Up-Link C/N_o for Operation in the Spade Transponder Up-link Frequency 6320.2275

Item	Description	Value	Comment
1	Saturation flux density (dBW/M^2)	-67.5	
2	Input backoff	-11.0	Based on a "nominal" of 320 carriers accessing the transponder (see note a)
3	Area of unit gain antenna assuming frequency of 6.32 GHz	-37.5	
4	$(G/T)_{SAT}$	-18.6	
5	1/k (Boltzmann's Constant)	+228.6	
6	Total up-link C/N_o	94.0	
7	Sharing among N carriers (1/N)	-25.0	Based upon a nominal loading of 320 carriers (see note (a)
8	$(C/N_o)_{UP}$ available to one carrier (no loss)	69.0	
9	Atmospheric loss and pointing loss	-0.7	
10	Up-link loss L_u	$-L_u$	(see note b)
11	$(C/N_o)_{UP}$	68.3 $-L_u$	

NOTES:

a. with fewer than 320 up-link carriers, input backoff increases but the "sharing" loss is reduced so that $(C/N_o)_{UP}$ tends to remain constant. If the nominal operating point is -10 dB input backoff rather than -11 dB, then $(C/N_o)_{UP}$ will increase by 1 dB.

b. Loss applies to the single SATNET up-link carrier. Any loss will decrease $(C/N_o)_{UP}$ and reflect directly into the down-link as a power sharing loss.

Table 6-3. Down-Link C/N_o for Operation in the SPADE Transponder Down-Link Frequency 4095.2275 MHz

Item	Description	Value	Comment
1	Satellite e.i.r.p. at beam edge	+22.0 dBW	One dB more e.i.r.p. may be available (+23)
2	Output backoff (dB)	-5.0	May be slightly less (-4.2 dB) (see note a)
3	Free space loss	-196.6	(see note b)
4	Antenna pointing and atmospheric loss	-0.7	
5	Receive earth station (G/T)	+40.7	
6	1/k	+228.6	
7	Total $(C/N_o)_{DN}$	89.2	
8	Sharing among N carriers	-25.0	Based on 320 carriers in transponder (see note a)
9	$(C/N_o)_{DN}$ available to one carrier	64.2	
10	Down-link loss	$-L_D$	Applies only to the SATNET carrier
11	Up-link loss	$-L_u$	Applies only to the SATNET carrier
12	$(C/N_o)_{DN}$ with losses	$64.2-L_D-L_u$	

NOTES:

a. For fewer than 320 carriers, operating point tends to change but sharing also changes, so these terms tend to compensate. For worst case sharing with 400 active carriers, $(C/N_o)_{DN}$ would be reduced by 1 dB.

b. Applies to reception at Tanum with approximately a 17° elevation angle. Will experience 0.1 dB more loss at Etam, 0.2 dB less loss at Goonhilly.

A final term that should be included in the link budget is the ratio of carrier power to intermodulation noise. It is difficult to estimate this value exactly at a particular frequency slot in the transponder. For the output backoff level of -5 dB, with the transponder loaded with 320 carriers, the ratio of total carrier power to intermodulation power would be approximately 18.5 dB. If the intermod noise power were spread uniformly over the 36 MHz transponder bandwidth, a total C/N_o of 94 dB-Hz would result, or 69 dB-Hz per carrier with 320 equal carriers. This $(C/N_o)_{IM}$ value of 69 dB-Hz will be used as a "worst-case" value for IM levels. Note that this ratio would be reduced by any up-link loss L_u.

Overall C/N_o can now be obtained as

$$\left[\left(\frac{C}{N_o}\right)_T = \left\{\left[\left(\frac{C}{N_o}\right)_{UP} \cdot \ell_u\right]^{-1} + \left[\left(\frac{C}{N_o}\right)_D \ell_u \ell_d\right]^{-1} + \left[\left(\frac{C}{N_o}\right)_{IM} \ell_u\right]^{-1}\right\}^{-1}\right.$$

where ℓ_u and ℓ_d are the up- and down-link losses expressed as numbers <1.0 and the C/N_o values are expressed as numbers, not as dB quantities. Total C/N_o values are given in Table 6-4.

Table 6-4. Total C/N_o for One Carrier*

L_u (dB)	L_D (dB) → 0	1	2	3	4
0	61.8(62.8)	61.2(62)	60.6(61.2)	59.8(60.4)	59.1(59.6)
1	60.8(61.8)	60.2(61)	59.6(60.2)	58.8(59.4)	58.1(58.6)
2	59.8(60.8)	59.2(60)	58.6(59.2)	57.8(58.4)	58.1(57.6)
3	58.8(59.8)	58.2(59)	57.6(58.2)	56.8(57.4)	56.1(56.6)
4	57.8(58.8)	57.2(58)	56.6(57.2)	55.8(56.4)	55.1(55.6)

*Numbers in parenthesis assume that intermods are negligible

The C/N_o values can be converted to C/N ratios in the 38 kHz noise bandwidth as $C/N = (C/N_o)_T - 45.8$ dB. Alternately, the values can be converted to E_b/N_o as $E_b/N_o = (C/N_o)_T - 48.1$ dB.

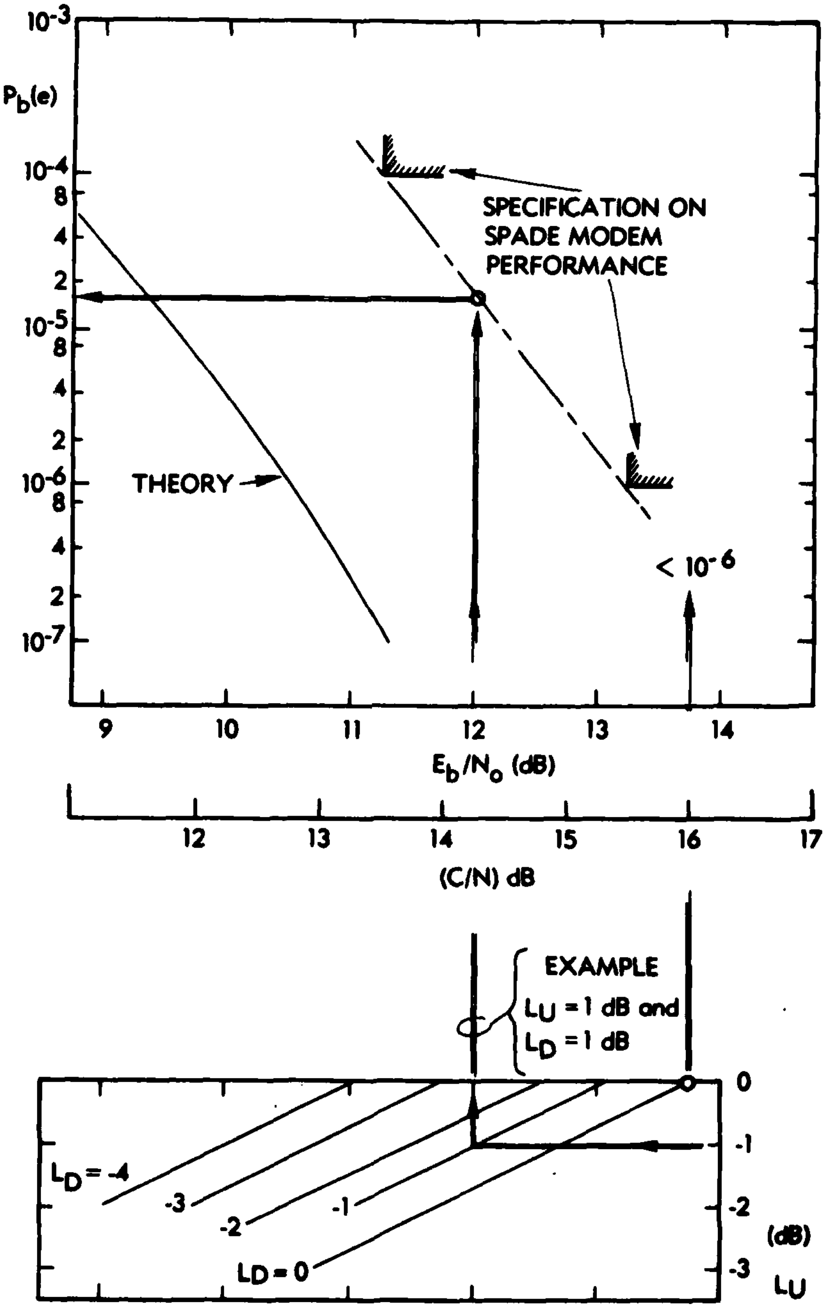

Figure 6-4. Impact of Up-Link and Down-Link Losses on SATNET Bit Error Rate

6.4 MISSED-HELLO PACKET DATA

The actual missed hello-packets experienced on SATNET since the beginning of the year 1979 are reviewed in this subsection. The raw data for January and February 1979 is summarized in Table 6-5. The nine SATNET links are consistently coded by the numbers 1 through 9 as follows: (E = Etam, G = Goonhilly, T = Tanum):

		From E	From G	From T
To:	E	1	2	3
	G	4	5	6
	T	7	8	9

All of the numbers given in Table 6-5 are the total number of hello packets missed in a day, by link. The data for links 1, 2, 3, 5, 8, and 9 are plotted in Figure 6-5.

The SATNET error data for March 1979 is summarized in Table 6-6. Included in this table is an "outage correction factor" to account for the fact that on eight of the ninteen days, data was not collected over the whole 24 hour period. Knowing these outage times, one can "correct" the averaging interval used to calculate "fraction of hello packets missed."

The correction factors are included on Figure 6-6 as the heavy black arrows on certain days. All data points on that particular day should be moved upward by the length of the arrow. Note for example, that on the 2nd, 9th, and 18th of March this would reduce somewhat the seemingly drastic improvement from one day to the next.

Table 6-5. Errors Reported on SATNET; January 1979 and February 1979

January 1979

Link	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
1	36	55	95	80	82	515	3	26	15	4	695	0		1	0
2	1607	193	157	163	347	877	937	104	53	46	640	0		91	30
3	116	64	116	133	421	794	111	56	38	16	621	0		4	43
4	21	82	145	29	59	416	2	18	16	27	145	0		70	13
5	1793	168	259	110	243	692	1653	164	144	154	877	244		495	579
6	69	49	59	19	115	546	69	33	31	24	87	138		3	48
7	18	40	149	26	44	461	5	16	22	26	16	0		35	35
8	1062	172	96	54	157	664	501	65	68	75	140	26		42	162
9	90	29	145	58	284	686	122	42	39	27	234	23		59	52
Daily Avg.	534	94	135	74	194	633	378	58	47	44	383	47		88	106

Link	18	19	20	21	22	23	24	25	26	27	28	29	30	31
1				1	3		3	6	0	3	78	11	6	4
2				29	43		37	55	22	21	53	132	82	31
3				5	11		2	7	3	3	3	13	7	9
4				1	5		0	2	2	0	16	3	55	16
5				11	14		19*	24	8	18	26	69	41	39
6				4	12		0	3	5	4	10	9	54	28
7				1	4		9	1	142*	4	11	4	8	100
8				9	7		20	15	81*	24	27	62	70	54
9				4	14		3	3	67*	6	7	21	3	33
Daily Avg.				7	12		10	12	38	9	25	36	36	35

February 1979

Link	1	2	3	4	5	6	7	8	9	10	11	12	13	14
1	2	17	16	10		345	24	24	8		19	22	63	49
2	35	30	23	34		2130	132	26	53		0	18	99	596
3	9	11	15	11		139	17	20	27		14	16	15	44
4	5	11	5	8		104	9	44	23		0	26	39	87
5	18	22	25	29		4487	91	22	134		0	39	101	1406
6	7	12	23	10		29	7	21	3		0	28	5	15
7	10	13	61	10		66	7	97	16		14	202	113	89
8	16	16	31	18		331	85	75	22		0	49	210	79
9	9	9	30	16		41	6	82	12		5	184	105	7
Daily Avg.	12	15	25	16		852	42	45	33		6	64	83	263

Link	15	16	17	18	19	20	21	22	23	24	25	26	27	28
1	40	10			1	65	33	37	534*	84	89	147	121	
2	9985	0			0	85	93	60	562*	50	65	247	60	
3	394	160			7	109	102	47	865	42	42	159	75	
4	395	0			0	74	9	55	47	56	60	119	91	
5	14962	10			0	63	243	57	41	26	78	74	72	
6	3068	0			0	88	65	82	65	11	9	98	60	
7	13	11			1	83	91	27	32	4	8	189	107	
8	4432	0			0	171	98	26	32	0	0	131	36	
9	728	216			9	75	99	156	50	0	0	87	40	
Daily Avg.	3778	45			2	90	92	60	247	30	39	139	73	

*errors sensitive to position in hello subframe
Blank entries; no data available

Average over 47 days 193 (2.95×10^{-3})

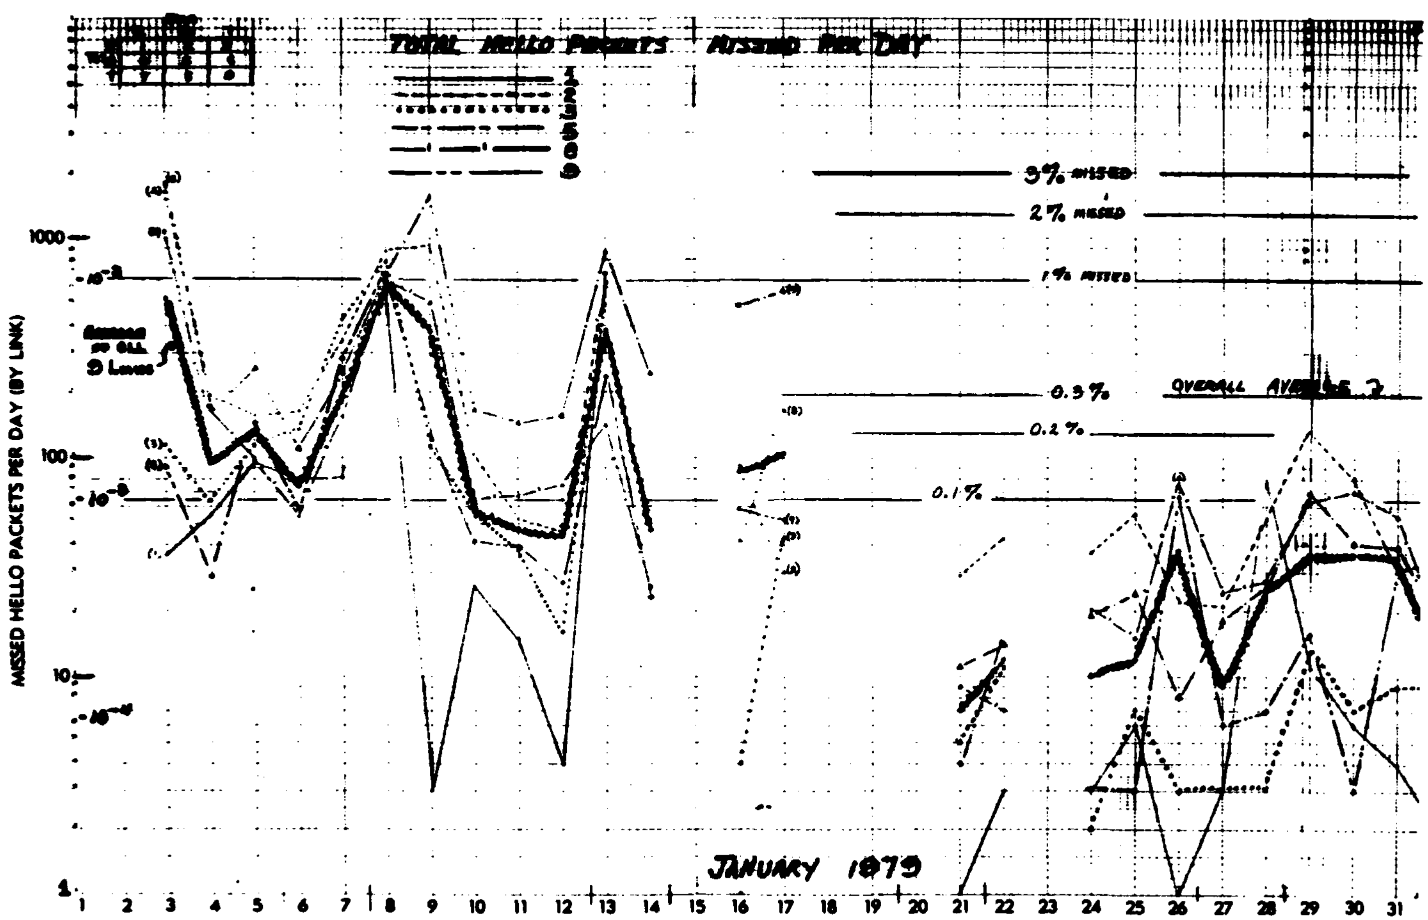

Figure 6-5a. SATNET Missed Hello Packet Data; January 1979

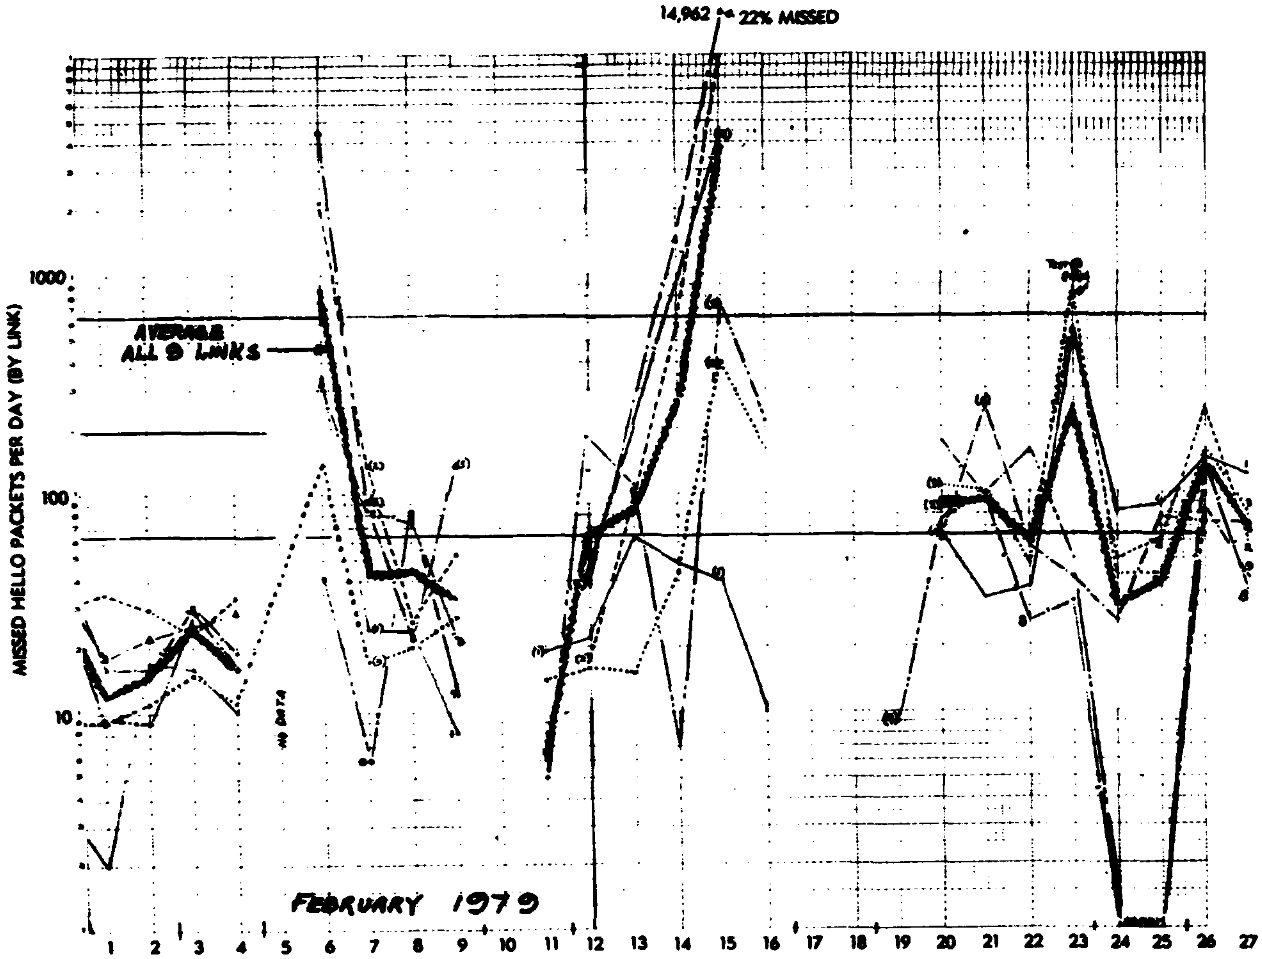

Figure 6-5b. SATNET Missed Hello Packet Data; February 1979

Table 6-6. Errors Reported on SATNET, March 1979

Link	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20†	21†	22	23	24	25	26	27	28	29	30	31
1	73	18*	160*	201*		215		186	52	133	21		14	237	87	77		37				23	256	143	63				899*		
2	139	50*	211*	273*		444		203	129	177	54		70	241*	185	86		29				61	167	99	798				1149*		
3	123	13	113	448		189		179	126	134	15		17	117	109	40		9				21	27	545*	743*				100		
4	94	6	69	65		180		103	30	169	59		71	44	95	62		26				20	75	126	778*				37		
5	102	262	618	518		1346		916	719	1863	788		1067	739	418	422		269				232	740	1311	4054				527		
6	99	20	97	308		343		154	118	177	43		81	43	82	43		17				26	45	928	3251				25		
7	340*	16*	130	123		27		46	33	52	8		19	31	69	196		9				74	142	646	167				55		
8	271*	33	208	162		48		70	63	86	9		45	38	125	525		6				121	286	1169	385				9		
9	211*	13	86	186		24		90	2	32	10		14	20	23	148		4				93	192	843	201				16		
Daily Avg.	161	47	188	253		312		216	141	313	111		155	167	132	177		45				74	214	645	1160				313		
Outage** Corr. Factor	2.43	2.0				1.06			1.32							1.17		2.71				1.18	1.45								

* Errors sensitive to position in hello subframe.

† Testing at Etam on 21 and 22 March (no data from 9:00 a.m. on 20 March until 5:30 p.m. on 21 March)

** Outage correction factor ≐ 24 ÷ total hours in day for which data is available

Average, weighted by outage factor, all 9 links over 19 days 283 errors daily (4.33×10^{-3})

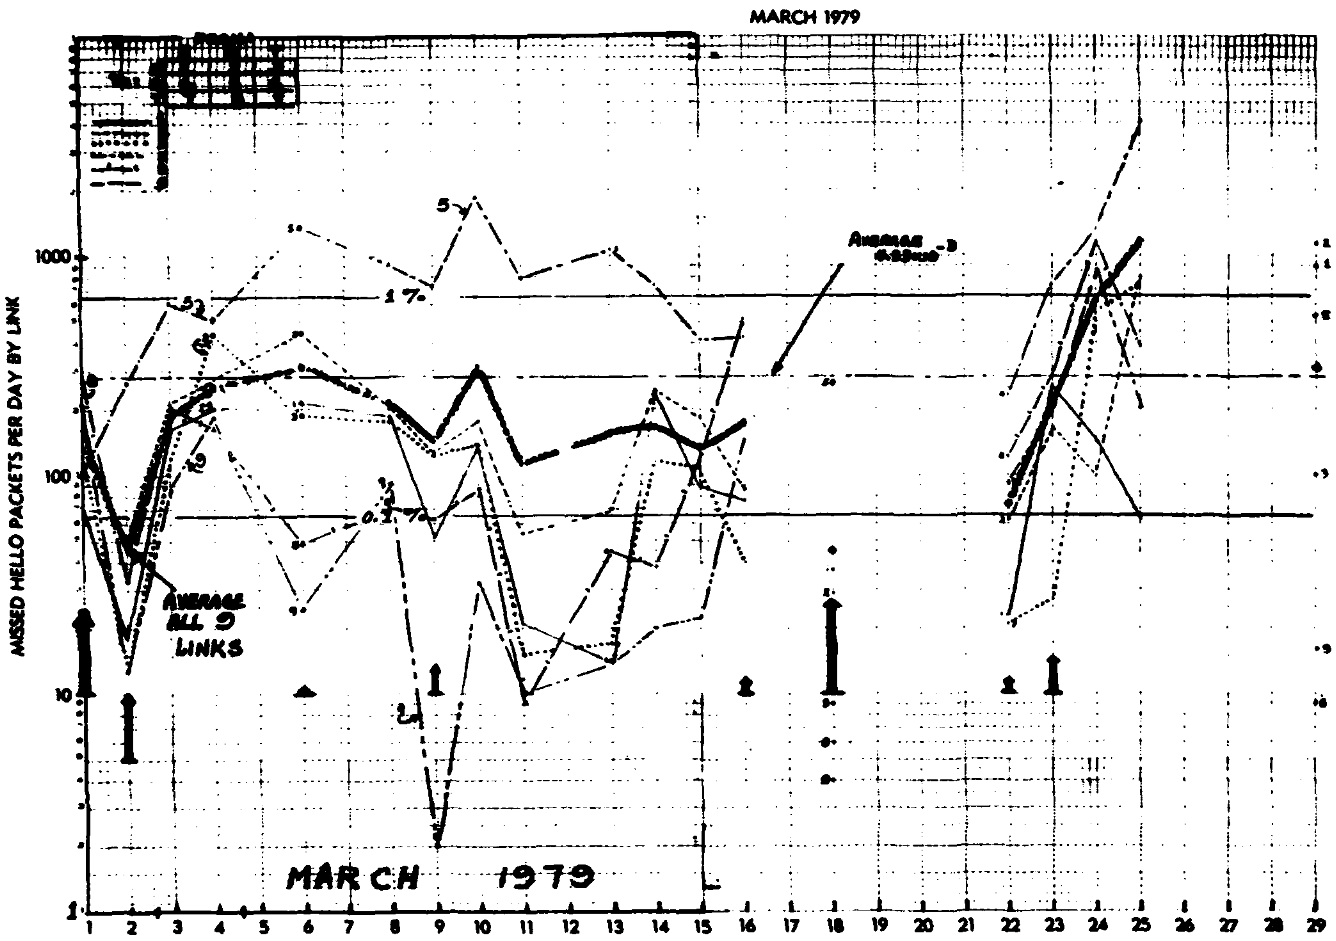

Figure 6-6. SATNET Missed Hello Packet Data; March 1979

The error data for April 1979 are summarized in Table 6-7 and Figure 6-7. During this period, some changes were made which are reflected in the data. First, on 3 or 4 April, a correction to the transmitted power level was made at Goonhilly. This level was low, by at least 1 dB, and had evidently been low compared to the transmissions from the other earth stations, for several months. Secondly, the format of the hello bursts was changed on 9 April. The first change seemed to stabilize the performance of the links originating from Goonhilly. The effect of change is shown more clearly in Figure 6-8 which plots missed hello packets on link No. 5 (Goonhilly). The test report from the Goonhilly earth station also recounted problems with antenna tracking and the HPA which could have accounted for additional losses of up to 2 dB or certain days during the early part of 1979. These variations would likely explain the highly erratic behavior of the links originating from Goonhilly. Antenna tracking problems, in particular, could explain why the Goonhilly→Goonhilly link seems to have extremely high packet miss rates on certain days. Note from Figure 6-4, for example, that an up-link and a down-link loss of 2 dB results in a bit error rate significantly higher than 10^{-4} which gives a missed-packet percentage greater than 10 percent.

Data for May 1979 is summarized in Table 6-8 and Figure 6-9. Data for June 1979 is given in Table 6-9 and Figure 6-10, and for July 1979 in Table 6-10 and Figure 6-11.

6.5 CONCLUSIONS

The data presented in this section summarizes the missed-hello packet experience on SATNET during the first half of 1979. The data during this period was punctuated by high miss percentages in the transmissions originating from Goonhilly. Several requests to the earth station, and routine checks of power levels, failed

Table 6-7. Summary of SATNET Data for April 1979

Link / Day	31 March	1	2	3	4	5	6	7	8	9	10	11	12	13	14
1	54	141	1901	899	95	2430	114	89	683	1308†	54	66	78	90	126†
2	193	779	1711	718	74	1039	5	3	3	201†	21	100	59	245†	1070†
3	5	59	70	0*	0*	325	4	65	91	200†	621	15	5	13	36
4	96	60	72	39	31	264	15	1	30	90	47	22	4	36	44
5	211	931	597	163	104	232	49	7	99	248	110	44	162	99	190
6	9	37	118	0*	0*	60	55	2	72	124	357	3	22	16	61
7	112	7	109	0*	0*	114	37	4	49	9	47	0	1	51	9
8	52	26	259	0*	0*	46	40	2	48	95	69	29	22	36	46
9	35	18	206	0*	0*	70	42	7	57	6	85	34	19	56	10
Daily Average	85	229	560	444	76	508	40	20	126	258	156	35	41	71	177
Outage Correction Factor	1.47	1.12	-	1.42		1.30	1.66	1.39	-	-	(2.4)	1.37	1.40	1.44	-

Link / Day	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30
1	71	49	88		35	80	72	220	265	160†	68†	228*	56	2	53	12
2	922†	128	35		64	120	17	28	17	91†	39†	171†	22†	1	109	22
3	16	136	67		1	230	9	19	5	22	23†	30†	4	1	93	4
4	52	174	233		32	279	49	37	18	15	24	61	1	19	56	30
5	239	86	118		153	229	232	191	60	70	236	233	26	35	112	164
6	101	37	110		20	260	108	80	28	27	66	87	4	22	101	30
7	45	77	113		10	209	33	20	7	7	6	16	0	51	48	26
8	48	4	28		5	143	22	16	9	10	16	136	0	32	96	31
9	56	0	27		2	140	39	14	6	14	6	14	2	1	106	7
Daily Average	172	77	88		36	187	64	69	46	47	53	109	13	18	86	37
Outage Correction Factor	-	1.26			1.43					1.25	109	1.07	7.27	6.0	1.06	1.18

*TAMUN down for repairs
†Errors sensitive to position in hello subframe

Weighted Average 30 Day 162 (2.5×10^{-8})

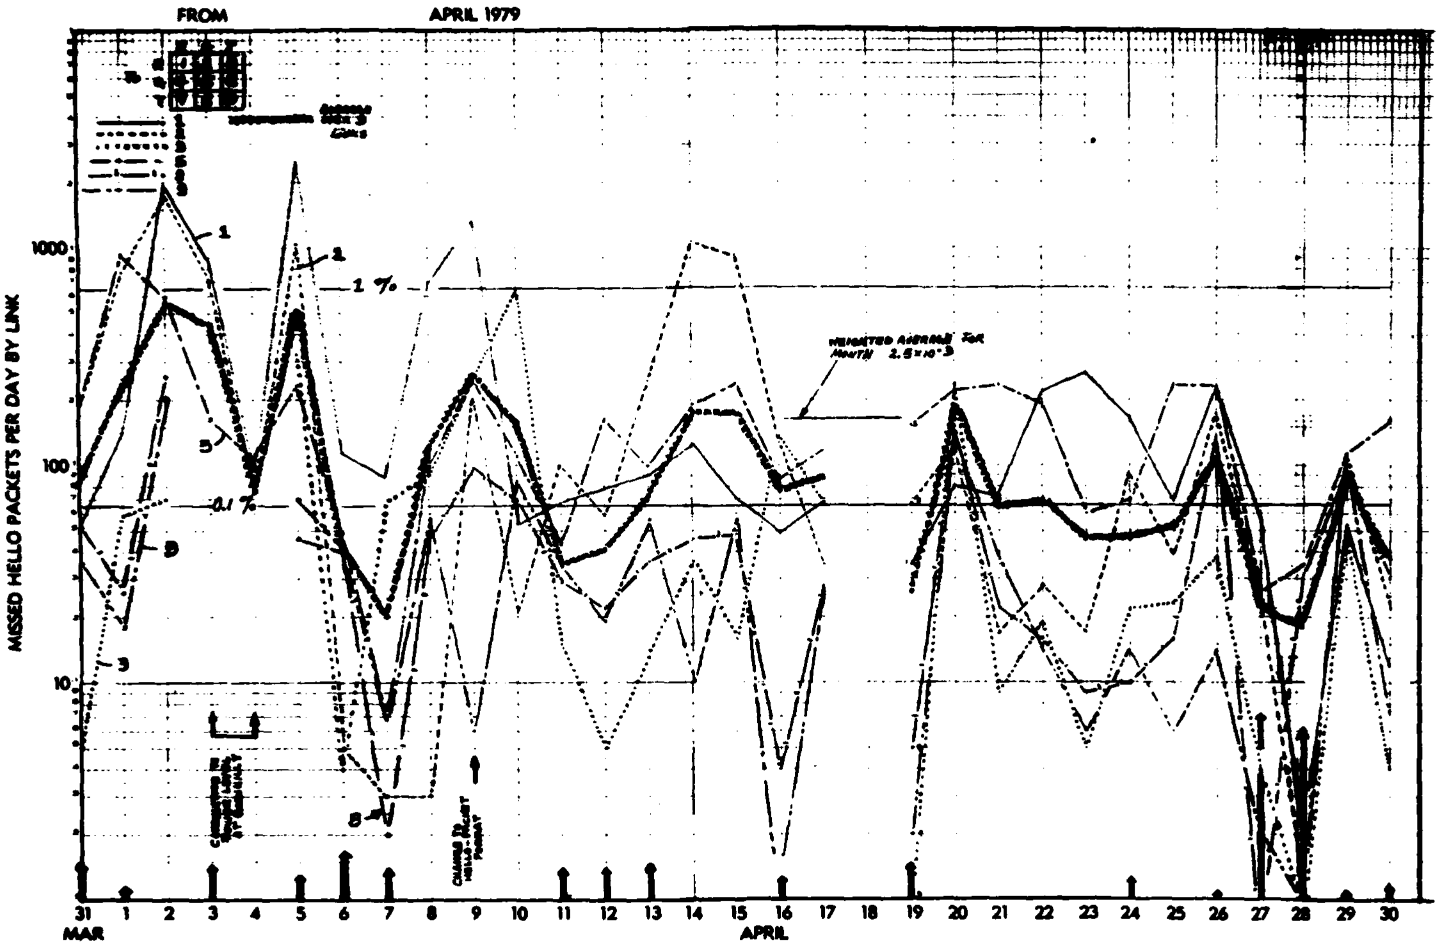

Figure 6-7. Missed Hello Paackets Per Day By Link, April 1979

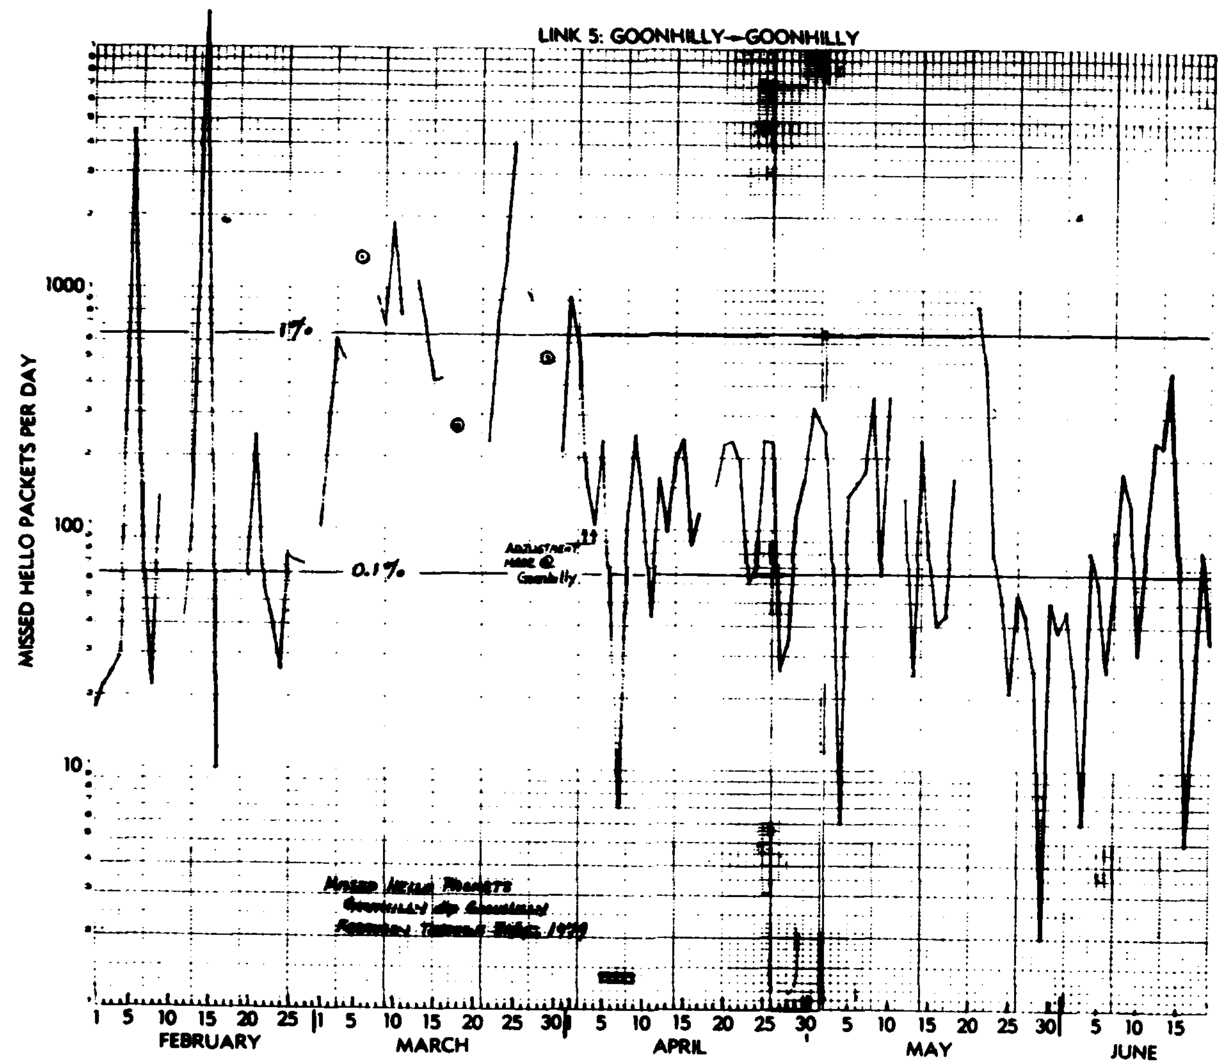

Figure 6-8. Missed Hello Packets, Goonhilly→Goonhilly, February through April 1979

Table 6-8. Summary of SATNET Data for May 1979

Table 6-8. Summary of SATNET Data for May 1979

Link	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
1	183	259	7	57	6	43	131	265	9	103		8	15	72	974	57	235	79			178	116	215*	811*	313*	200*	244*	2156*	21*	400	9
2	137	257	41	52	20	30	118	306	17	143		87	1378	1677	911	41	43	1185			1390	176	220*	1783*	185*	499*	103*	534	30*	98	12
3	235	232	17	48	1	9	114	250	2	982		8	14	4	281	20	97	45			135*	76	89	42	514	30	42	155	4	1	24
4	302	69	76	49	27	73	217	220	42	357		8	12	1259	1173	109	220	82			266*	101	100	22	21	15	5	47	1	52	6
5	322	252	68	6	140	158	175	356	65	354		134	25	237	77	41	44	167			851	460	78	50	21	54	44	26	2	50	38
6	474	115	86	5	22	40	124	190	42	342		17	9	68	8	30	5	75			231*	101	100	49	566	29	44	33	3	11	9
7	257	198	15	41	3	65	200	301	8	295		4	8	1248	1142	107	375	67			150	93	58	42	49	14	44	171	1	52	9
8	273	233	17	7	60	36	139	120	16	195		90	4	98	6	37	13	421			82	148	134	42	43	26	43	53	0	9	12
9	260	165	17	1	2	2	119	68	4	225		8	5	71	1	28	63	34			118	81	19	47	43	13	43	60	2	5	10
†	271	197	38	30	31	51	148	230	23	332		40	163	526	508	52	121	239			378	150	106	320	195	98	72	359	7	75	14
††	-	-	-	-	-	-	1.06	-	-	2.96		-	-	-	-	1.23	1.07	1.07			1.46	-	1.11	-	1.4	-	-	-	-	-	-

†Daily Average
†† Outage Correction Factor

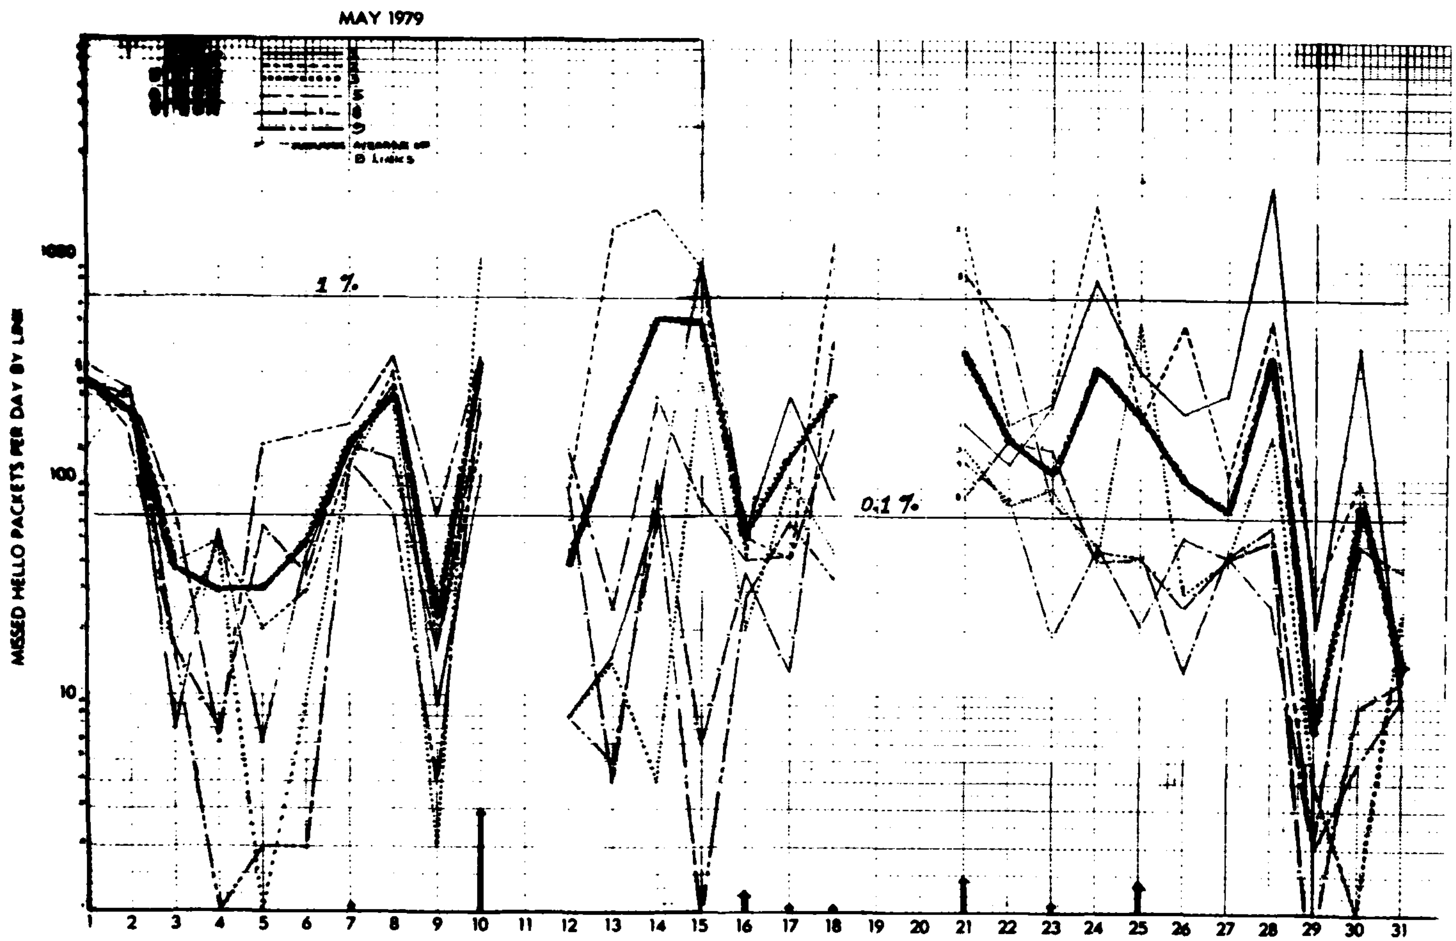

Figure 6-9. Missed Hello Packets Per Day By Link; May 1979

Table 6-9. Summary of SATNET Data for June 1979

Link	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30
1	204	35	10	13	25	47	49	10	2	17	27	27	15	1328	22	10	29	30	1	3	36	8	11	31	95	1		411	601	
2	65	23	9	19	3	41	39	23	5	10	17	30	21	773	19	10	31	145	39	142	54	95	4	15	84	0		75	269	
3	84	17	8	15	398	41	114	14	21	11	11	11	18	783	62	20	75	22	6	134	20	778	49	1132	60	350		63	281	
4	437	23	9	20	28	128	36	13	6	17	32	42	38	2756	28	8	17	4	2	3	21	12	15	33	-	17		149	625	
5	46	23	6	81	58	26	57	173	126	31	98	232	227	460	65	5	17	81	35	24	29	52	11	25	16	2		42	76	
6	49	26	5	32	66	16	110	78	41	12	26	69	64	263	125	44	66	11	10	16	26	57	48	42	97	2		4	4	
7	403	68	8	23	69	19	42	43	7	20	29	28	85	2687	106	24	70	32	5	5	40	103	78	64	130	17		90	235	
8	33	22	5	25	74	3	5	63	11	12	41	34	83	588	97	23	91	64	30	11	43	123	69	51	99	3		42	54	
9	39	21	7	18	21	4	78	55	10	16	20	40	63	538	109	33	75	8	12	2	35	112	77	69	143	70		13	11	
†	151	29	7	27	82	36	59	52	25	16	33	57	58	1130	70	20	52	44	16	38	34	141	40	162	91	51		99	240	
††	1.05	-	-	-	-	-	1.04	1.24	-	-	-	-	1.08	-	-	-	-	-	-	1.14	-	-	-	-	-	-		-	-	

†Daily Average
††Outage Correction Factor

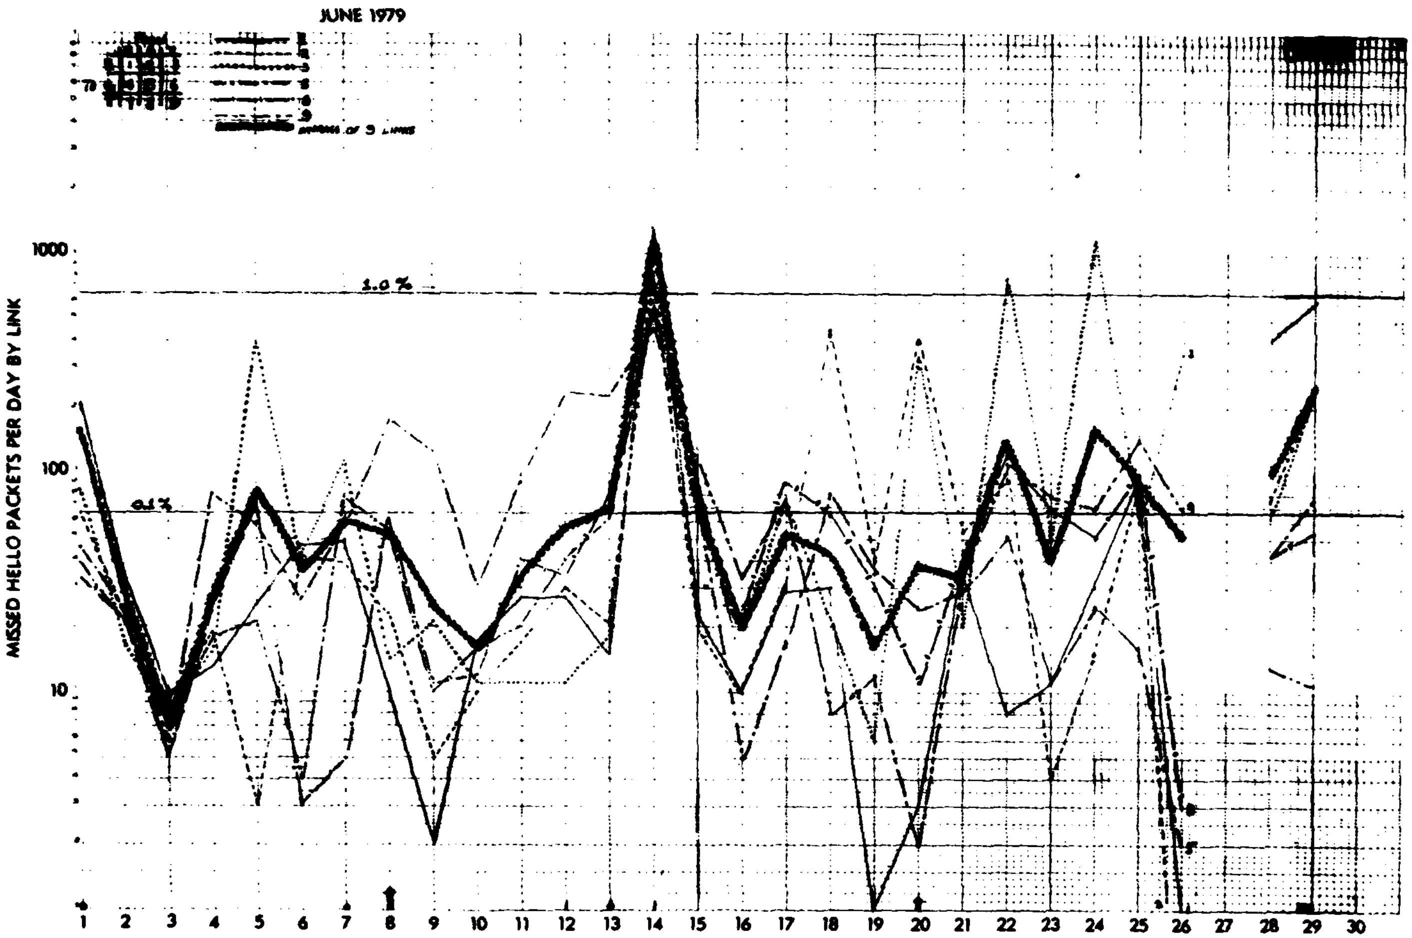

Figure 6-10. Missed Hello Packets Per Day By Link; June 1979

Table 6-10. Summary of SATNET Data for July 1979
Missed Hello Packets

Day / Link	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
1	1	136	306	38	1	1	0	0		1	17	568	625	6	169	204	2	282		38		20	14	5	4	0	3		0	54	49
2	0	299	206	26	12	1	0	0		15	13	175	210	3	47	137	47	18	147	29		1901	275	66	31	111	662		1	95	120
3	0	118	64	26	3	1	3	5		15	11	236	270	14	135	2	2	11	169	24		28	62	44	2	1	15		-	85	112
4	2	221	336	12	5	4	1	5		16	33	311	182	0	69	228	30	68	172	20		17	16	30	3	2	0		1	204	106
5	2	4	295	15	53	56	28	22		46	46	50	206	6	10	26	61	32	26	23		579	60	43	73	113	63		14	70	62
6	2	9	92	2	6	17	16	13		36	15	126	88	16	10	18	11	14	8	11		94	68	60	16	19	26		5	16	66
7	0	226	106	14	1	2	0	4		1	31	239	277	1	69	119	6	0	149	20		4	1	91	3	2	19		0	29	9
8	0	29	103	5	21	0	1	1		16	4	78	249	5	13	45	45	6	9	8		29	9	110	11	65	31		3	5	7
9	1	29	152	6	9	5	3	7		0	0	110	101	2	35	9	2	3	12	0		12	33	103	7	8	26		2	5	8
Daily Average	1	119	195	16	13	11	6	6		16	19	210	250	6	54	98	23	48	845	20		298	60	61	17	36	94		3	63	60
Outage Correction Factor	1.53	2.81	1.06	1.06	-	-	2.12	1.95		1.30	-	-	1.20	1.09		1.45	1.45	1.02	-	-		-	-	-	1.10	1.33	1.30		1.00	1.03	-

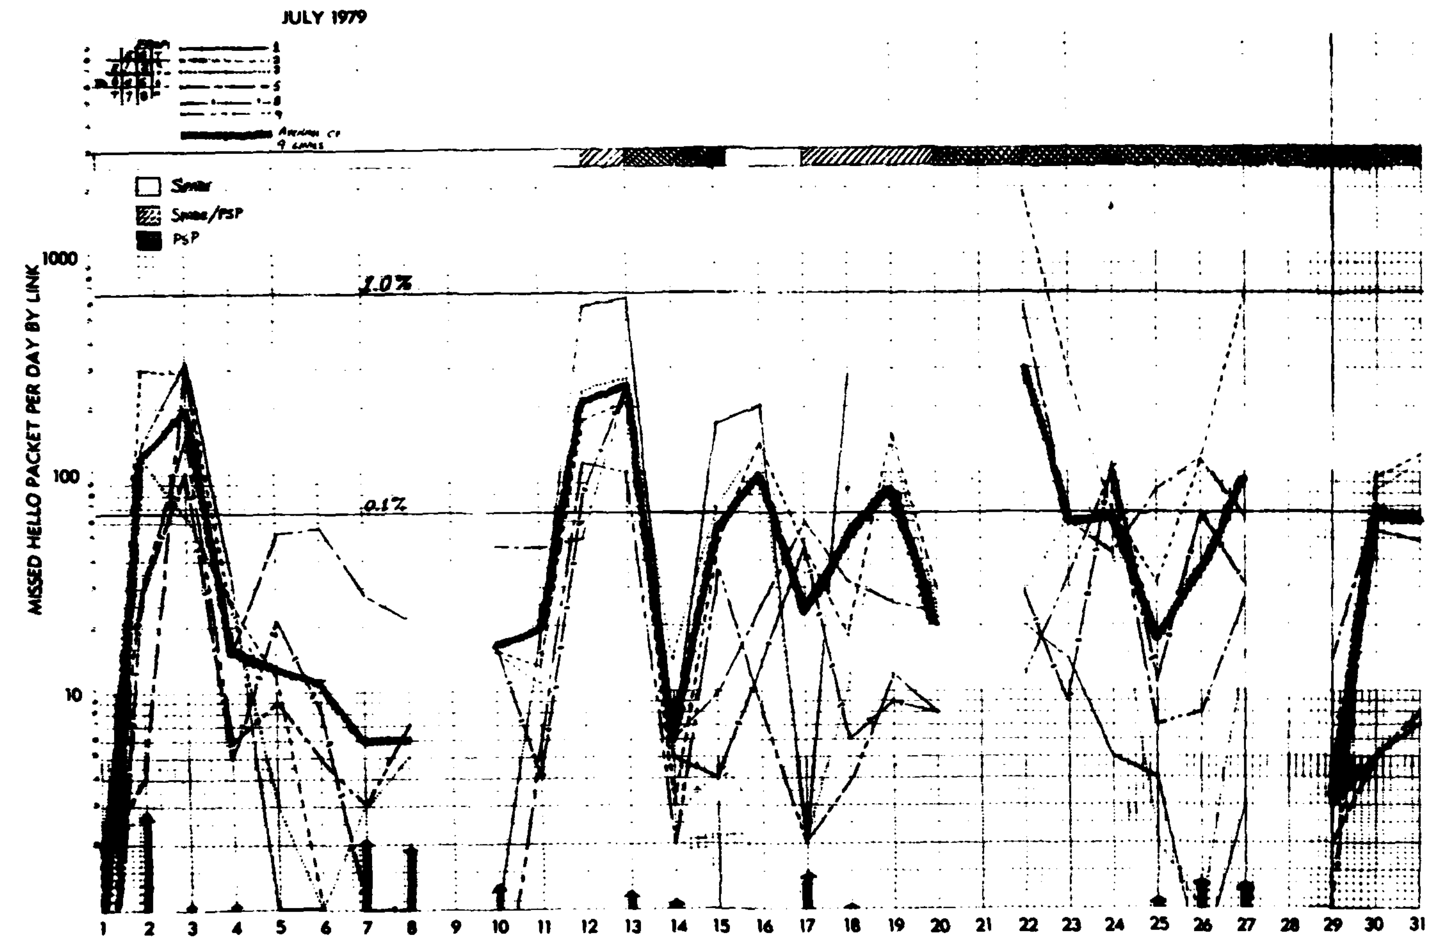

Figure 6-11. Missed Hello Packets Per Day By Link; July 1979

to reveal any abnormality in the operating conditions at Goonhilly. Finally in early April 1979, a power level adjustment was made plus extensive checks of the SPADE equipment. This action resulted in more stable performance during the following weeks.

Along with the data review, an approximate model has been developed for SATNET performance which allows average channel bit error rate to be inferred from the missed-packet percentages. This model indicates that the missed-packet behavior (except for some of the extremes) is normal and to be expected. A trend not shown in the data presented here is a tendency for the missed packets to occur during that period of each day when the SPADE transponder in the INTELSAT IV-A satellite is carrying other (mostly voice) traffic. Missed packets are relatively rare in the early morning and late evening hours. There is clearly some correlation, within each day, with the loading of the transponder. This "loading" has also increased with time and more carriers are utilizing the transponder now (Spring 1979) than they were a year ago, say in the Spring of 1978. The experience with SATNET does seem to reflect both this growth in traffic, and the diurnal variations in transponder loading.

7. SATELLITE COMMUNICATIONS SYSTEM TRADEOFF STUDIES

7.1 INTRODUCTION

In addition to the other activities at COMSAT related to the Atlantic Packet Satellite program, some effort has been devoted to a study of future applications of packet transmission by satellite. This study has focused on two questions:

a. how can future communication satellite systems benefit from the packet transmission format?
b. how might packet transmission formats impact the architecture of future communication satellite systems?

These questions are closely related; the first seeks to determine the benefit to be derived by utilizing packet transmission techniques, as opposed to other techniques, for transmitting digital data via communications satellites. The second question asks how future communication satellite systems might be different, given that packet-transmission techniques are being used. Such differences could arise in the terrestrial part of the satellite system, i.e., the earth stations, or in the spacecraft itself.

To address these questions, it is necessary to hypothesize a network, and to consider various quantities and categories of data to be handled by that network. The latter area is important since packet-transmission concepts are clearly not applicable to all types of data transmission requirements (digital TV for example). This description of the type of network being considered and possible volume, type, and mixture of traffic to be handled are outlined in Section 7.2. These two items together, a traffic forecast, and a hypothesized network serve as general guidelines for the study.

In Section 7.3, the implications of the traffic scenarios on the gross characteristics of a communications satellite system are determined, and a range of system concepts are defined that match these characteristics. In Section 7.4, future satellite system concepts and architectures are described that take advantage of the packet transmission concept and exploit the flexibility and adaptability that is inherent to this type of data transmission. A comparison of the relative cost effectiveness of circuit switching and packet transmission is given in Section 7.5.

7.2 NETWORK AND TRAFFIC SCENARIOS

One of the difficulties in judging the impact of packet transmission formats on communication satellite networks is that the answers obtained are highly dependent upon the assumed network topology and upon the traffic to be carried by the network. Satellite communications system architecture for providing packet communication services is highly dependent on the level, distribution, and type of traffic to be handled by such a system. Accurate forecasts of future user requirements are difficult to obtain although two such projections are available. In a report by Ovum L. [9], estimates are given for future computer communications traffic in a few countries. For example, the estimates for annual traffic in the United States are 2,480 and 7,000 millions of kilopackets in 1980 and 1985, respectively. These estimates correspond to an annual growth rate of 23 percent. If this growth is extended still further into the future, the estimated annual computer communications traffic in the U.S. is 19,700 millions of kilopackets per year, in year 1990. These traffic estimates are based on the assumption that an average packet will contain 500 bits of data. Assuming a uniform distribution for this traffic results in an average data rate of 312 Mbit/s in the year 1990. These are estimates of the United States' computer communications traffic only and do not include any voice-data traffic. The Department of Defense (DOD) has estimated future data and voice traffic requirements for the late 1980s and early 1990s as 36 Mbit/s of data traffic combined with 2700 Erlangs of voice traffic [2]. Assuming an average of 32 Kbit/s for each voice channel results in an integrated voice and data traffic requirement of 122 Mbit/s.

Since these traffic estimates for 10 years in the future are subject to significant variations, they can be used only as a baseline from which a range of user requirement scenarios can be defined. This range of possible scenarios can then be used to bound the requirements on future digital satellite communications systems for packet switching applications.

Three packet traffic scenarios are defined in terms of the total network traffic in kilopackets per year, the total number of earth stations participating in the network, and the mixture of traffic between data and voice. Based on the total traffic level, the three scenarios are designated as nominal, low, and high.

Nominal Scenario. In this scenario, it is assumed that the total DOD "packet" traffic requirement will reach a level of 8×10^9 kilopackets per year (125 Mbit/s) by 1990. It is assumed that a total of 1,000 earth stations will access the satellite, and that 20 of these stations will carry 30 percent of the traffic, another 80 stations will carry another 30 percent, and 900 earth stations will carry the remainder of the traffic (40 percent), with the traffic in each group divided equally among the corresponding stations.

Low Scenario. This scenario assumes that the total packet traffic does not grow as rapidly as in the nominal scenario so that by 1990 the total will be 8×10^8 kilopackets per year. It is assumed that the total number of earth stations reaches a level of 100, and that 30 percent of this traffic will be carried by 5 of the stations (on an equal basis), another 30 percent by another 15 stations, and the other 40 percent by the remainder of the stations (80) on an equal basis. In this scenario, it is assumed that stream (voice) traffic does not grow rapidly and, hence, most of the traffic will be data.

High Scenario. In the high scenario, it is assumed that packet satellites will carry a significant amount of voice (stream) traffic and the total traffic by 1990 reaches 8×10^{11} kilopackets per year. It is assumed that the number of stations reaches 10,000, and that 10 of these stations will carry 25 percent of the traffic (on an equal basis), 100 of the stations will carry another 25 percent of the traffic, another 1,000 stations sill carry another 25 percent, and the rest of the stations (8,890) will carry the remaining 25 percent of the traffic.

These projections of future traffic are summarized in Table 7-1 where an additional breakdown is included within each category giving several mixes of voice and data traffic. The scenarios cover a range of four orders of magnitude from the low to the high estimates of annual traffic volume. This extremely wide range of possible requirements obviously will have a large impact upon the communications satellite system that provides such a service. These impacts are examined in the next subsection.

Table 7-1. Summary of "Packet" Traffic Scenarios

Scenarios	Level of Traffic in Packets/Year (P)	Number of Earth Stations (N_{ES})	Mix of Traffic	
			Percent Data	Percent Voice
Nominal				
Ia	8×10^{12}	1,000	50	50
Ib	8×10^{12}	1,000	10	90
Ic	8×10^{12}	1,000	1	99
Low				
IIa	8×10^{10}	100	100	0
IIb	8×10^{10}	100	50	50
High				
IIIa	8×10^{14}	10,000	10	90
IIIb	8×10^{14}	10,000	1	99
IIIc	8×10^{14}	10,000	0.1	99.9

7.3 IMPLICATIONS OF THE PROJECTED TRAFFIC ON SPACECRAFT

7.3.1 Required Spacecraft Capacity

Packet satellite networks are to be considered that are characterized by the number of users served by the network, N_u, the number of earth stations in the network, N_{ES}, and the total annual network traffic, P (packets per year). For now, it will be assumed that "earth stations" and "users" denote the same thing so that $N_u = N_{ES}$.

The peak system data rate $\hat{R}_b$ can be obtained from P if an average of 500 bits/packet are assumed as

$$\hat{R}_b = \alpha \bar{R}_b = \alpha \frac{P \text{ packets/yr} \cdot 500 \text{ bits/packet}}{365 \cdot 24 \cdot 3600 \text{ s/yr}}$$

$$= 1.58 \times 10^{-5} \alpha P \qquad (7\text{-}1)$$

where

$\hat{R}_b$ = peak data rate in bits/s

α = peak-to-average ratio of the traffic

$\bar{R}_b$ = "average" data rate in bits/s

P = total annual network traffic in packets

From Table 7-1, values of P range from 8×10^{10} to 8×10^{14}. Assuming $\alpha = 4$ gives peak data rates ranging from 5 Mbit/s to 50,000 Mbit/s. Between these two extremes, a "nominal" system is postulated with $P = 8 \times 10^{12}$ ($\alpha = 4$) giving a peak data rate $\hat{R}_b$ = 500 Mbit/s. The implications of this range of traffic will be determined in the following.

Transmission of the peak data rates, $\hat{R}_b$, through a single spacecraft implies a total spectrum $\hat{B}$

$$\hat{B} = \frac{\hat{R}_b \text{ bits/s}}{\eta \text{ bits/s/Hz}} \tag{7-2}$$

where η is the bandwidth efficiency in bits/s per Hz. This efficiency would typically range from 0.5 for coded PSK operation to 1.5 for uncoded QPSK modulation. For bandwidth efficiencies of 0.5 and 1.5 the total satellite bandwidth required for the network can be estimated as shown in Table 7-2.

The three scenarios for traffic, i.e., the low, nominal, and high traffic levels call for considerably different spacecraft configurations to satisfy the requirements. These requirements can be summarized as follows:

a. The low traffic level can probably be satisfied in a fraction of a single transponder of a satellite that provides wide-area (i.e., CONUS) coverage. Such service could be provided at 6/4,14/11, or possibly 30/20 GHz.

b. The nominal traffic level requires approximately twenty 40-MHz transponders which would correspond to all, or a significant part of a satellite of the complexity of INTELSAT V. Operation in more than one of the frequency bands, 6/4, 14/11, or 30/20 GHz may be required.

c. The high traffic level imposes severe (possibly unrealistic) requirements on the space segment and may have to be reduced somewhat; however, as an upper bound, the high secnario would require several thousand 40-MHz transponders. The frequency reuse required would imply up to one hundred individual antenna beams formed at the satellite. Even if reduced 10-to-1, this scenario may require multiple satellites, or a very complex space station.

Table 7-2. Total Bandwidth Required by Packet Satellite Network

	Scenario					
	Low		Nominal		High	
$\hat{R}_b$ (Mbit/s)	5.0		500		50,000	
Bandwidth Efficiency	$\eta = 1.5$	$\eta = 0.5$	$\eta = 1.5$	$\eta = 0.5$	$\eta = 1.5$	$\eta = 0.5$
Bandwidth Required (MHz)	3.33	10.0	333	1000	33,333	100,000
Spectrum Use at 6/4 or 14/11 GHz (500 MHz available)	0.006	0.02	0.67	2.0	66.6	200
Spectrum Use at 30/20 GHz (2500 MHz available)	0.0013	0.004	0.133	0.4	13.3	40

7.3.2 Communications System Design Implications

Experiments with packet-switch transmission via communications satellites have proved the concept at low-data rate ($\leq$64 kbit/s) between a limited number of earth stations. Efficient demand-assignment multiple-access techniques have been developed during these experiments which consider such factors as priority and maximum delay time in making assignments in a variable TDMA message frame. These techniques are directly applicable to higher transmission rates such as 1.5 Mbit/s or even higher.

The future traffic projections for packet-satellite systems given earlier in this section, cover an an extremely wide range so that satellite architectures to serve these future needs could vary significantly. For example, to meet the most moderate traffic projection, the packet satellite network could still be served by a small fraction of an existing spacecraft.* (Possibly a complete transponder or a large fraction of a complete transponder.) As the higher traffic projections are considered, more and more of the spacecraft could be utilized by the packet satellite network until, for the most optimistic estimates of future traffic growth, very complex spacecraft (or even multiple spacecraft) are needed which reuse the available frequency spectrum many times to provide all of the necessary communications capacity.

*As a point of reference, SATNET uses one 64-kbit/s channel in the SPADE transponder of the INTELSAT IV-A spacecraft. This present use thus consumes 0.125 percent of the bandwidth and approximately 0.3 percent of the power of this single 40-MHz transponder which is, itself, only about 5 percent of the total communications capability of the spacecraft (i.e., the INTELSAT IV-A spacecraft contains twenty 40-MHz transponders).

These projections for complex spacecraft follow from some rather simple approximate calculations. Consider for example, a typical and not unreasonable, forecast of a future network that carries 8×10^{12} packets per year. Assuming a peak-to-average ratio of four, such a projection for traffic gives a peak rate of approximately 10^6 packets/s. If an average packet is assumed to contain 500 information bits, this average packet rate can be converted to an information bit rate of 500 Mbit/s. Recognizing that a significant fraction of the packet transmission is devoted to synchronization preambles, the peak transmission rate could easily double (i.e., the total transmission is half data and half "overhead"), giving 1000 Mbit/s. Such a rate would require 17 40-MHz transponders assuming QPSK transmission which achieves approximately 1.5 bits/Hz. On an average basis (i.e., if the transmissions could be spread uniformly through each day) one-fourth of this capacity would be required.

The point of calculations such as these is that future applications of communications satellites for packet transmission could evolve into complex spacecraft which serve networks with extremely large capacities. Such spacecraft must reuse the available frequency spectrum, either through multiple disjoint antenna coverage areas or dual polarization, or both, and could also utilize more than one of the frequency bands which have been allocated for satellite communications. These particular technological areas provide the potential for increased bandwidth capability in future spacecraft, but this increase is generally accompanied by attendant problems of maintaining the required connectivity between users; such problems become particularly acute in complex networks serving many earth stations. Concepts for satellite-switched time-dividion multiple-access (SS-TDMA), and intersatellite links (ISLs) interconnecting multiple, relatively closely spaced satellites serving a given region, have emerged as possible solutions to these problems of connectivity and proliferation of earth stations. These concepts have been considered for future satellite systems that carry multichannel telephony and television transmissions.

To date, future satellite architectures have been driven almost exclusively by the steady growth in "conventional" communications traffic. International traffic of this type will undoubtedly continue to emanate from large, expensive earth-station instllations such as the 30-m "Standard-A" station used in the INTELSAT system. The trend in domestic satellite communications systems on the other hand, is to "move" the earth stations nearer to the ultimate users of the communication service (ultimately collocating the earth station with the user). Such a trend calls for smaller earth stations. These domestic networks, which will eventually be connected either indirectly or directly to the international satellite system (INTELSAT), are being designed to serve a variety of future requirements for digital data transmission services. Such services range from "real-time" transmission of digital voice and either conventional or slow-scan television, through the "almost-real-time" requirements such as interactive data transmission, facsimile transmission, and some computer-to-computer data transmission applications to finally, the less "delay sensitive" applications such as electronic mail, and bulk data transfer between computers.

For all commercial applications of communications satellites, a major issue is the efficient utilization of the resources of the spacecraft, specifically the power and bandwidth in the various transponders. In the past, less capable satellites have required the use of large earth stations to produce strong enough signals at the input to the satellite, and to receive sufficient signal strength from the relatively weak satellite transmitters, to yield an overall signal-to-noise ratio that would allow effective use of the overall bandwidth. These early systems were "power limited" and generally used the available bandwidth inefficiently. The situation has changed with current satellites with the advent of directive antennas on the spacecraft and higher power transmitters. Todays satellites are generally capable of operating at the cross-over between power- and bandwidth-limited operation when moderately large earth stations are used, and modulation techniques such as quaternary phase-shift keying (QPSK) are used on the communication links. In the future, continued advances in spacecraft antenna design will provide coverage that is tailored to the needs of particular communication networks. One of the benefits of this technology is a lowering of the cross-over between bandwidth and power-limited operation to smaller earth stations. It is possible in future spacecraft, therefore, to utilize the spacecraft efficiently, in the sense that a bandwidth efficient modulation is used (i.e., a large value of η bits/Hz) with relatively small earth stations. Minimizing the size, and hence cost of earth stations assumes major importance in networks with a large number of earth stations.

A second major issue in efficient utilization of the space segment is the choice of multiple-access technique. These techniques allow multiple earth stations to share a particular part of the spacecraft resource, usually the e.i.r.p. and bandwidth of a particular transponder. The conventional multiple-access techniques, namely FDMA, TDMA, and CDMA (code-division multiple access), are well known and will not be elaborated (except to note that CDMA is relatively inefficient except in extremely power- or interference-limited situations, so that it has very limited application in commercial communications satellite systems). Both FDMA and TDMA are used in commercial satellite systems (TDMA will soon be used), and both are presently utilized on a pre-assigned basis. The "user" in this case bundels his traffic into large groups (of telephone channels, for example) and delivers this builk traffic to the earth station. This bundeling can be done with some sophistication (for example, the statistics of certain sources of traffic can be exploited such as is done in TASI or DSI for processing speech), but generally the assignment of satellite resources is relatively static and fixed. Users are pre-assigned a segment of the satellite, i.e., a certain fraction of its power and bandwidth, or a certain number of channels, and these assignments are changed only rarely. The service is paid for typically on a monthly basis, and such pre-assigned service is generally cost-effective for large users, i.e., those with a large amount of traffic.

It became clear over ten years ago that the pre-assigned use of satellite communications capacity was definitely not cost-effective for an important and growing class of users. These were the new users who had relatively little traffic and wanted to join the network. Although these "thin routes" would reasonably be expected to grow with time, it was not cost effective to install a large earth station, equipped with an FDM multiplexer capable of handling 12 voice channels (the minimum increment available in INTELSAT) and then use such a station to handle only a few telephone calls per day (say a few hours of active speech traffic per day). In this case, users paid for full-time use of channels in the spacecraft but these channels were underutilized. The SPADE system was developed specifically for these "thin-route" users. The SPADE* system operates in a global transponder and provides for demand-assignment of 800

*Note that the same transmission format as used in SPADE, i.e., 64-kbit/s QPSK, can be used without the demand assignment feature. Such systems are usually referred to simply as single-channel-per-carrier (SCPC) systems.

45-kHz frequency slots which fill this transponder. One of the 64-kbit/s channels in the transponder serves as a common signaling channel and is used for distributed control of the frequency channels. Each user has access to a certain time segment of the 64-kbit/s TDMA control channel, and in this brief period, he makes requests for vacant frequency slots (which are known to all users). All users monitor the channel-request transmissions and are aware of the requests of other users, and the successful assignment of pairs of frequency slots. In the present SPADE system, the common signaling channel is operated as a 64-kbit/s BPSK TDMA channel with a 50-ms frame time. This frame is divided into 50 1-ms slots so that 50 SPADE users (i.e., 50 earth stations) can participate in the common signaling channel. Note that the actual establishment of a "call" generally consumes several back-and-forth transmissions between the transmitting and receiving earth stations via the signaling channel. Each must traverse the earth-satellite-earth path with its 0.25-s delay. Generally, a total delay of 5-10 seconds is required to establish a connection. Once a channel pair is assigned, the parties retain the channels for the duration of the call (i.e., generally for several minutes). Voice activity detectors are used in SPADE (and most SCPC systems, as well) to turn off a carrier during the period when one speaker is silent. This on-off switching follows a typical pattern of a few seconds on, and several seconds off. Several features of the SPADE demand-assignment system are worth emphasizing since some parallels will arise when packet-transmission systems are discussed. These can be summarized as follows:

a. The common-signaling channel is a global TDMA channel with sufficient time slots to allow 50 users to participate in the network in an ordered or noninterfering basis.

b. Control is "distributed" in the sense that all earth stations have equal precedence in requesting channels, and in making assignments.

c. The "set-up" time for a channel is relatively rapid for the type of traffic being handled (i.e., speech). However, for other demand-assignment applications, such as those that arise in packet-transmission applications, the speed of the channel request and assignment functions used in SPADE would be judged to be much too slow.

With this background, a key question is how packet transmission formats might impact future satellite communication systems. These impacts can focus upon the earth stations in the network, upon the spacecraft itself, and upon the control subsystem that dynamically allocates and reallocates spacecraft resources among the various network "users" (who in the following will be taken as synonomous with the earth stations). To provide as realistic a framework as possible, a CONUS (Continental United States) setting will be used with a single spacecraft serving a diverse set of earth stations that are more or less uniformly distributed over the coverage area. A multibeam spacecraft is assumed which provides the potential for frequency reuse. The volume of total network traffic is assumed to be such that reuse is needed to satisfy the network requirements. Implicit to the multibeam configuration is the increased up-link sensitivity, and down-link e.i.r.p. of the spacecraft which favors small, and hence inexpensive, earth stations. Such a trend is necessary in networks with hundreds or even thousands of such stations.

In this hypothetical future environment, some interesting questions arise about the future directions of the technology of future spacecraft and earth stations. The past trend of relatively simple reliable spacecraft has relegated most of the complexity to the participating earth stations. This trend is presently changing and such changes will undoubtedly be accelerated by the availability of the space transportation system (STS or shuttle) which will remove many of the constraints which placed inherent limitations on spacecraft designs in the past. Removal of these restrictions, coupled with the unique characteristics of packet-transmission networks, and the packet formats themselves, can lead to some interesting speculation about future spacecraft which might serve these networks in the future.

The intent here is not to describe and analyze particular networks, or even to design satellites to serve hypothetical networks. Such steps would not only be premature, but would be well beyond the scope of this limited effort. The following sections will deal only with concepts that might find application in future packet-switched satellites, and where possible, point out the advantages and disadvantages of these concepts.

7.4 FUTURE PACKET SATELLITE CONCEPTS

Present trends in communications satellites emphasize higher capacity spacecraft that utilize one or more of the frequency bands allocated to satellite communications. By means of directive spacecraft antennas or polarization diversity, or both, a large portion of the allocated spectrum can be reused giving the potential for significantly more digital transmission capacity than is available with the basic allocation. This trend toward frequency reuse results in complex spacecraft antenna structures, consisting for example of large reflectors and multiple feed elements, which can form separate isolated beams directed at specific regions of the earth's surface. Upon reception at the spacecraft, these narrow coverage areas provide high sensitivity (large gain) to signals transmitted from the earth's surface.

Likewise, transmission from the spacecraft is able to concentrate high effective radiated powers (e.i.r.p.) into the covered region. These high-gain spacecraft antennas allow smaller earth stations to be used in networks that contain a large number of dispersed users, each with a relatively low average volume of traffic.

If these concepts are applied to the Continental United States (CONUS), the antenna beam coverage from the spacecraft would appear as shown in Figure 7-1. The maximum simultaneous frequency use from such a spacecraft would be:

$$F = W_A \times N_B \times f_B \tag{7-3}$$

Where W_A is the allocation (e.g., 500 MHz), N_B is the number of simultaneous up- and down-link beams that can be formed at the spacecraft, and f_B is a factor that may be less than one to account for the fact that touching spot beams cannot support simultaneous transmissions on the same frequency at the same instant of time. Generally, the number of simultaneous beams formed at the spacecraft (N_B) could be less than the number of contiguous spots, N_s, that are needed to cover all of CONUS.

A second concept that has been proposed for future communications satellites is the "processing" repeater. With few exceptions, current and planned satellites act simply to frequency translate, amplify, and retransmit the signals that appear at their inputs. Any noise that is present at the input to the transponder is amplified and radiated along with the desired input signals. A processing repeater, on the other hand, would demodulate incoming signals and then remodulate a separate downlink transmitter with the recovered data. Although the processing repeater results in a more complex spacecraft design, this increased complexity may be justified in networks with many relatively small earth stations.

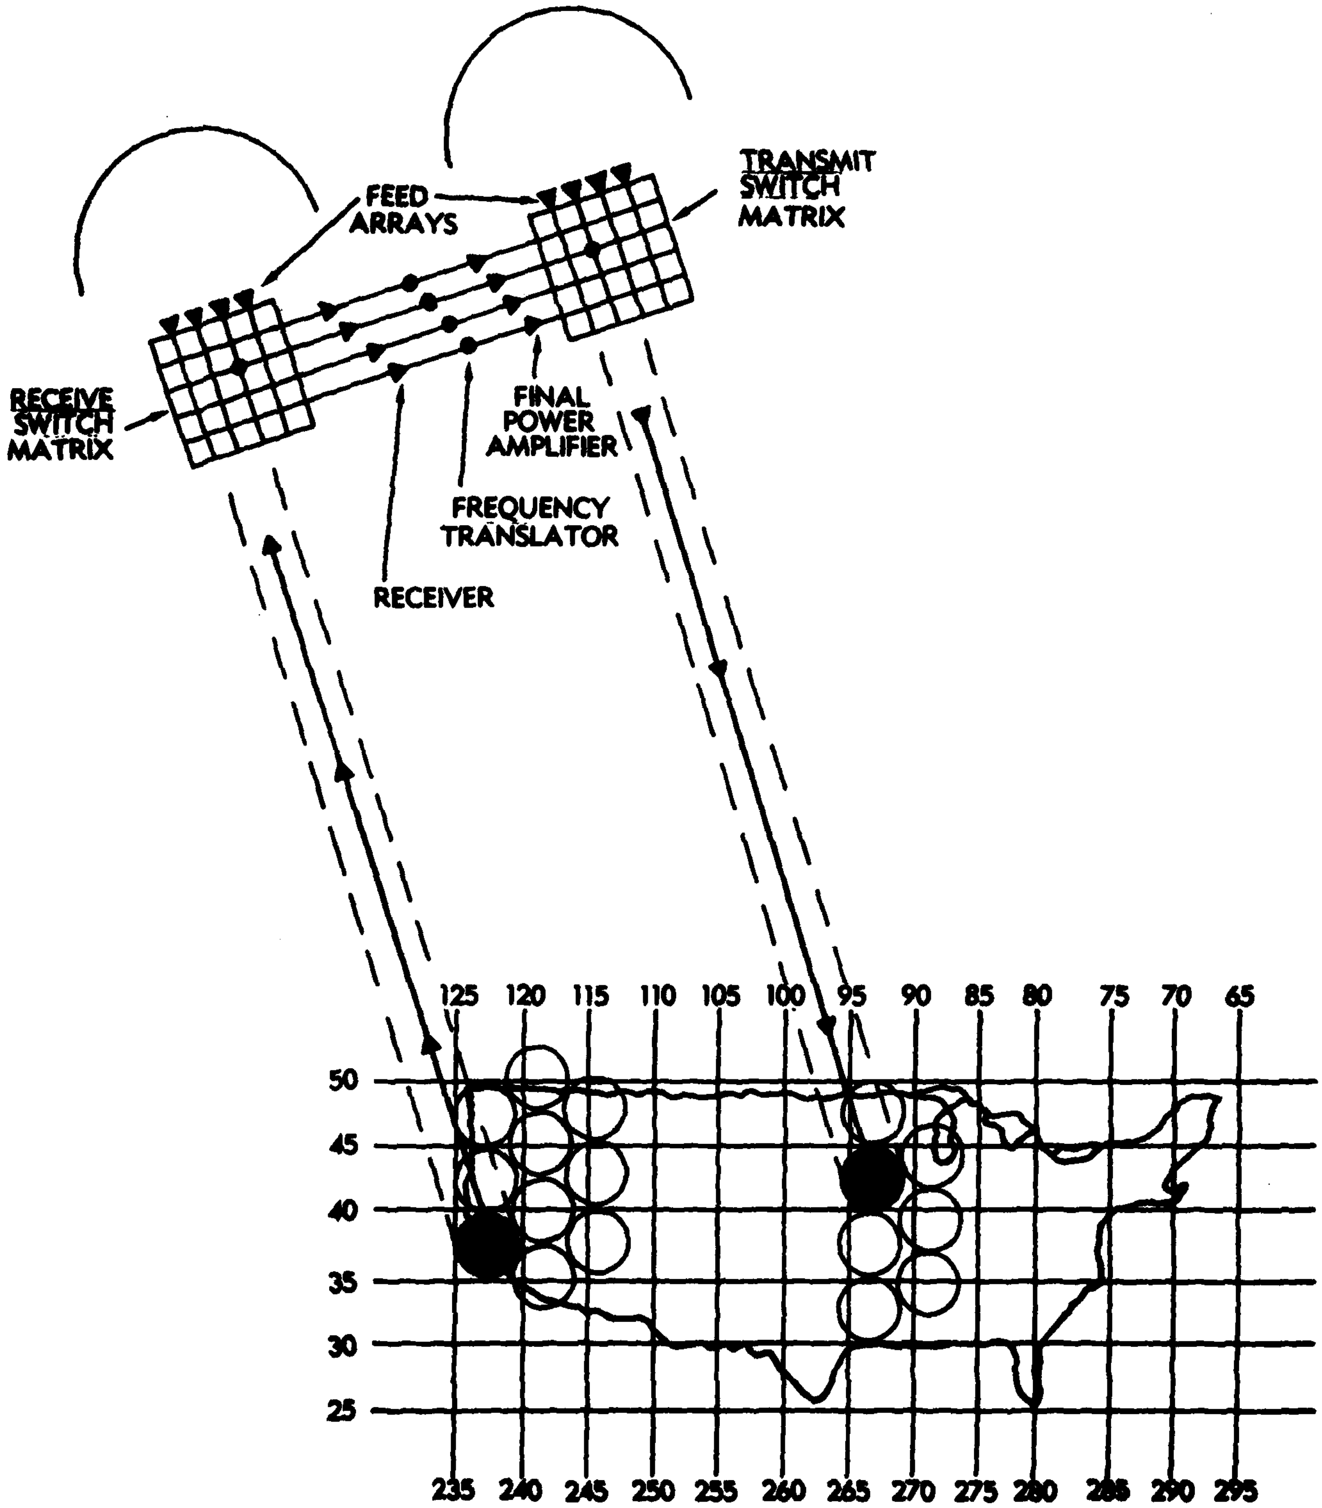

Figure 7-1. Multibeam Satellite Concept

With packet transmission formats, it is possible to envision different options for onboard processing where the more complex techniques include on-board storage of signals. If we consider a standard packet as shown in Figure 7-2, then some of the options for on-board processing can be listed briefly as shown in Table 7-3. These concepts range from the conventional "straight-through" transponder (A), to a conventional demodulation/remodulation repeater (D), which in turn is followed by more complex alternatives for processing, some of which are unique to the packet transmission format. The final concept (I), which includes on-board storage, could include the queueing of packets at the spacecraft to await down-link transmission. With such storage available, the spacecraft ceases to be a "real time" transmission medium and can be envisioned more as a node in a packet switched network. The complexity implied by some of this processing may make it impractical to consider it for inclusion in a spacecraft. As concepts, however, the alternatives listed in Table 7-3 are worth listing. The conclusion is that "processing satellite" can have a much broader interpretation when "packetized" digital transmission is being considered and this is especially true for concepts that provide for temporary storage of data at the spacecraft.

7.5 ANALYSIS OF NETWORK CONFIGURATIONS

7.5.1 Model and Assumptions

In this section, we consider in more detail the tradeoffs associated with serving a network of users by a communications satellite system that utilizes one of several alternative

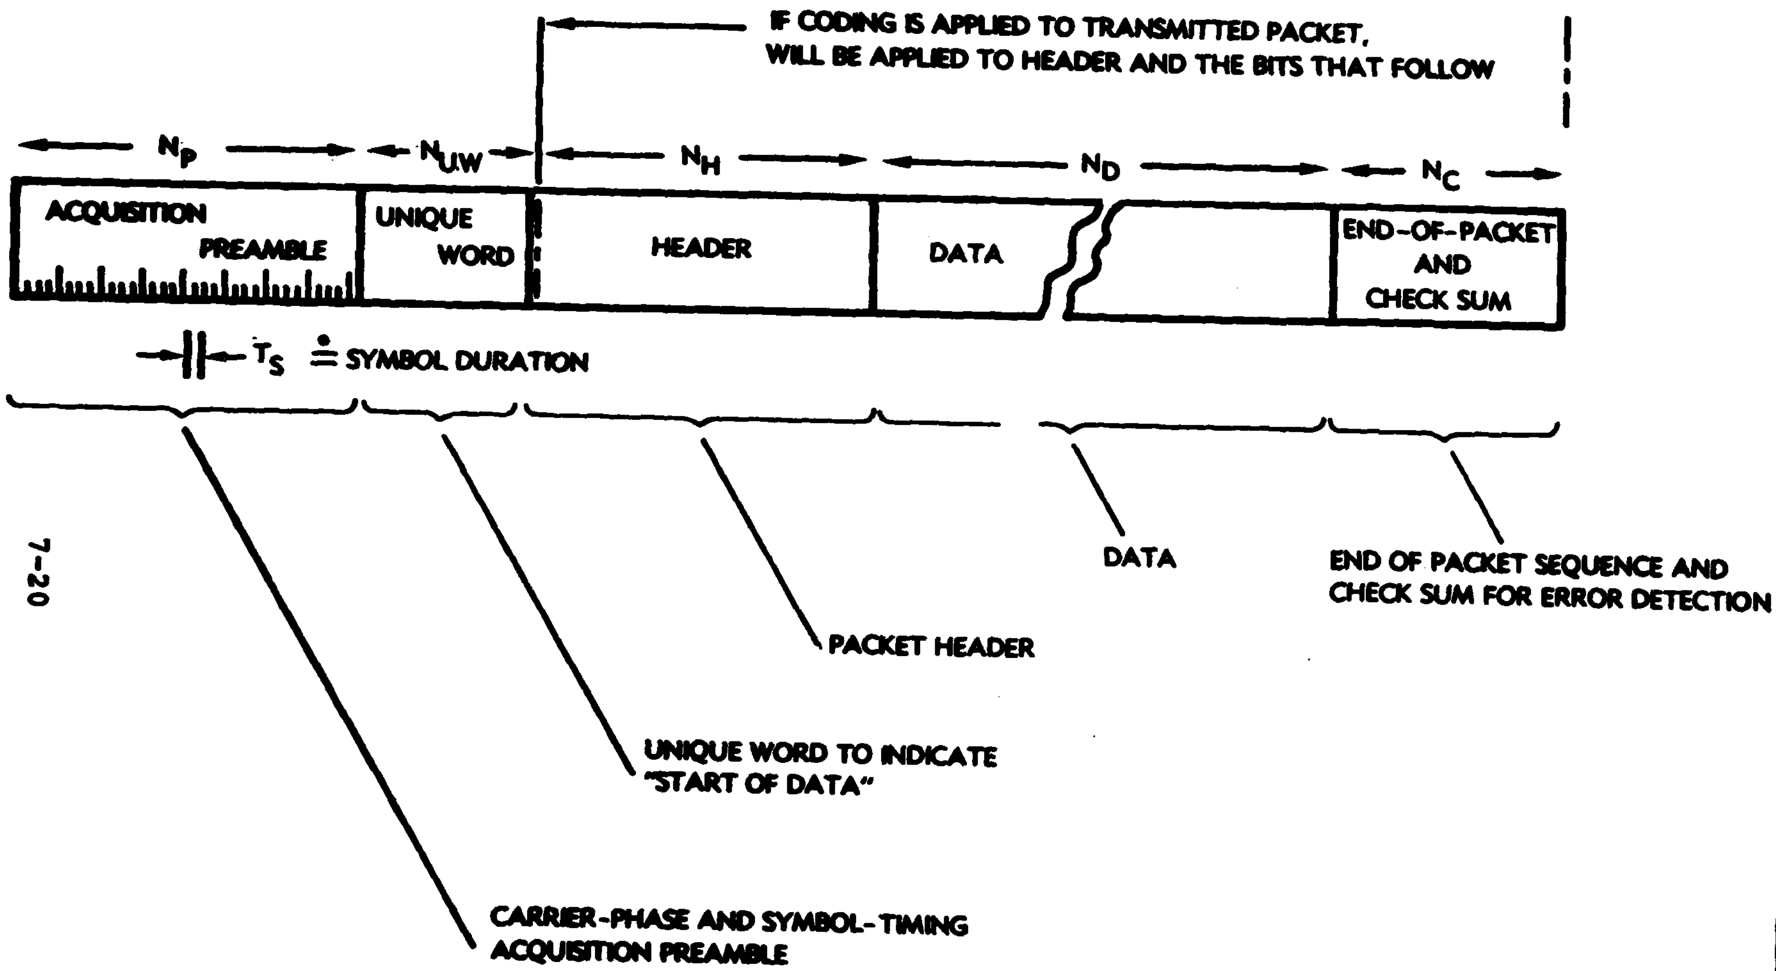

Figure 7-2. The Standard "Packet"

Table 7-3. Options for Processing of Packet Transmissions in the Spacecraft

	Is Preamble Detected/ Acquired?	Is Unique Word Detected?	Is the Header Demodulated?	Is the Data and Postamble Demodulated?	Is Decoding for Error Correction Performed?	Detect E.O.P. and Compare Check Sum ?	Process the Header ?	Store All or Part of the Packet?	Comments
A	No	No	No	No	No	No	No	No	Conventional "STRAIGHT-THROUGH" TRANSPONDER
B	No	Yes[a]	No	No	No	No	No	No	Could be used to transfer time to spacecraft
C	Yes	Yes	No	No	No	No	No	No	
D	Yes	Yes	Yes	Yes	No	No	No	No	Conventional Demodulation/ Remodulation Repeater
E	Yes	Yes	Yes	Yes	No	No	Yes	No	Header information could be used for down-link beam switching
F	Yes	Yes	Yes	Yes	Yes	No	No	No	Error correction at the satellite (down-link packets shorter than up-link)
G	Yes	Yes	Yes	Yes	Yes	Yes	No	No	Error detection at the spacecraft, packets with errors could be discarded
H	Yes	Yes	Yes	Yes	No[b]	Yes	No	No	
I	Yes	Yes	Yes	Yes	No/Yes	No/Yes	Yes	Yes	Processing of Header information combined with on-board storage and reformatting of packets

[a]Noncoherent.
[b]None included.

techniques for allocating the spacecraft resource. The analysis will apply strictly to the "nominal" scenario which was given in Subsection 7.2. The total traffic volume per year, $\hat{P} = 8 \times 10^{12}$ packets, and the total number of users, $\hat{N}_u = 1000$, will be kept constant for the analysis. Certain other assumptions in the scenario will be treated as variables, however. First, we will consider the total population of users to be divided into three classes; "High" rate, "Medium" rate, and "Low" rate users. The total volume of traffic is divided such that the high rate users generate a fraction P_H of the total traffic, the medium users P_M, and the low rate users a fraction P_L (where $P_H + P_M + P_L = 1$). Each user is further identified by a utilization factor u_i, indicating the fraction of the 24 hours per day that he is actually using the network, and by an activity factor a_i which indicates the fraction of time that the user actually generates data, given that he is using the network. For example, a user who uses the network for an average of 3 hours per day would have a utilization factor $u_i = 1/8$. When using the network, the user could generate data sporadically with a ratio of active data generation time to total time of five percent. In this case, the activity factor of the user would be 0.05.

The user population is assumed to be subdivided so that a fraction f_H, f_M, and f_L of the 1000 users are in the high, medium, and low categories, respectively. For example, if 2 percent of the users are in the high category, then the network would have 20 high rate users. The number of users per category will be denoted n_H, n_M, and n_L. These variables pertaining to the user population are summarized in Table 7-4.

Table 7-4. Summary of Network Variables
Traffic Volume, $\hat{P} = 8 \times 10^{12}$ packets/year
Total Number of Users, $\hat{N}_u = 1000$

User Class	Fraction of Users	Number of Users	Fraction of Total Traffic	Utilization Factor	Activity Factor
High	f_H	$n_H = f_H\hat{N}_u$	P_H	u_H	a_H
Medium	f_M	$n_M = f_M\hat{N}_u$	P_M	u_M	a_M
Low	f_L	$n_L = f_L\hat{N}_u$	P_L	u_L	a_L

NOTE: $f_H + f_M + f_L = 1.$
$P_H + P_M + P_L = 1.$

We must next make some assumptions about the traffic matrix that exists between the different users. This matrix will have the general form shown in Figure 7-3. It is assumed that the traffic matrix is ordered such that high-rate users appear first ($i = 1, n_H$), medium-rate users appear next, and low-rate users appear last. The matrix can then be partitioned to have the mutual traffic between high-rate users in the upper left corner, and the mutual traffic between low-rate users in the lower right corner. It is assumed that all traffic is restricted to mutual traffic between users of the same class. If we assume that a "packet" contains 500 information bits, and that the traffic submatricese are filled with identical elements (full connectivity required, uniform distribution of traffic), then several traffic measures can be computed as follows.

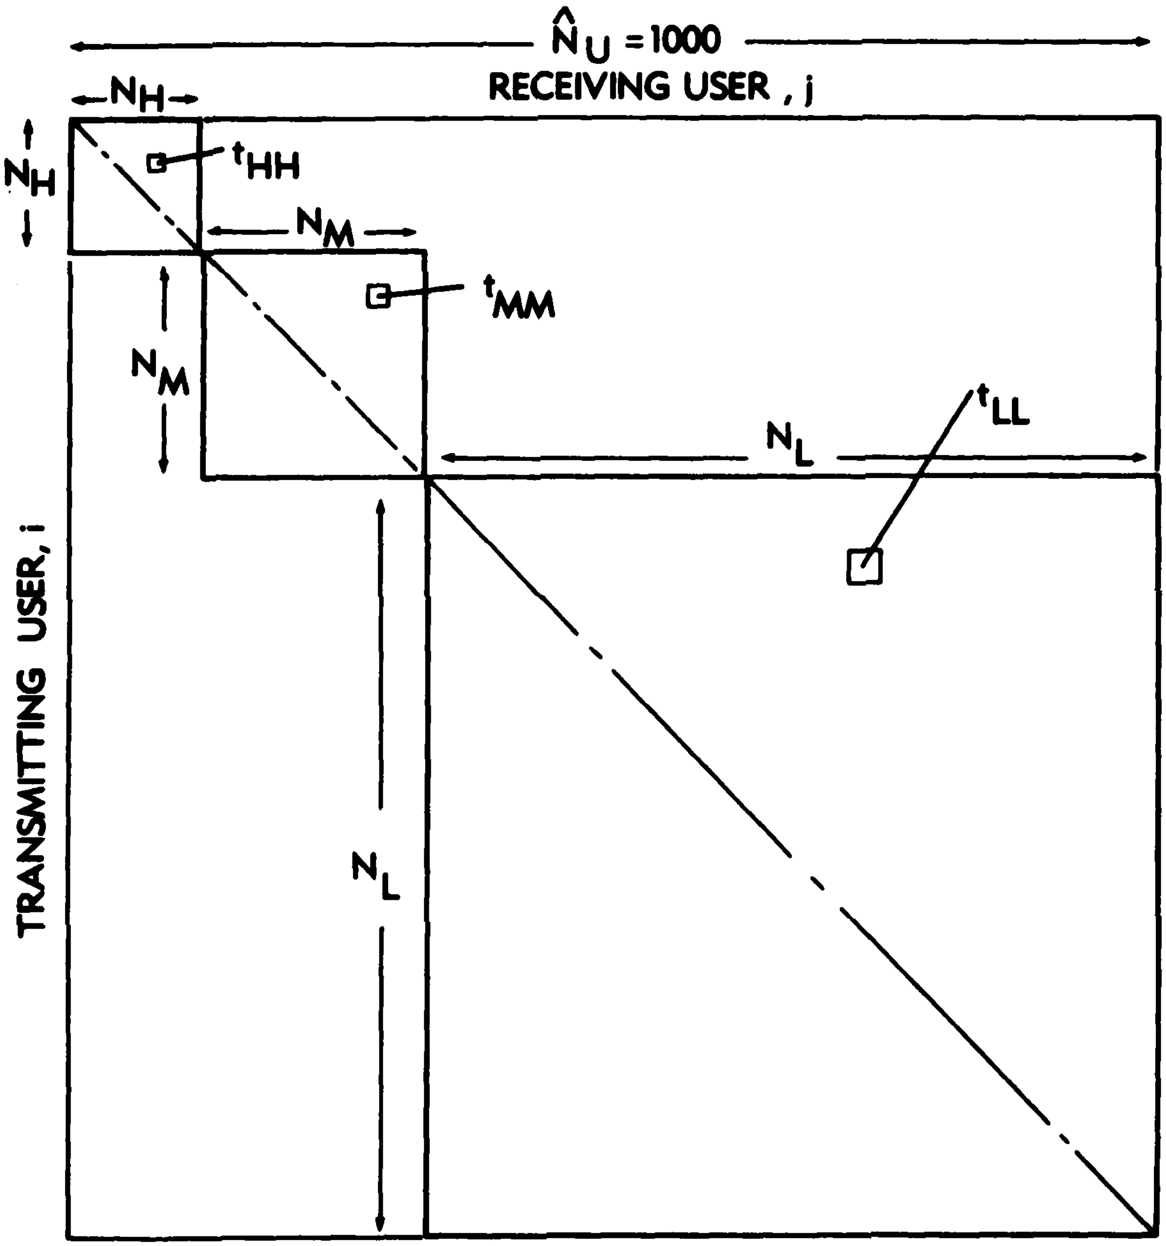

Figure 7-3. General Traffic Matrix

Annual volume of packets - one user of class "x" to another

$$t_{xx} = \frac{P_x \times \hat{P}}{n_x(n_x - 1)} \text{ packets/year} \tag{7-4}$$

Average information bit rate, one user in class "x" to another

$$\bar{R}_{xx} = \frac{P_x \hat{P}}{n_x(n_x - 1)} \text{ packets/year} \times 500 \text{ bits/packet}$$

$$\times \frac{1 \text{ year}}{3.15 \times 10^7 \text{ sec}} \text{ (bits/second)} \tag{7-5}$$

Peak information bit rate, one user in class "x" to another

$$\hat{R}_{xx} = \frac{1}{u_x a_x} \bar{R}_{xx} \text{ (bits/second)} \tag{7-6}$$

As a typical example of the application of these relationships, consider a network with the following characteristics:

a. 20 high-rate users have 60 percent of the traffic, operate continuously throughout the day ($u_H = 1$) with 25 percent activity.

b. 80 users have 30 percent of the traffic, and operate 12 hours/day with 20 percent activity.

c. 900 users have 10 percent of the traffic, and operate an average of 3 hours/day with 5 percent activity.

These assumptions yield peak information bit rates of 0.8 Mbit/s, 0.06 Mbit/s, and 2.5 kbit/s among the high-, medium-, and low-rate users, respectively.

Given that the description above provides sufficient flexibility in defining the network traffic between classes of users, we must next assume a geographic distribution of these users within the satellite coverage area. It is assumed that users are uniformly distributed over CONUS and that the satellite antenna provides S disjoint "spot" coverage areas over this same region as shown in Figure 7-4. The number of spot coverage areas varies inversely as the half-power beamwidth θ_s of the satellite antenna. As the antenna beamwidth is narrowed (providing more gain at the satellite), more spot regions are needed to cover CONUS and fewer users would be included in each region. It is further assumed that each spot coverage region contains N_{ES} earth stations (N_{ES} is at least one) to serve the users in that region. For example, if S = 20, then for the scenario above, each spot coverage region would contain one high-rate user, four medium-rate users, and 45 low-rate users.

With the assumed network configuration, we now want to examine three different techniques for establishing the required connectivity between users. Two cases of "circuit" switching will be considered, one without traffic concentration at the earth stations, and a second where such traffic concentration is possible. The third technique considered consists of idealized packet switching. For all three of these techniques, beam switching is assumed at the satellite where multiple 40 MHz transponders in the satellite are shared among various up-link/down-link beam combinations to provide the required connectivity.

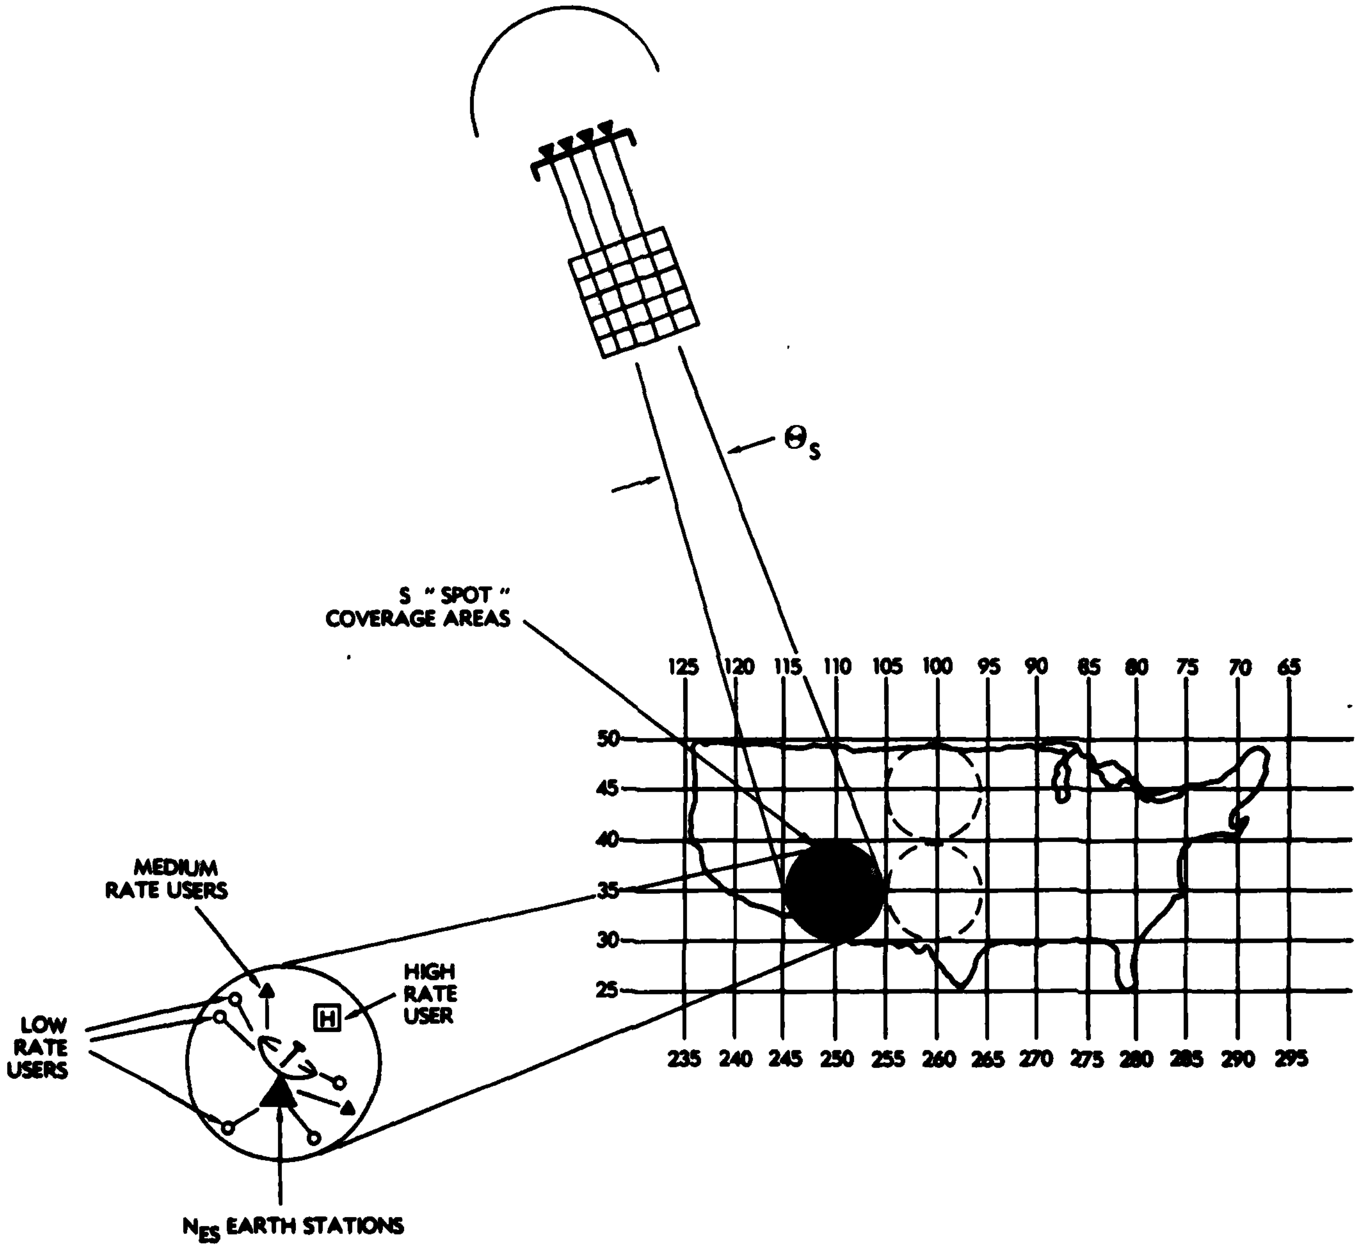

Figure 7-4. Assumption About Distribution of Users

The beam switching can be illustrated by a simple example where two transponders are shared among eight up- and down-link beams. Possible switching patterns are shown in Figure 7-5. The up-link beam is switched slowly among the eight coverage regions and for each connection, the down-link connection of the transponder output is sequenced rapidly through each of the possible down-link coverage areas. The second transponder is likewise cycled through a sequence of slow up-link connections and more rapid down-link connections. The switching patterns are chosen so that touching coverage areas are never illuminated by the antenna at the same instant of time. The total switching frame T_F repeats every second. If the burst transmission rate of each earth station is 60 Mbit/s, then the average transmission rate from spot to spot is simply

$$R_b(\text{average}) \leq 60 \times 10^6 \times \frac{\text{Burst Time}}{\text{Frame Time}} \times \text{Number of Transponders}$$

$$\leq \frac{60 \times 10^6}{64} \times 2 \leq 1.875 \text{ Mbit/s}$$

where the inequality indicates some loss in efficiency due to preambles, guard times, and time needed to synchronize the earth stations to the satellite switch.

The simple example of beam switching in Figure 7-5 points out a restriction on the number of times frequency spectrum can be "reused" among a contiguous set of spot coverage areas. For the example shown, it is possible to find only two <u>complete</u> switching sequences among nontouching spot beams (it is assumed that beams separated from one another by one half-power beamwidth are sufficiently isolated to avoid co-channel interference). A third, incomplete, pattern can be found that allows a third transponder to be used some fraction of the frame period.

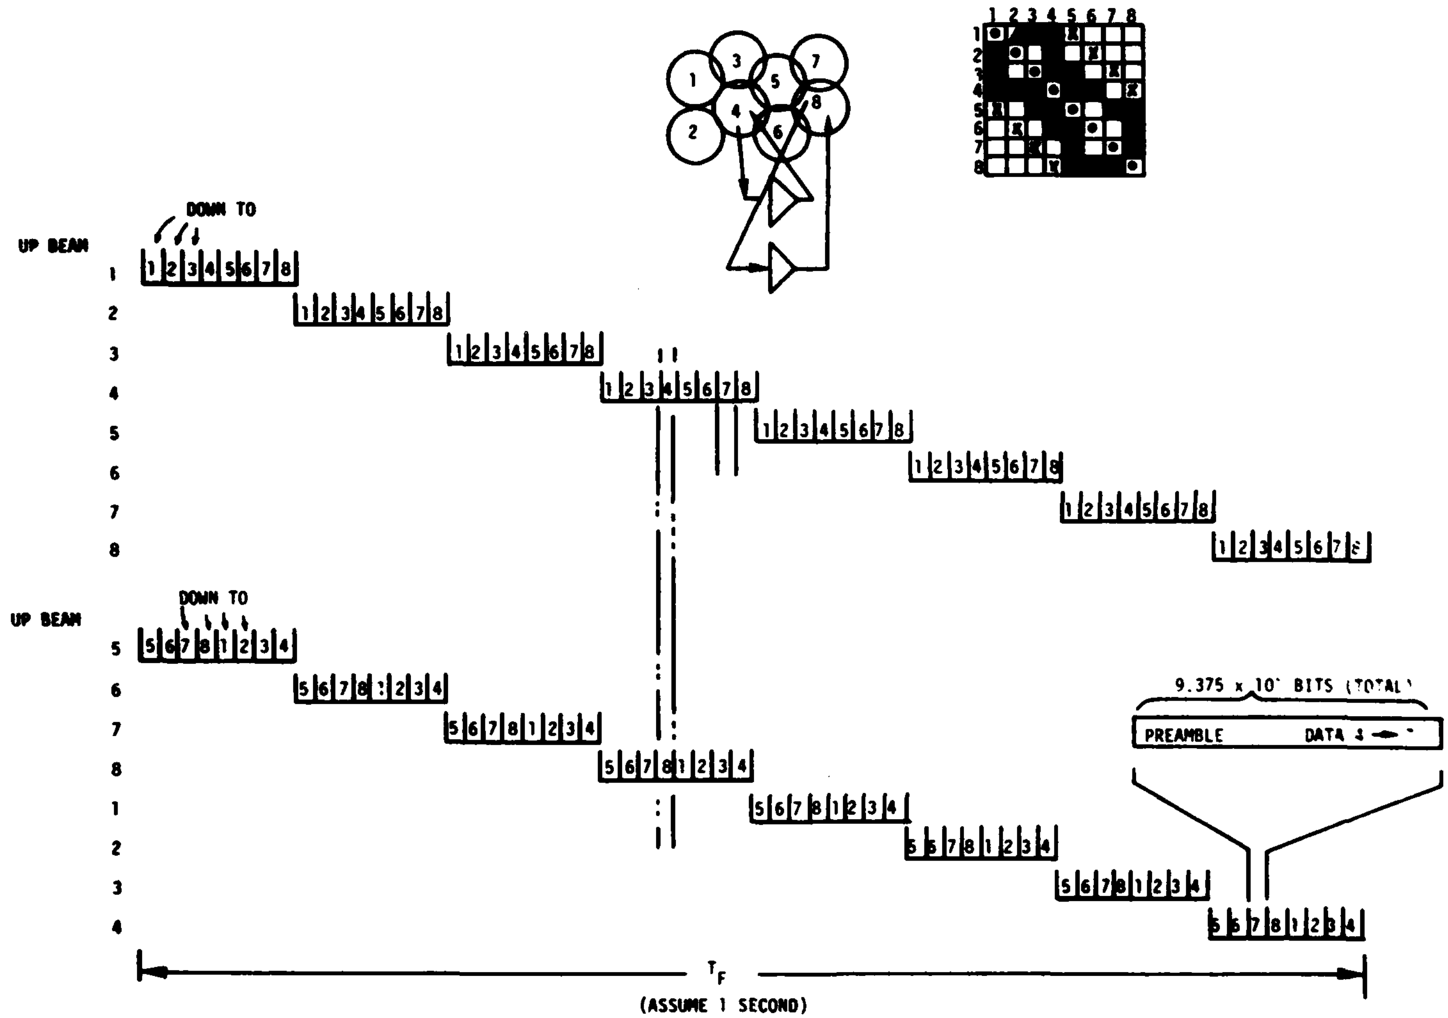

Figure 7-5. Possible Switching Patterns; 2 Transponders Among 8 Beams

Generally, as the number of spot beams, N_s, becomes large, it is possible to find approximately 1/3 N_s switching patterns so that this number of transponders could be utilized to interconnect pairs of beams.

The switching in Figure 7-5 equally distributes transponder capacity among the eight hypothetical coverage areas. The pattern is fairly straightforward because an equal distribution of traffic was assumed. When the traffic between coverage areas is nonuniform, the beam connections must dwell longer on the high-traffic areas. For these cases of unbalanced traffic, it may be difficult to find switching patterns that provide for maximum frequency reuse.

For the perfectly balanced case, we can now consider that portion of the switching frame that interconnects two coverage areas, say region-2 to region-7 in Figure 7-5. We will consider the traffic from the mix of users in region-7. This situation is shown in Figure 7-6. For the representative parameters assumed earlier, each region has one high-rate user, four medium-rate users, and 45 low-rate users. Recall that the characteristics of these users were as follows:

Category	Average Use	Activity	Peak Information Bit Rate
High Rate (1)	24 hrs/day	0.25	800 kbit/s
Medium Rate (4)	12 hrs/day	0.20	60 kbit/s
Low Rate (45)	3 hrs/day	0.05	2.5 kbit/s

The earth station in Region 2 can transmit two bursts in each frame each of duration approximately 1/64 s. Assuming a 60 Mbit/s TDMA system, and neglecting preambles and guard times, the total rate transmitted from the earth staion would be 1.875 Mbit/s. We can think of each burst as providing satellite "channels" of rate equal to the number of information bits put

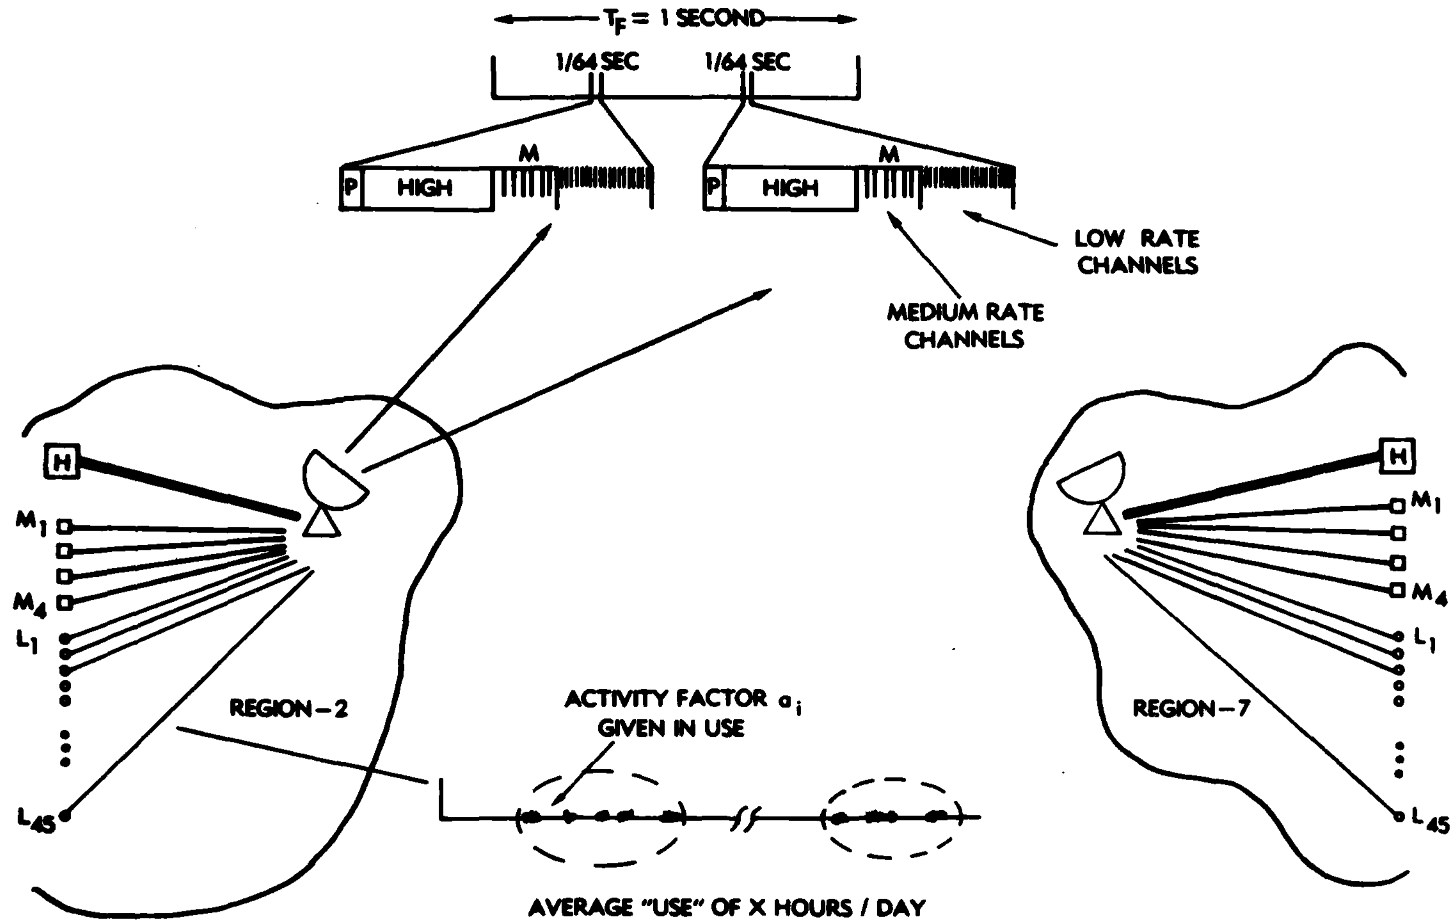

Figure 7-6. Circuits in TDMA Frame

into each burst. For example, a block of 1000 bits in one burst would provide a 1-kbit/s channel.

A simple but inefficient method of providing channels from region 2 to region 7 would be to simply dedicate segments of the burst for all possible circuit combinations. This would result in the following required capacity.

One high rate channel	0.8×10^6 bits
4 x 4 = 16 medium rate channels	0.96×10^6 bits
45 x 45 = 2025 low rate channels	5.063×10^6 bits
(Fully dedicated)....	6.823×10^6 bits

which exceeds the available capacity by almost four-to-one. A second, more practical approach would be to assign a "pool" of circuits and assign these as needed to the user pairs for the duration of their "call" or data transmission session. For example, if we establish one (dedicated) high-rate circuit, 12 medium-rate circuits, and 300 low-rate circuits, the following total capacity is required.

One high rate channel	0.8×10^6 bits
12 medium rate channels	0.72×10^6 bits
300 low rate channels	0.75×10^6 bits
(Pool of switched CKTS)....	2.27×10^6 bits

which exceeds the available capacity by only 20 percent. Of course, there is some possibility that a user in region 2 with traffic for a user in region 7 will find all circuits "busy", in which case he would be unable to send his traffic. Such a "blocked" situation implies that all time slots in the bursts leaving the earth station in region 2 have been assigned. Although the circuits are assigned, this does not necessarily mean that each burst is carrying all information bits. In fact, with low activity data sources, most of the bursts are "empty' or are carrying "filler" data.*

To this point, we have noted that a significant saving in satellite capacity can be realized if circuit connections are provided from a pool of circuits to serve users with relatively low utilization (i.e., those that use the channel only a few hours per day). This saving was realized at the expense of additional equipment at the earth station, and extra time (or transmission capacity) on the channel to establish the temporary connection. Also, such circuit switching is characterized by some probability that all satellite circuits are "busy" or occupied when a particular user wishes to establish a connection. This "blocking probability" must be traded off against efficiency in using the available spacecraft channels.

Given that such a switched network is set up to handle temporary connections between pairs of users, there is one additional step that can be taken to improve efficiency if the channels are not always active after being assigned. Continuing with our representative example, assume that sufficient satellite channels are provided for one high rate, 12 medium rate, and 300 low rate users where the connections by class have activity factors 0.25, 0.2, and 0.05, respectively. The ensemble of input circuits could appear as shown in Figure 7-7. During any one second interval, the high rate channel is "active" with probability 0.25, each of the 12 medium rate channels is active with probability 0.2, and each of the 300 low rate channels is active with probability 0.05. If this activity and lack of activity

*In practice, the lack of activity on a particular channel would probably manifest itself by "all zeros" in the data stream. The data to the modulator would be scrambled (and removed at the receiver) to avoid spans where the carrier is unmodulated.

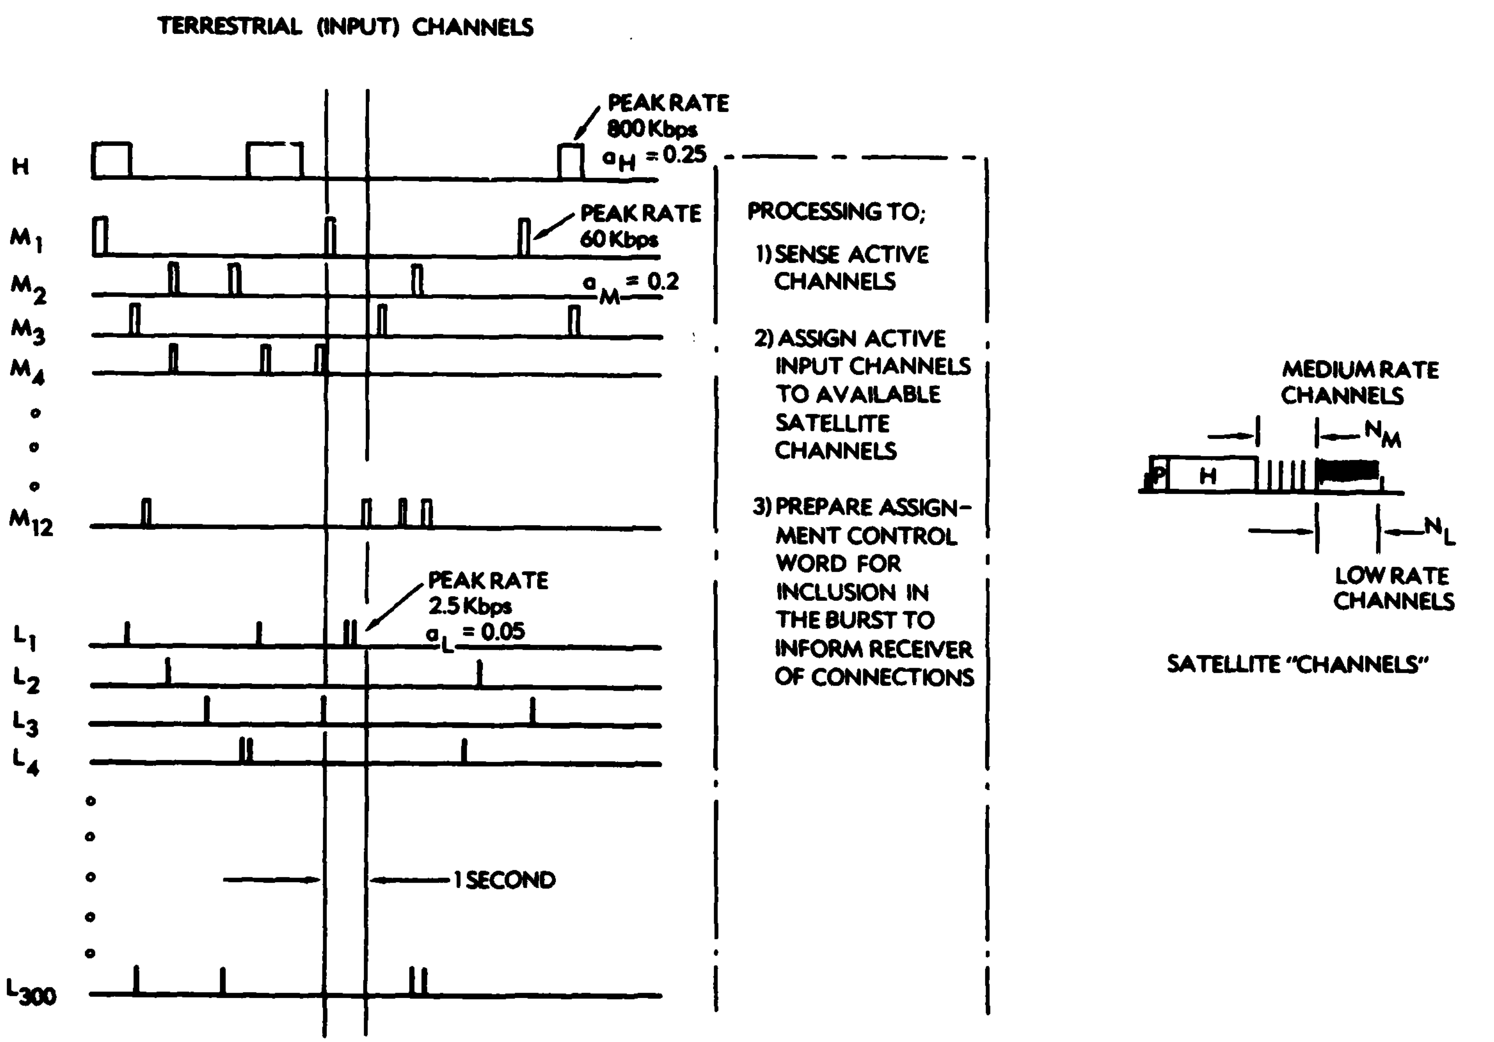

Figure 7-7. Ensemble of Circuit Switched Sources

could be sensed, the satellite channels could be rapidly reconfigured on a frame-by-frame basis to carry only those channels that are currently active. Of course, additional overhead must be put into the transmitted burst to inform the receiver of the changes in the connections.

This technique is commonly used for speech transmission with various names such as TASI (time assignment speech interpolation) and DSI (digital speech interpolation). The technique used to convey the connection information to the receiver generally depends on the relationship between the frame time (in our example assumed to be one second) and the average "speech-spurt" duration. Two extremes are possible:

a. If the frame duration is very short compared to the average duration of an active source interval, then a sequence of "connect/disconnect" messages may be sufficient to keep up with the changes in the various sources from "off" to "on" or from "on" to "off." For example with speech, where the mean spurt duration is about 1 s, the use of a frame duration of 750 μs results in an active channel remaining "on" for an average of more than a thousand frames. A single connect/disconnect message per burst is adequate in this situation to track, or keep up with, the changes in activity of an ensemble of "busy" input circuits.

b. At the opposite extreme, if the frame time is long compared to the average "on" interval, many changes in the active/inactive status of channels could occur in one frame interval. In this case, a longer control word (more burst overhead) may be needed to, in effect, completely define the connectivity between satellite and terrestrial channels for each burst.

Note that this form of dynamic time assignment monitors the activity on a group of circuits "in use" (but each with <100 percent activity) and assigns satellite channels only to those that are active at that instant of time. With such a scheme, a number of satellite channels N_s could serve a number N_T of terrestrial circuits where $N_s \geq aN_T$, where a is the average activity of the input channels. At times, a particular channel may be "frozen out" in this process (data is lost),* and also misconnections can occur due to bit errors in receiving the channel assignment control words. Nevertheless, such systems can achieve a gain that approaches 1/a when the number of terrestrial channels is large. The equipment needed to perform this processing adds to the cost of the earth station.

For the example discussed previously, it is clear that nothing can be done with the single high rate channel to recover the four-to-one loss due to its activity factor. With the 12 medium rate channels, it is likely that some of the five-to-one loss could be recovered and, finally, a large part of the activity loss could probably be recovered for the low rate channels. Assuming that 6 medium rate and 30 low rate satellite channels would provide sufficiently low data loss probabilities, our earlier example now becomes

*With speech, these freeze-out probabilities are typically made 0.01 to 0.05 since the loss of a frame's worth of speech (for short frames) may not be serious. For data transmission, much lower probabilities of lost data may be necessary.

One high rate channel	0.8×10^6 bits
6 (interpolated) medium rate channels	0.36×10^6 bits
30 (interpolated) low rate channels	0.075×10^6 bits
(Pool of switched CKTS with interpolation of low activity CKTS)	1.235×10^6 bits

which is only two-thirds of the available satellite capacity.

Thus, by adding more complexity at the earth stations, the required satellite capacity is reduced. Whether this is a favorable tradeoff, or not, depends on the relative costs involved.

7.5.2 Packet Transmission Assumptions

The discussion to this point has been oriented toward a circuit-switched network where temporary connections are established between pairs of users. The assumption was made that the mix of users was uniformly distributed over CONUS and that these users were grouped into small regions which were served by a single earth station. The satellite provided earth station-to-earth station connections (region-to-region connections when there is only one earth station per region), via a regular beam switching pattern associated with a satellite-switched TDMA multiple access system.

To consider a packet transmission system, it is useful to consider first a more general model, as shown in Figure 7-8, consisting simply of a single very wide bandwidth spacecraft which serves the total population of users. In this case, these users are confined to CONUS. We can first consider the total aggregate of users who generate an average of 8×10^{12} packets/year. As before, we will assume that a packet consists of 500 information bits. The composite input rate from all users is

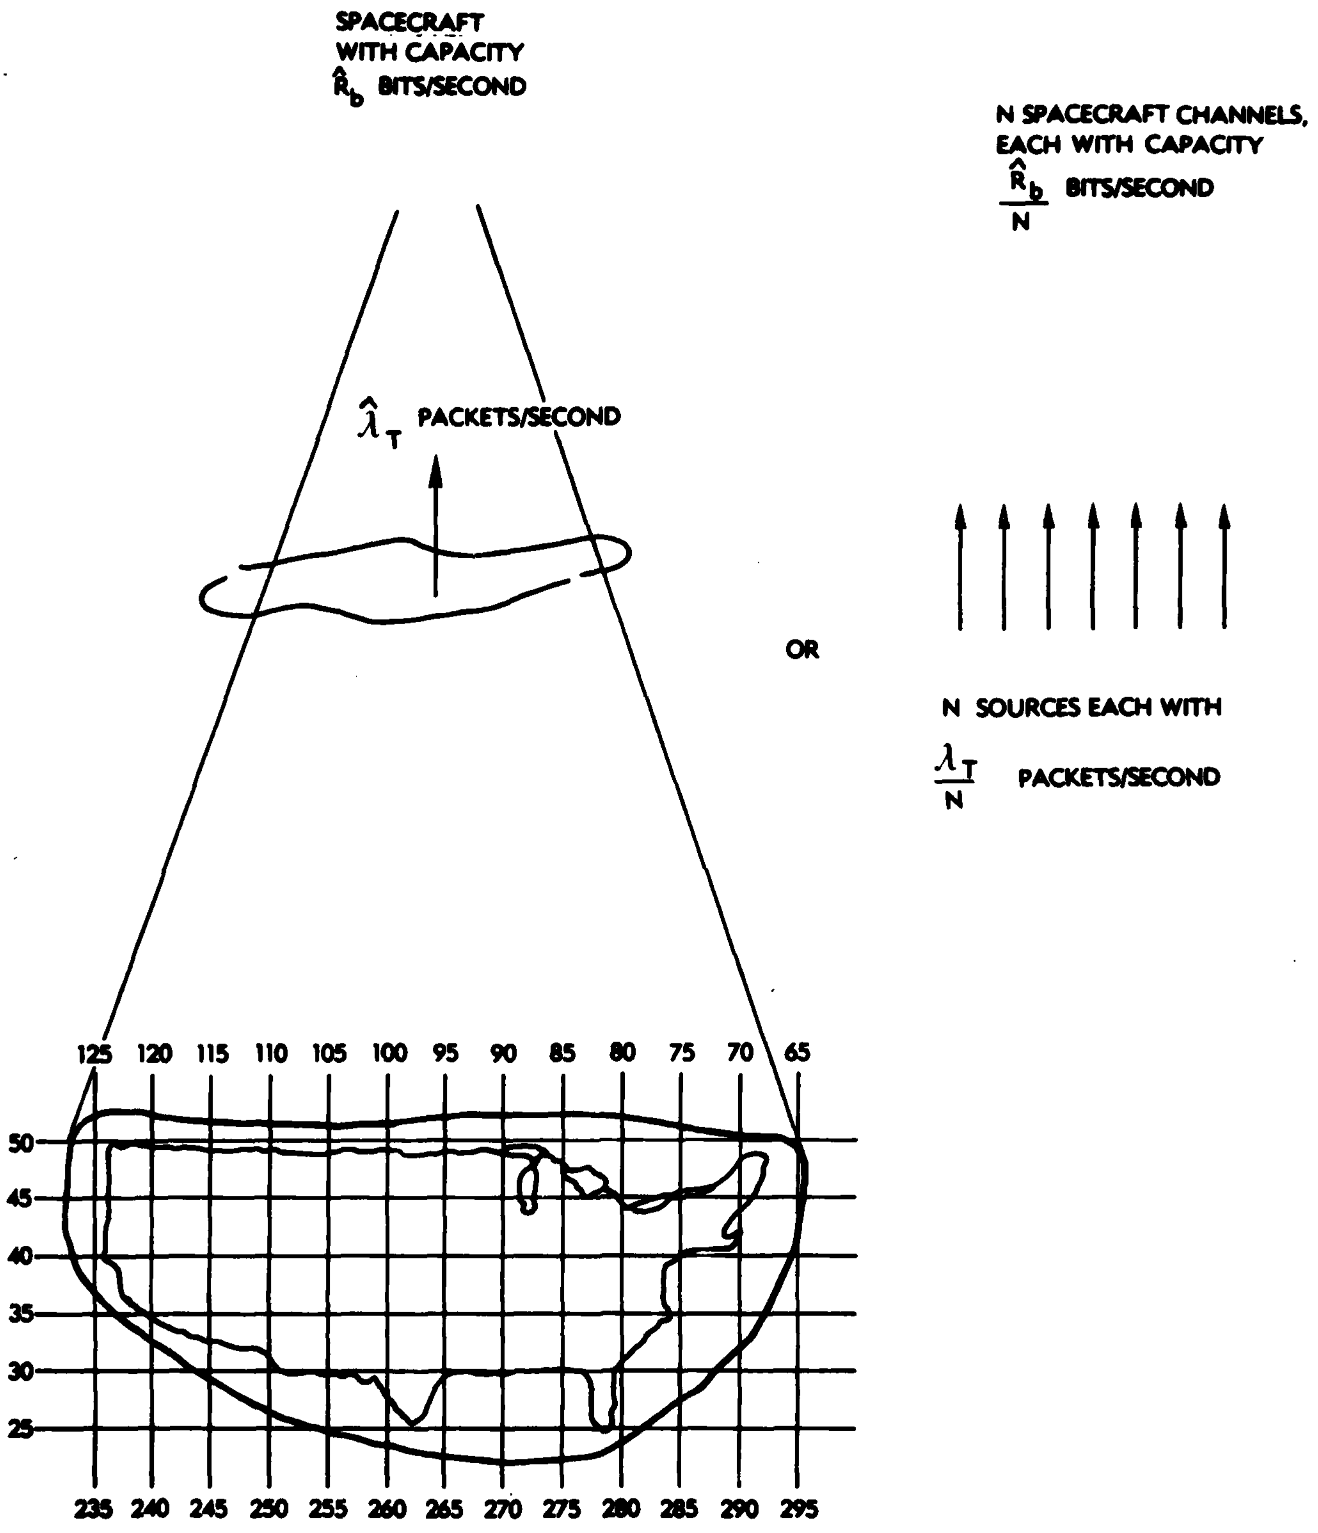

Figure 7-8. Packet Transmission via Satellite

2.537×10^5 by a single packet transiting the path from transmitting earth station to receiving earth station can be written as;

$$\bar{D}_p = \bar{d}_s + \bar{d}_T + \bar{d}_Q \tag{7-7}$$

where

$\bar{d}_s$ = average propagation time over the satellite path (0.27 s)

$\bar{d}_T$ = average "transit" time of the packet through the earth station (packet length in bits/divided by channel transmission rate in bit/s)

$\bar{d}_Q$ = queueing delay due to the assumption that the packets arrive at random times and must wait to be transmitted from an earth station (either due to queues in the earth station, or contention for the spacecraft capacity, or both).

For a general model of Poisson arrivals, the queueing delay can be written as

$$\bar{d}_Q = \frac{500\ f}{\hat{R}_b} \cdot \frac{1}{2} \frac{\rho}{1 - \rho} \tag{7-8}$$

where

$\hat{R}_b$ = channel rate provided by the satellite circuit (bits/s)

f = factor to account for additional overhead that is added to the 500 information bits that are presented to the earth station before the "packet" is sent over the channel

ρ = ratio of packets/s in $\hat{\lambda}_T$ to the capacity of the channel for transmitting packets [$\hat{\lambda}_o = \hat{R}_b/(500\ f)$]

The "transit" delay $\bar{d}_T$ is simply 500 f/$\hat{R}_b$ s. Combining the three delays gives;

$$\bar{D}_p = 0.27\ s + \frac{500\ f}{\hat{R}_b} + \frac{500\ f}{\hat{R}_b} \cdot \frac{1}{2}\left[\frac{500\ f\hat{\lambda}_T}{\hat{R}_b - 500\ f\hat{\lambda}_T}\right]$$

$$= 0.27\ s + \frac{500\ f}{\hat{R}_b}\left\{1 + \frac{1}{2}\left(\frac{500\ f\hat{\lambda}_T}{\hat{R}_b - 500\ f\hat{\lambda}_T}\right)\right\} \qquad (7\text{-}9)$$

Clearly, $\hat{R}_b$ must be greater than 500 $f\hat{\lambda}_T$ (>126 Mbit/s if f = 1) or the queueing delay becomes extremely large. For a channel transmission rate of 126 Mbit/s, the packet transmission interval is 4 · f μs which is totally negligible compared to the 0.27 s propagation time. Thus, it appears that queueing delays can be made completely negligible by oversizing the channel slightly, say by making $\hat{R}_b = 1.1 \times f \times 500\ \hat{\lambda}_T$. In this case, the channel capacity required is directly proportional to the factor f which denotes the multiplicative factor due to the overhead which must be put into each packet.

As an extension of the wide band packet transmission system, we can consider a uniform subdivision of the input traffic $\hat{\lambda}_T$ into N sources. These sources could be considered as earth stations, each serving a complement of users, or as the users themselves. The total system capacity $\hat{R}_b$ is likewise divided by N indicating that each separate source shares the channel capacity. Such sharing would apply to the earlier example where earth station-to-earth station traffic was assigned a certain segment of a TDMA frame. For a general subdivision by N, the average delay experienced by a packet would be

$$\bar{D}_p = 0.27\ s + \frac{500\ fN}{\hat{R}_b}\left\{1 + \frac{1}{2}\left(\frac{500\ f\hat{\lambda}_T}{\hat{R}_b - 500\ f\hat{\lambda}_T}\right)\right\}s \qquad (7\text{-}10)$$

In a network with 20 earth stations, for example, there are 20 x 19 = 380 earth station-to-earth station paths so that the delay increases more rapidly as $\hat{R}_b$ is reduced toward 500 $f\hat{\lambda}_T$ than for the case where the channel is not subdivided. Subdivision of the channel is avoided, of course, for the situation of a single CONUS beam with a single wide band transponder. Here, all users have access to all other user's transmissions. This is not true if spacecraft capacity is divided into segments, either through TDMA, FDMA, or spatially through satellite switching of beams. The latter type of multiple access allows reuse of the frequency spectrum.

In the remainder of this section, we will consider the cost tradeoffs involved with these three systems; namely

a. packet switching,
b. circuit switching, and
c. circuit switching with processing at the earth station to "concentrate" low activity sources.

One measure of system effectiveness will be the cost per packet.

7.5.3 Cost Analysis

Our overall goal in the following will be to determine the average cost per "packet" for various mixes of user types, user utiliztion factors, and user activity factors. Other variables will include the frequency band of operation, earth station characteristics, and the spacecraft parameters. Stated generally, the overall goal is to compute a per-packet cost of the form;

$$\left\{\begin{matrix}\text{Packet}\\\text{Cost}\end{matrix}\right\} = \frac{\{\text{Annual earth station cost}\}}{\{\text{No. of packets transmitted per year from earth station}\}}$$

$$+ \frac{\{\text{Total annual cost of the space segment to support the network}\}}{\{\text{Total packets transmitted by all earth stations in one year}\}} \tag{7-11}$$

If we denote the earth station cost by $\$_{ES}$, and assume N_{ES} earth stations in the network, then this general relationship reduces to;

$$\{\text{Packet Cost}\} = \frac{N_{ES}A_{ES}\$_{ES} + N_{SC}\$_{SC}}{\hat{P}} \tag{7-12}$$

where

N_{ES} = number of earth stations in the network

$\$_{ES}$ = cost of an earth station

A_{ES} = annualization factor (e.g., 0.3) to convert earth station cost to annual cost

$\$_{SC}$ = annual cost of a "unit" of spacecraft capacity. The unit will be assumed as 40 MHz

N_{SC} = number of units of spacecraft capacity required to support the network

$\hat{P}$ = total network traffic in packets/year (for the nonimal scenario, $\hat{P} = 8 \times 10^{12}$)

A tradeoff is inherent to this relationship in that two extreme situations can be postulated. First, a network can be envisioned with small inexpensive earth stations that use the spacecraft capacity rather inefficiently. In this situation, earth segment cost would be minimized but the annual cost for the space segment would be large. In contrast to this, large expensive earth stations could be used in the network, and more complex multiple-access equipment could be utilized in order to use the spacecraft capacity as efficiently as possible. Here, the division of costs would be shifted toward increased earth segment costs and decreased cost for satellite resources. A "least-cost" solution may exist as the relative costs are varied between these extremes.

The tradeoff analysis deals with equation (7-12). The total annual volume of network traffic, $\hat{P}$ (packets/year), is assumed to be fixed. To estimate space segment cost, we will simply assume a range of 0.5, 1.0, 2.0, and 4.0 \$M per year as the "cost" of a 40 MHz segment of bandwidth. This segment of bandwidth corresponds to a typical transponder so that, in effect, we assume that transponders can be rented on an annual basis, with an eight-to-one spread in the rental cost.

The number of earth stations in the network, N_{ES}, will be retained as a variable. The annualization factor, A_{ES}, is simply a constant. Thus, the major part of the analysis deals with the calculation of $\$_{ES}$ and N_{SC}. Both of these variables are dependent upon many factors. The relationships used to obtain these costs are given in the following.

7.5.4 Earth Station Cost

Earth station cost is determined as two parts; one independent of the transmission technique (i.e., circuit switching, circuit switching with activity concentration, or packet switching), and a second part that estimates the cost of the equipment needed to implement each technique. Total earth station cost is therefore dependent upon the transmission technique, i, as

$$\$_{ES}(i) = \$_{ES1} + \$_{ES2}(i) \tag{7-13}$$

where i = 1, 2, 3, corresponds to packet switching, circuit switching, and circuit switching with activity concentration, respectively.

The first element of earth station cost, $\$_{ES1}$, is dependent on such variables as antenna diameter, HPA size, frequency band of operation, LNA temperature, and burst transmission rate. Given these input variables, the first element of earth station cost is obtained from the subroutine*;

CES3(KBND, D, PES, RB, TES, A1)

where

KBND = frequency band coded as 1 = 6/4 GHz; 2 = 8/7 GHz; 3 = 14/11 GHz; 4 = 30/20 GHz

D = antenna diameter in meters

PES = power density (watts/Mbit/s) radiated from the earth station

*See Appendix B.

RB = burst bit rate from/to the earth station (Mbit/s)

TES = LNA noise temperature

A1 = cost of ANTENNA, HPA, LNA and modem in dollars

The second cost element in equation (7-13) includes the cost of the equipment needed to convert the incoming data lines at the earth station to the data that is actually transmitted over the satellite channel. These costs are calculated by the subroutine ESCOST( )* which returns the total cost $\$_{ES}(i)$ in equation (7-13). The assumed cost of the additional earth station equipment has a significant impact on the results of the tradeoff analysis. The model used to estimate this equipment cost is shown in Figure 7-9. Note that the assumed cost relationships are composed of fixed costs plus costs that are functions of the number of high, medium, and low-rate users that are connected to the earth station. For the *circuit-switched case*, cost of the switching equipment is assumed to be a function of t_k, the total potential circuits that can be established at each earth station. For example, with 80 medium-rate users in a network with 20 earth stations, each earth station serves 4 medium-rate users. In this case, assuming full connectivity, there are

$$t_2 = 4 \cdot 80 - 4 = 316$$

potential circuits that could be required at the one earth station (each user may require a connection to every other user, excluding himself). With less than 100 percent utilization, it is unlikely that all of these circuits would be busy (or "off hook")

*See Appendix B.

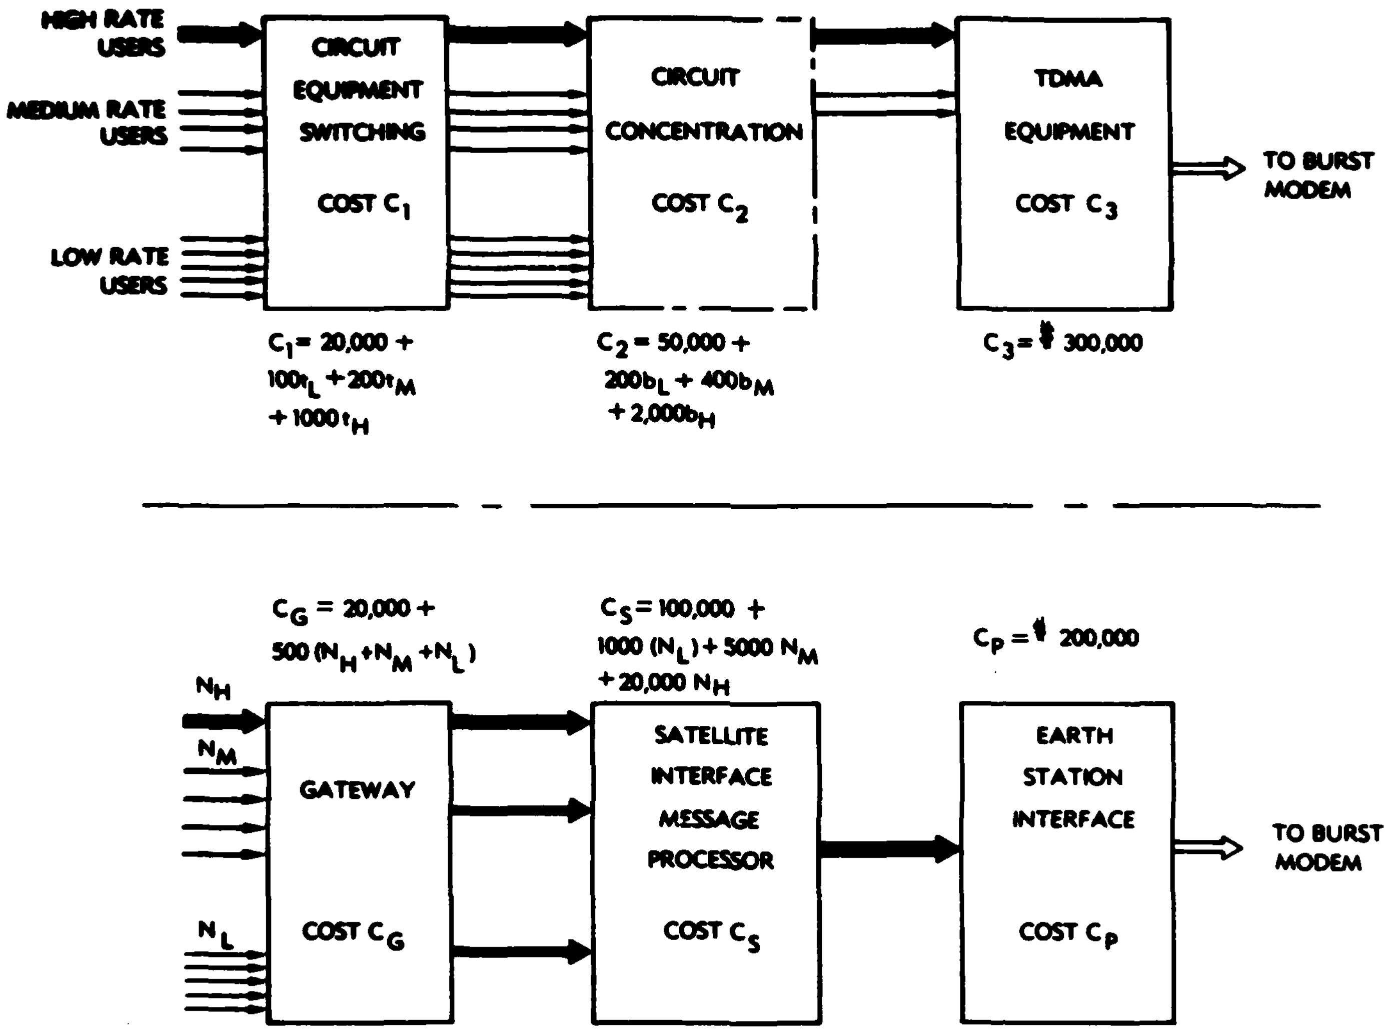

Figure 7-9. Elements in Earth Station Interface Equipment Model

at any given time. However, the circuit switching equipment must provide for the possible connections. When circuit concentration is possible, the concentrator deals only with the "off-hook" circuits that have been established. Note, therefore, that concentrator costs are a function of r_k, the number of "busy" circuits.

The cost of packet switching equipment is estimated using the model at the bottom of Figure 7-9. These costs are assumed to be a function of the number of high-, medium-, and low-rate users, n_H, n_m, n_L, respectively, that interface the earth station.

7.5.5 Required Spacecraft Capacity

The required spacecraft capacity is determined as the number of 40 MHz units of bandwidth that are required to carry the total network traffic for a given transmission technique. This procedure has three steps.

a. Overall carrier to noise ratio, $(C/N)_T$, is determined as a function of the earth station and satellite parameters.

b. Given $(C/N)_T$, the best combination of modulation technique and error control coding is determined that maximizes the bandwidth efficiency obtained, η, (in bit/s/Hz).

c. Given η, it is possible to obtain the number of transponders needed to carry the total network traffic. This calculation accounts for the preambles, overhead bits, and control bits that are necessary to implement the transmission technique.

A homogeneous network is assumed with N_{ES} identical earth stations. Adjacent satellites in the orbit are spaced

ϕ_{SAT} degrees and adjacent satellite interference is estimated by assuming that an identical network uses these adjacent satellites. Up-link carrier to noise ratio is obtained as

$$(C/N)'_{UP} = \frac{4 \cdot 14 \times 10^4 \ P_{ES} L_U D^2}{\theta^2_{SAT} \times T_{SAT}} \qquad (7\text{-}14)$$

where

P_{ES} = peak power density radiated from the earth station in watts/MHz

L_U = up-link loss

D = earth station antenna diameter (m)

θ_{SAT} = spot size of the antenna coverage from the satellite

T_{SAT} = equivalent noise temperature of the satellite receiver

The effects of co-channel and adjacent channel interference are included to obtain overall up-link C/N as

$$(C/N)_{UP} = \left(1/(C/N)'_{UP} + \Gamma^{UP}_{CCI} + 2\gamma^{UP}_{ASI}\right)^{-1} \qquad (7\text{-}15)$$

where

$(C/N)'_{UP}$ = carrier-to-noise ratio computed in equation (7-14)

Γ^{UP}_{CCI} = total up-link co-channel interference experienced. This level will be zero with no frequency reuse

γ^{UP}_{ASI} = up-link interference experienced from signals using the adjacent satellite

Down-link C/N is obtained from

$$(C/N)'_{DN} = \frac{4 \cdot 14 \times 10^4 \; P_S L_D D^2}{\theta^2_{SAT} T_{ES}} \tag{7-16}$$

where

P_S = spacecraft radiated power density in watts/MHz

L_D = down-link loss

T_{ES} = effective systems noise temperature of the receiving earth station

The effects of co-channel and adjacent satellite interferences are added as on the up-link using

$$(C/N)_{DN} = \left(1/(C/N)'_{DN} + \Gamma^{DN}_{CCI} + 2\gamma^{DN}_{ASI}\right)^{-1} \tag{7-17}$$

and the overall C/N is finally obtained as

$$(C/N)_T = \left((C/N)^{-1}_{UP} + (C/N)^{-1}_{DN}\right)^{-1} \tag{7-18}$$

Given the overall operating carrier-to-noise ratio, the subroutine (see Appendix B),

BPSHZ(CNT, IPE, R, MM, PZ)

is used to obtain the achievable bandwidth efficiency on the link. The routine assumes PSK modulation and uses CNT $\doteq$ $(C/N)_T$ and required error rate (10^{-IPE}) as inputs, and determines code rate (R), type of PSK (2^{MM} phase, MM = 1, 2, or 3), and the bandwidth efficiency (η) in bits/s per Hz.

The next step is to use the bandwidth efficiency in a calculation of the number of 40-MHz bandwidth increments (i.e., the number of transponders) that are needed to carry the traffic in the satellite. To perform these calculations, it is useful to define three types of "circuits".

a. Terrestrial circuits - potential user-to-user connections as established at each earth station. Since we are assuming that full connectivity is required, i.e., any user of a class can potentially send traffic to any other user of the same class (except himself), there are a large number of these "potential" terrestrial circuits. Of course, with utilization <1.0, these circuits are not all active at any one time.

b. "Busy" circuits - these are the circuits that are being utilized or are "off-hook" at any given time. If the activity factor is less than 1.0, these circuits, although busy, carry actual traffic less than full time.

c. Satellite circuits - these are the slots provided in the TDMA frame to carry blocks of information bits provided by the high-, medium-, and low-rate sources. For example, if the low-rate user has a peak information bit rate of 2.5 kbit/s, and the TDMA frame, T_F, = 1 s, the satellite circuit would consist of a 2,500-bit block in the TDMA burst.

For a network with N_U users and N_{ES} earth stations, the total number of users of type k is simply,

$$n_K = f_K N_U \qquad (7\text{-}19)$$

(k = 1, 2, 3 for high, medium and low rate, respectively). The number of users per type by earth station is simply

$$e_K = n_K/N_{ES} \qquad (7\text{-}20)$$

since we are assuming that users are equally divided among the available earth stations. We now assume that the network is covered by N_B "beams" formed at the satellite. If $N_B = 1$, we have a CONUS-coverage beam, and all earth stations send single long multidestinational TDMA bursts with all of their traffic. With multiple beams, it is assumed that satellite switching takes place, similar to the example in Figure 7-5, so that earth stations must split their traffic according to the beam-to-beam traffic matrix. In this case, the beam-to-beam matrix is uniform. We will assume further that the beam-to-beam traffic matrix is the same as the earth station-to-earth station traffic matrix. This latter assumption simply means that the spot beams are narrow enough, and the earth stations sufficiently dispersed, that the satellite antenna can be pointed at the earth stations individually.

The potential terrestrial circuits can be obtained as

$$t_K = t_{K1} + (N_B - 1)\ t_{K2} = e_K\left\lceil\frac{n_K}{N_B}\right\rceil - 1 + (N_B - 1)\ e_K\ \frac{n_K}{N_B} \qquad (7\text{-}21)$$

with a single beam, these relationships reduce to simpler forms. In the multibeam case, it is necessary to separate the individual terms t_{K1} and t_{K2} since these circuits represent the individual beam-to-beam (or earth station-to-earth station) paths. With circuit switching, the groups of circuits t_{K1} and t_{K2} may be reduced if the utilization is low. This reduction is calculated by the subroutine (see Appendix B).

BLOCK(t, UTIL(k), NPBK, r)

where

t = number of input circuits
UTIL(k) = fractional utilization of each circuit
NPBK = allowable blocking probability
r = required number of circuits ($r \leq t$)

For the circuit switched case, therefore, the number of satellite channels required is obtained as

$$r_K = e_K\left[\frac{n_K}{N_B} - 1\right] + (N_B - 1)\ e_K \frac{n_K}{N_B} = r_{K1} + (N_B - 1)\ r_{K2} \qquad (7\text{-}22)$$

The total number of transponders required to serve the network is obtained by first obtaining the required information bit rate

$$R_{ES} = \text{bit rate/earth station} = \sum_{K=1}^{3} r_K \hat{R}_{KK} \qquad (7\text{-}23)$$

where

r_K = number of "busy" circuits by class
$\hat{R}_{KK}$ = peak bit rate from equation (7-6)

The bit rate per earth station is converted to the total number of transponders needed as follows

$$\hat{C}_{\substack{CKT \\ SW}} = N_{ES} \left\{ \frac{R_{ES}}{\eta B_T} \left[1 + \frac{N_B \rho}{R_{ES} T_F} \right] \right\} \text{ transponders} \qquad (7\text{-}24)$$

where

N_{ES} = number of earth stations in the network

R_{ES} = information bit rate per earth station

η = bandwidth efficiency in bit/s/Hz

B_T = transponder bandwidth

N_B = number of beams (also the number of bursts transmitted from the earth station)

ρ = number of bits in the TDMA preamble

T_F = TDMA frame duration

For circuit switching with activity concentration, a similar procedure is followed except that the statistical reduction of potential terrestrial circuits is performed twice, once accounting for the degree of utilization and a second time to account for activity. The routine "BLOCK" is thus used twice

BLOCK(t, UTIL(k), NPBK, r)

BLOCK(r, ACT(k), NFREZ, c)

where

t = potential terrestrial input circuits

UTIL(k), ACT(k) = fractional utilization and activity, respectively, of class k

NPBK, NFREZ = blocking and freeze-out probabilities, respectively

r = "busy" circuits $(r \leq t)$

c = concentrated busy circuits $(c \leq r)$

Additional overhead bits are put into each burst to account for the control needed to accomplish the concentration. Total satellite capacity is thus obtained as

$$\hat{C}_{\substack{CKT\\SW+\\CONS}} = N_{ES} \left\{ \frac{R'_{ES}}{\eta B_T} \left[1 + \frac{N_B \rho + C_{OH}}{R'_{ES} T_F} \right] \right\} \qquad (7\text{-}25)$$

which is similar to equation (7-24) except that

$$R'_{ES} = \sum_{K=1}^{3} C_K \hat{R}_{KK}$$

$$C_{OH} = \text{concentration overhead} = \sum_{K=1}^{3} \left[r_{K1} + (N_B - 1)\, r_{K2} \right]$$

For packet switching, we take a somewhat simpler view of the entire question of satellite capacity. We begin with the <u>average</u> network transmission rate

$$R_{AVG} = \frac{500 \ \frac{\text{information bits}}{\text{packet}} \times \hat{P} \ \frac{\text{packets}}{\text{year}}}{365 \times 24 \times 3600 \text{ s/yr}} \text{ information bits/s}$$

If we simply divided R_{AVG} by the product $\eta \times B_T$ (the capacity of one transponder), we would have an absolute minimum number of transponders needed to serve the network. Thus;

$$\hat{C}_{PCKT} > \frac{R_{AVG}}{\eta B_T}$$

We should now increase this absolute minimum capacity to account for the following known sources of inefficiency;

a. each "packet" of 500 information bits must be preceded by an acquisition preamble of length ρ bits (the same as for the TDMA bursts),

b. each packet contains "h" header bits which (among other things) signify the addressee of the packet,

c. the transmission channel must be oversized slightly [see equation (7-9)] to avoid large delays, and

d. a certain fraction of the capacity in the packet transmission system is consumed by control information. We will assume that this factor differs for the CONUS-coverage and multibeam cases.

The transponder capacity for packet switching is thus;

$$\hat{C}_{PCKT} = R_{AVG}\left[1 + \frac{\rho + h}{500}\right]\left[1 + f_D\right]\left[1 + f_c(\ell)\right] \qquad (7\text{-}26)$$

where

ρ = bits in the preamble of a packet

h = header bits in a packet

f_D = a factor <1 to account for the oversizing of the transmission channel to avoid delay

$f_c(\ell)$ = a factor to account for the control information; $\ell = 1$ CONUS coverage (all users receive all other users packets), $\ell = 2$ multibeam coverage. For example, if ten percent of the system capacity must be used for control, $f_c(\ell) = 0.1$.

7.6 TRADEOFF RESULTS

The tradeoff model described in the last section has been exercised to compare the cost effectiveness of the three transmission techniques as certain major network characteristics are varied. In this process, certain parameters are kept fixed. For example, spacecraft characteristics such as noise temperature and transmitted power level (watts/MHz) are kept fixed. We do, however, consider two conditions of antenna coverage; CONUS coverage with a single shaped beam, and spot coverage. For the latter, we assume that an upper limit on spacecraft antenna size of 4 m is imposed by the payload capability of the Shuttle launch vehicle. As noted earlier, we assume that the total traffic volume ($\hat{P}$ packets/year) is fixed, and also that the number of users in the network is fixed at 1000. The total number of users is divided into three classes, high-, medium-, and low-rate users. Although the high-rate users are few in number (generally one per earth station), these users carry a large fraction of the network traffic. At the opposite extreme, most of the network users are of the low-rate type, and each of these generates a very low proportion of the total network traffic.

The utilization and activity factors are likely to be significant variables in the analysis so particular attention is given to varying these parameters over typical ranges that might be encountered in practice. The tradeoff results presented in the following apply to the different cases shown in Figure 7-10. A basic variable in the analysis is the number of earth stations in the network, N_{ES}; values of 20, 40 and 100 have been used. Two cases of spacecraft antenna coverage have been considered, CONUS coverage, and "spot" coverage. Two frequency bands have been considered and typical rain losses have been included at 14/11 GHz to give better than 99 percent availability. A range of earth station antenna sizes have been considered for each frequency band and satellite antenna coverage type. Rather than attempt to determine the cost of the various spacecraft, a range of annual lease costs have been assumed for a 40-MHz increment of satellite bandwidth. These costs are assumed to range from $0.5M to $4.0M. Three different combinations of user "utilization" and "activity" are assumed. The first serves as a baseline where both utilization and activity = 1.0 for all user types. This somewhat unrealistic case corresponds to the situation where all user types send data continuously, 24 hours per day. A second case corresponds to moderate utilization and low activity. These assumptions about utilization and activity have a fairly significant impact on the results obtained. These assumptions are summarized in Table 7-5. Additional assumptions and fixed parameters for the analysis are listed in Table 7-6.

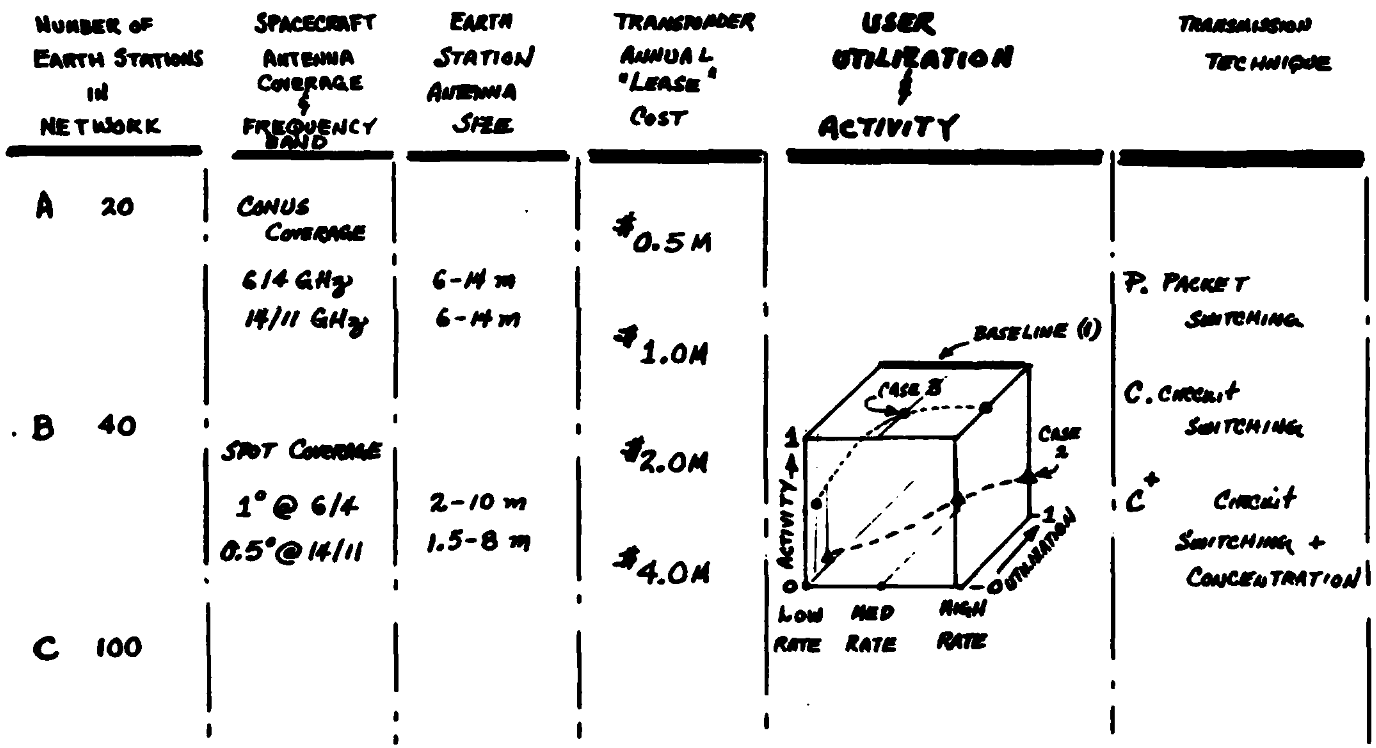

Figure 7-10. Variables in Tradeoff Analysis

Table 7-5. Assumed Values for Utilization and Activity

	Type of User		
	High-Rate	Medium-Rate	Low-Rate
Baseline (Case 1)			
Utilization	1.0	1.0	1.0
Activity	1.0	1.0	1.0
Case 2			
Utilization	1.0	0.5	0.125
Activity	0.25	0.10	0.05
Case 3			
Utilization	0.4	0.3	0.1
Activity	1.0	1.0	0.5

An initial set of tradeoff results is shown in Figure 7-11 where the tradeoff model has been exercised for 20, 40, and 100 earth stations in the network (from top to bottom), for cases 1 (baseline), 2, and 3 (from left to right). The results in this figure apply only to 6/4 GHz operation. Four different subsets of results are shown within each set; the CONUS coverage results are shown as dashed curves, results for one-degree spot beam coverage as solid curves. Results are also included for $0.5M and $4.0M annual cost for a 40-MHz increment of spacecraft bandwidth.

Although the baseline case (at the left of Figure 7-11) of full utilization, 100-percent activity is unrealistic, it does serve as a calibration of the model. The results in all of the curves in Figure 7-11 give cost (dollars per 10^6 "packets") versus earth station antenna diameter (in meters). Transmission

Table 7-6. Fixed Parameter Assumptions

Earth Station			
HPA Power	20 W/MHz		
LNA Temperature	120 K @ 4 GHz	300 K @ 11 GHz	
Burst Rate	60 Mbit/s		
Satellite:			
HPA Power	0.25 W/MHz		
LNA Temperature	2000 K		
Beam Size	7° CONUS	1° Spot @ 6/4 GHz	0.5° Spot @ 14/11 GHz
Link:			
Up-Link Loss	-0.27 dB @ 6 GHz	-4.95 dB @ 14 GHz	
Down-Link Loss	-0.04 dB @ 4 GHz	-3.01 dB @ 11 GHz	
Co-channel Interference	-300 dB CONUS	-24 dB Spot	
Satellite Separation	4° @ 6/4 GHz	3° @ 14/11 GHz	
Packet:			
Information Bits	500		
Preamble Bits	100		
Header Bits	275		
Capacity Overhead	10%		
Control Overhead	10% CONUS	50% Spot	
Circuit:			
Circuit Blocking Prob.	1%		
Concentration Freezeout Prob.	0.01%		

techniques are noted on these figures as "P" (packet), C (circuit switched) and C^+ (circuit switched with activity concentration). For the baseline case with 100-percent activity, activity concentration gives no advantage and both circuit-switched configurations give the same cost/packet. Note for the baseline case that packet transmission is more expensive than circuit switching. The difference is small when spacecraft annual cost is small ($0.5M/year) and when there are only 20 earth stations in the network. The difference in cost increases significantly as annual transponder cost increases to $4.0M. A final general observation from the baseline results in Figure 7-11 is that broad cost minima are observed as earth station size is varied. The least-cost earth station size is reduced as more stations are utilized in the network. With only 20 earth stations, the minima lie in the region 8-12 m; with 100 stations in the network, the least-cost size reduced to the 5-8 m range.

Case 2 in Table 7-5 assumes moderate utilizations and activities, with the activity factors reduced to low values for the low-rate users. Recall that total annual network traffic is a constant (8×10^{12} packets/year) so that as utilization and activity are decreased, peak transmission rate must increase to maintain the long-term average rate constant. It is for this set of source characteristics that the packet transmission format is cost effective. With 20 earth stations in the network, packet transmission is as much as 50-percent less expensive than circuit switching with activity concentration. The packet transmission technique has the largest advantage for spot-beam operation when transponder cost is maximum ($4.0M/year). These same general trends hold as the number of earth stations in the network increase to 40 and then to 100. Note that with 100 stations in the network, packet transmission retains a small but consistent advantage for the cases where transponder annual cost is low. With

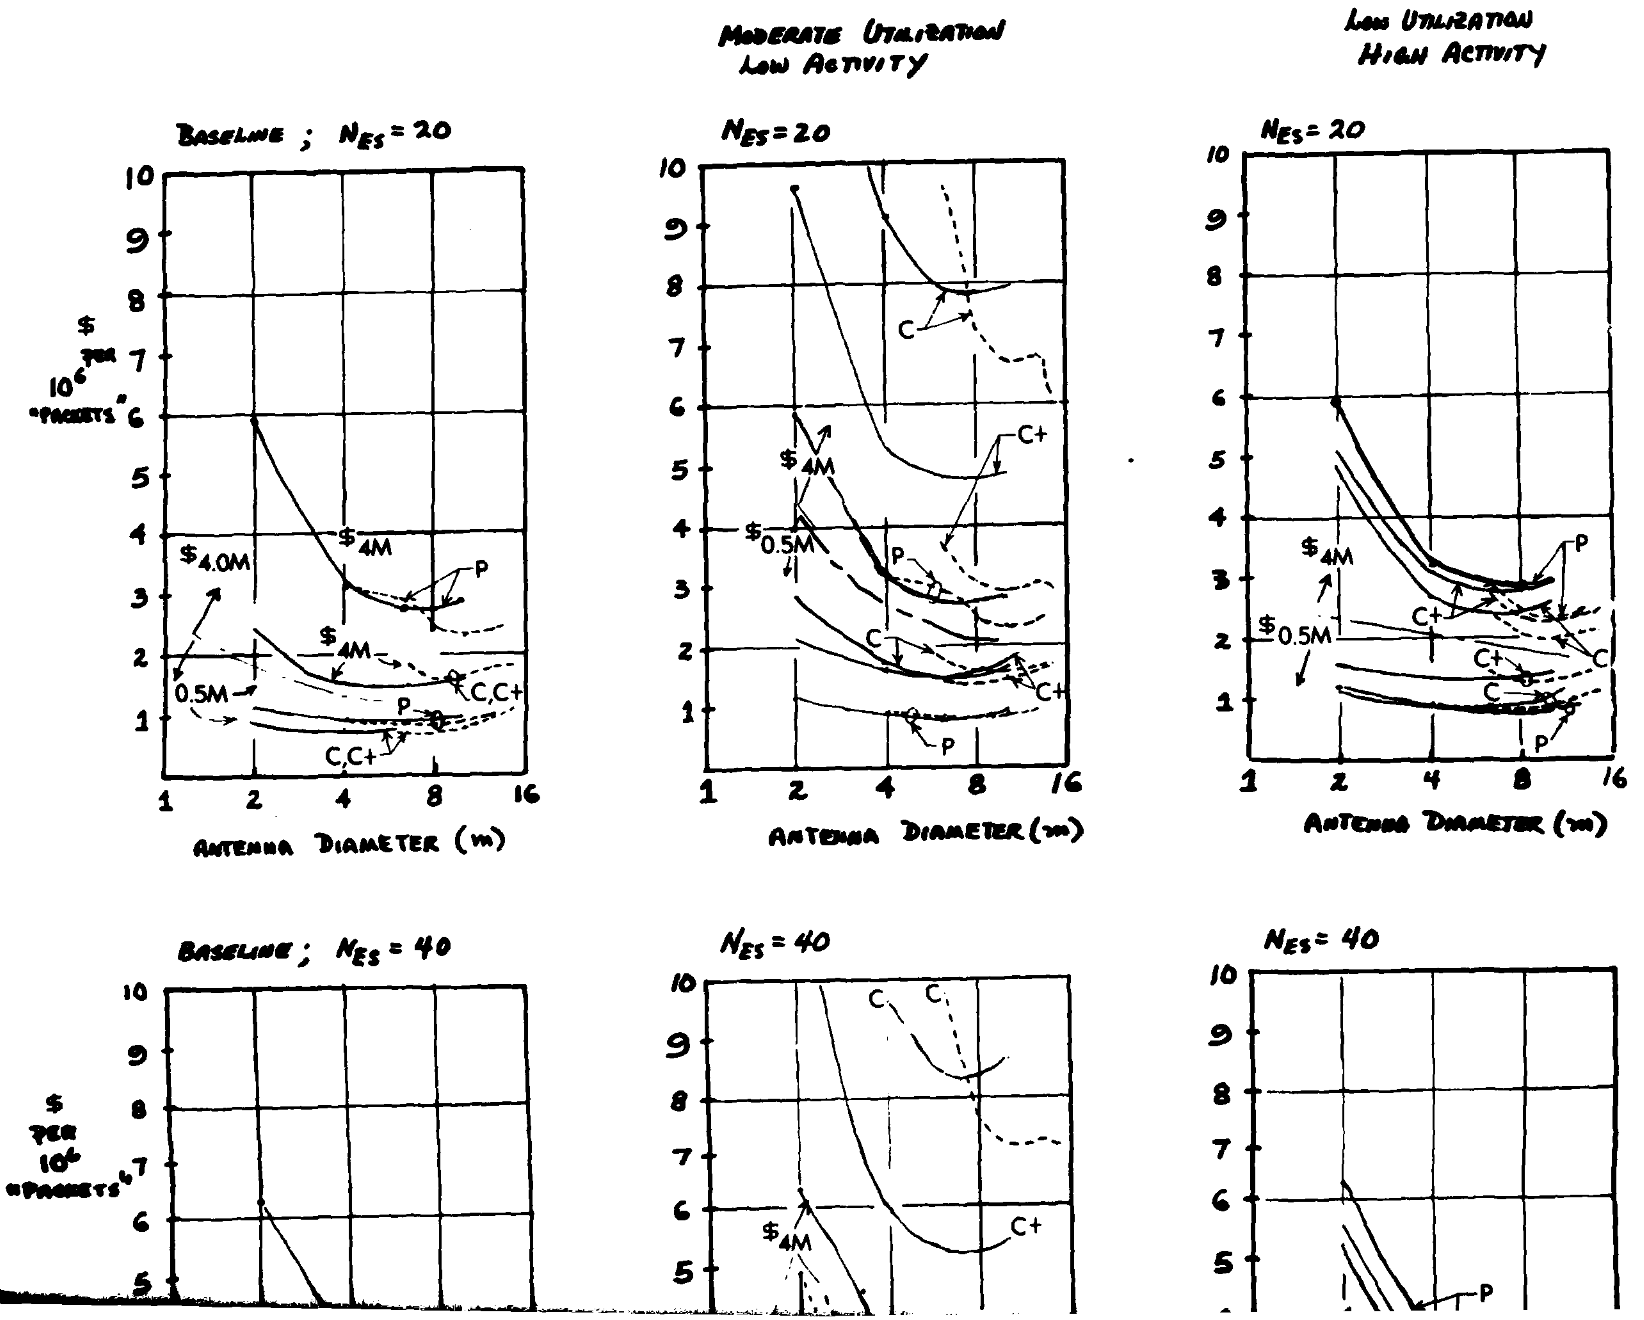

Moderate Utilization
Low Activity
Low Utilization
High Activity
Baseline ; N_{ES} = 20
N_{ES}=20
N_{ES}= 20
$ per 10^6 "Packets"
$4M
$4.0M
0.5M
$0.5M
P
C
C+
C,C+
Antenna Diameter (m)
Baseline ; N_{ES} = 40
N_{ES} = 40
N_{ES}= 40

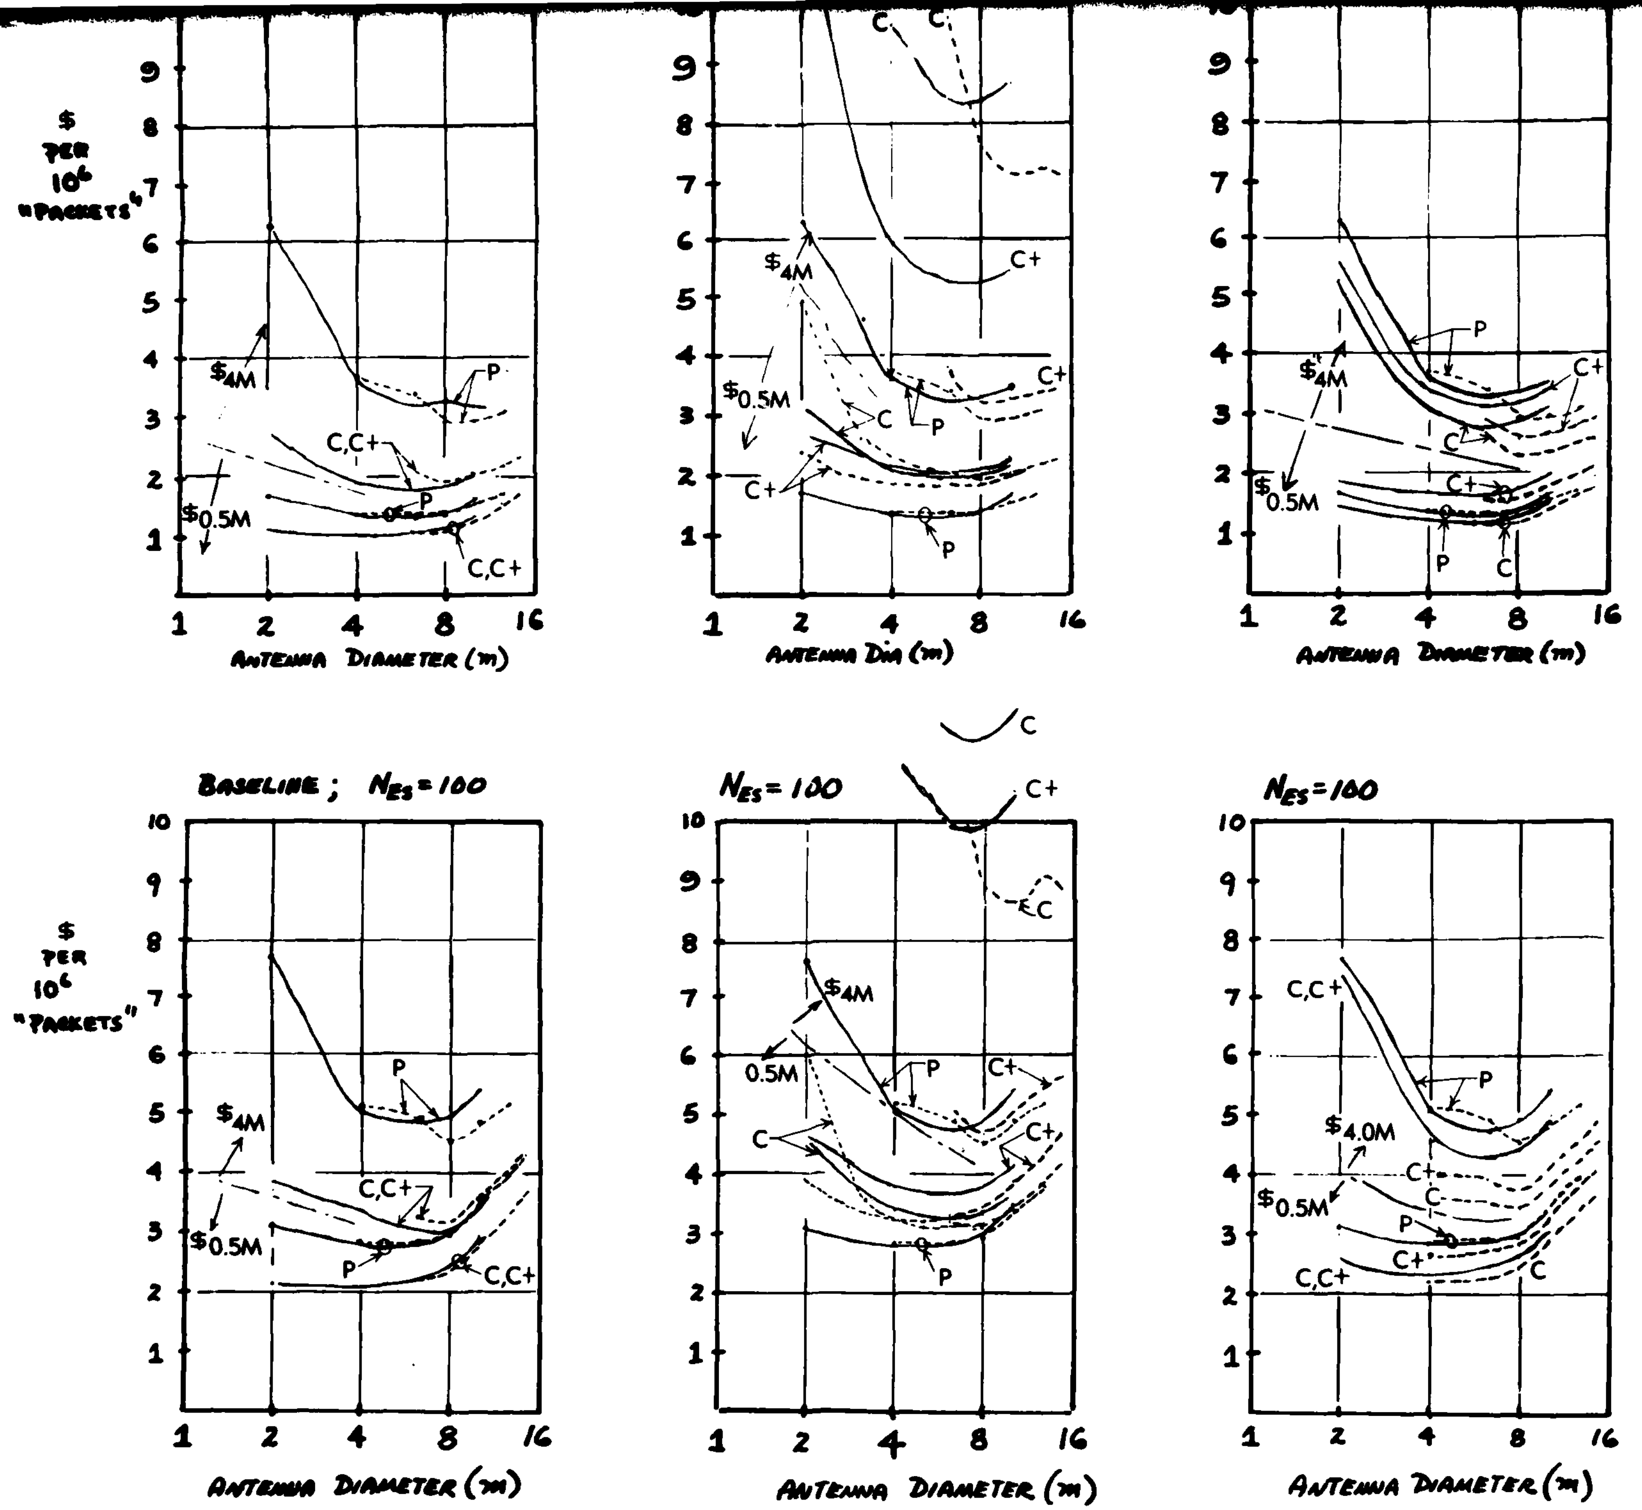

Figure 7-11. Cost/Packet vs Earth Station Antenna Diameter

high annual transponder cost, packet transmission has a significant cost advantage (almost two-to-one) for spot-beam coverage but gives approximately the same cost as circuit switching with activity concentration for CONUS coverage.

The final case shown at the right of Figure 7-11 applies to a situation where the traffic sources have moderate utilization but relatively high activity. For these cases, circuit switching is generally the least-cost transmission technique.

It is now useful to concentrate on case 2 of Table 7-5 and examine in more detail those cases where packet transmission is cost effective. From the results in Figure 7-11, we can focus on those earth station sizes that yield near-minimum transmission costs. At 6/4 GHz, these antenna diameters are in the region 5-8 m; at 14/11 GHz, the range is 4-6 m. Figure 7-12 plots cost per packet versus annual spacecraft cost ($M for a 40-MHz bandwidth increment). The nine different cases are identified by the number of earth stations in the network (from left to right), and from top to bottom; 6/4 GHz CONUS coverage, 6/4 GHz spot coverage, and 14/11 GHz spot coverage. A general conclusion from Figure 7-12 is that a packet transmission gives a lower cost per "packet" than circuit switching with activity concentration. For CONUS coverage, with few earth stations in the network, and for the lowest annual transponder cost, the packet transmission system gives about a 30-percent cost saving. With many earth stations, and transponder cost of $4.0M/year, this advantage is diminished relative to circuit switched system with activity concentration.

For the spot beam systems, the packet transmission system is consistently less costly on a per-packet basis. This advantage becomes more pronounced with many earth stations in the network, and when the transponders have an annual cost of $4.0M.

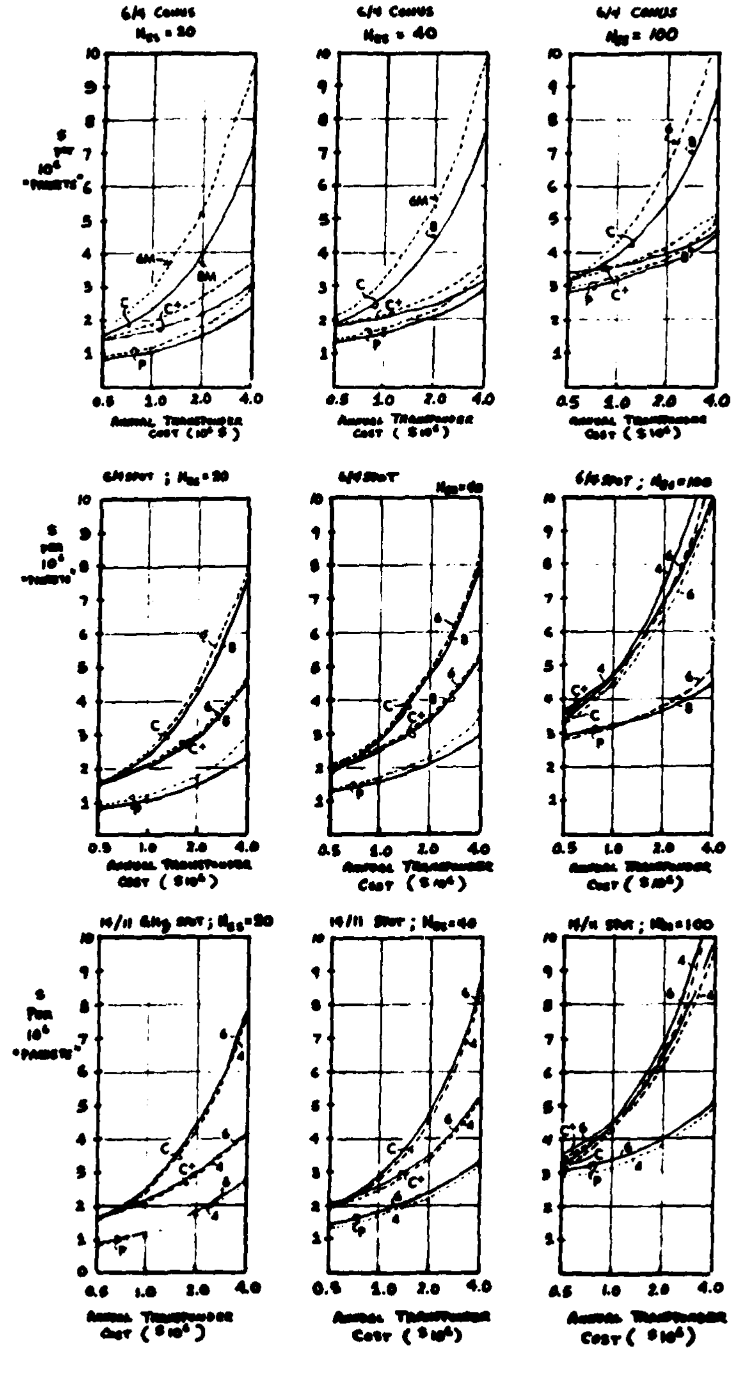

Figure 7-12. Cost/Packet vs Transponder Cost

Note that the trends for 6/4 GHz and 14/11 GHz spot coverage are very similar. (For the models used in this analysis, the differences between the two frequency bands is small. In the following, we will concentrate on 6/4 GHz; the conclusions should also apply to 14/11 GHz with only minor differences.) For spot-beam operation with 20 and 40 earth stations in the network, packet transmission decreases cost by 30 to 40 percent relative to circuit switching with activity concentration. With 100 stations in the network, the reduction is 10 to 50 percent.

As with any analysis of this type, the trends can be sensitive to certain assumptions made in the modeling. This is particularly true for the assumptions used in the cost models for earth station processing equipment. To test this sensitivity, several additional results are shown in Figure 7-13 where the following changes were made. First, the assumptions about packet overhead is reduced approximately 50 percent (header from 275 to 128 bits), and control information by 50 percent). These results are shown in Figure 7-13a. A slight reduction in the cost/packet is noted of from 5-15 percent. In Figure 7-13b, the cost of the packet transmission equipment is doubled and halved relative to its nominal value in the cost model. One can conclude from these results that the advantage of packet transmission relative to circuit switching with activity concentration disappears if we have underestimated the cost of the packet transmission equipment by two-to-one. Finally, in Figure 7-13c, the sensitivity to total network traffic $\hat{P}$ is shown for all the transmission techniques. These results show that the absolute values of cost (dollars per megapacket) is highly dependent upon the total network traffic. When these costs are dominated by the earth station processing equipment costs, as they are in this case for circuit switching with activity concentration, and

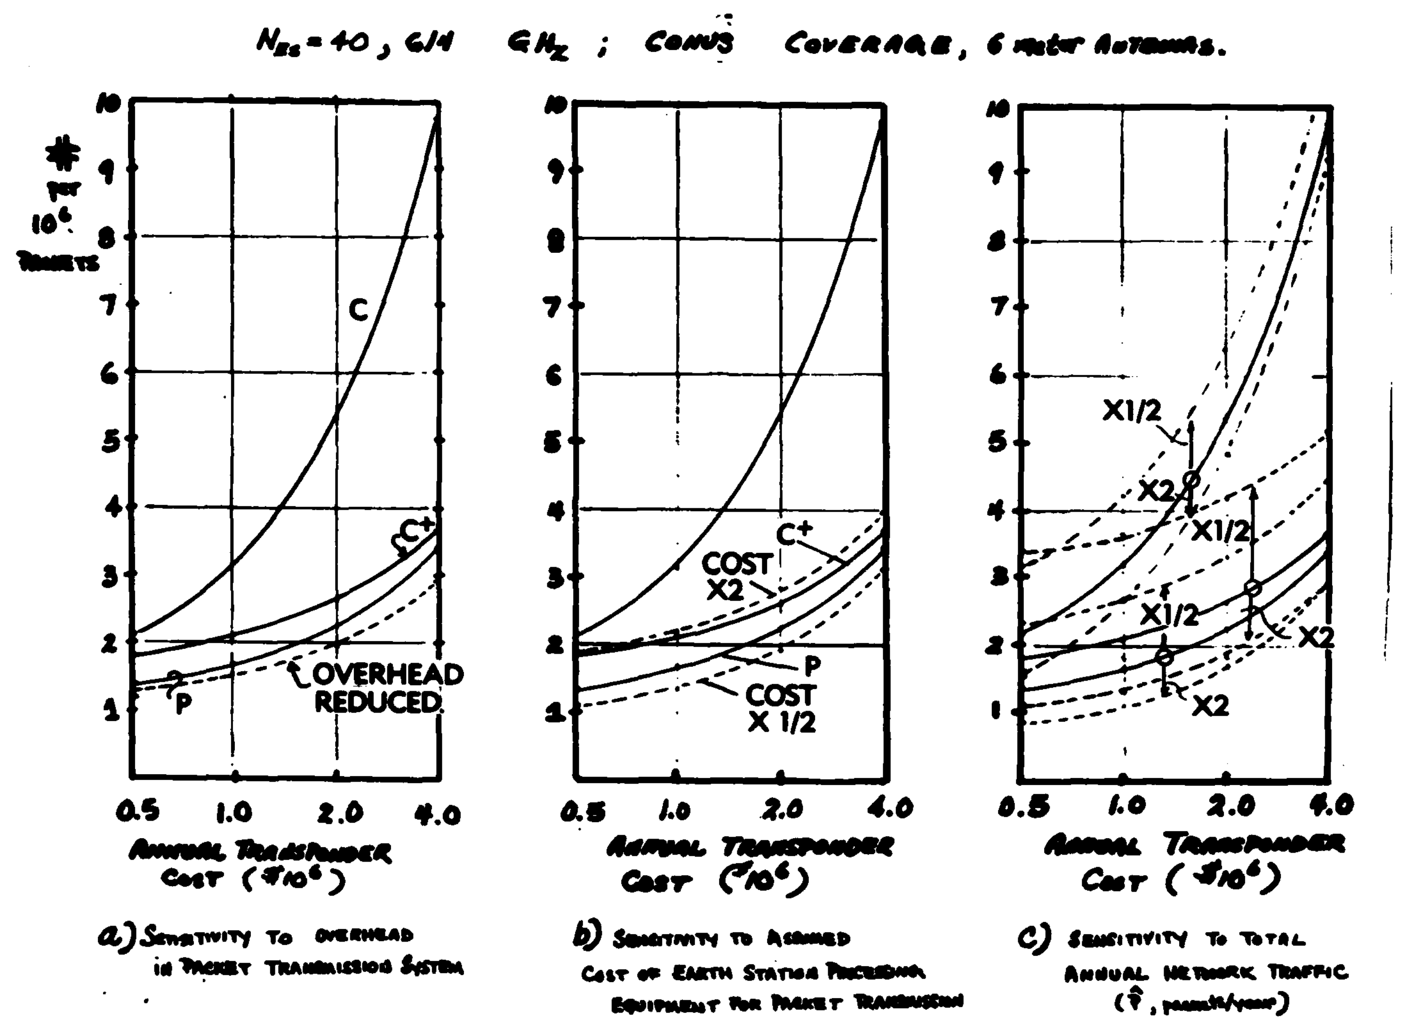

Figure 7-13. Results of Sensitivity Analysis

packet transmission, then cost/packet varies approximately inversely as traffic volume. Conversely, for circuit switching, cost is dominated by the cost of the space segment, particularly when the transponder cost is high. For this situation, cost/packet is much less sensitive to traffic volume.

Although this subsection has given "results" of the tradeoff study, these should be viewed as only representative of the type of tradeoffs that could be examined with further development of the various models. To yield definitive tradeoff results, the following extensions and refinements would be necessary.

a. More detailed cost models for the earth station processing equipment.
b. Improved modeling of the source characteristics (i.e., utilization and activity) and a more detailed examination of the delays expected in packet transmission systems.
c. Consideration of other multiple access techniques (FDMA) in addition to the high burst-rate TDMA technique considered in the model to date.
d. Addition of spacecraft onboard processing, including a more detailed treatment of antenna beam switching and traffic loading.
e. Consideration of the higher frequency bands (14/11 GHz) and the different levels of availability that can be achieved due to propagation disturbances.

8. REFERENCES

[1] Communications Satellite Corporation, First Interim Report (period 5 October 1976 to 30 April 1977) under Contract No. F04701-76-C-0240, date of issue 5 July 1977.

[2] E. Hoversten, "UET Implementation Planning," 1 February 1977 (PSPWN #58).

[3] E. Hoversten, "Phase II SSI Data Transmission Module Operation, Packet Formats, and Timing," October 1977, (PSPWN #91).

[4] Communications Satellite Corporation, Second Interim Report (period 1 May 1977 to 1 June 1978) under Contract No. F04701-76-C-0240, date of issue 15 February 1979.

[5] C. Heegard, J. A. Heller, and A. J. Viterbi, "A Microprocessor-Based PSK Modem for Packet Transmission Over Satellite Channels," *IEEE Transactions on Communications*, COM-26, No. 5, May 1978, pp. 552-564.

[6] I. Jacobs, L. Lee, and A. Viterbi, "CPODA, A Demand Assignment Protocol for SATNET," Fifth Data Communications, Symposium, Snowbird, Utah, September 1977, *Proc.*, pp. 2-5 to 2-9 (also available as PSPWN 75).

[7] L. Lee and I. Jacobs, "A Priority-Oriented Assignment (PODA) Protocol and an Error Recovery Algorithm for Distributively Controlled Packet Satellite Communication Networks," Electronics and Aerospace Systems Conference (EASCON), September 25–28, 1977 (also available as PSPWN 77).

[8] W. Wu, "Synchronization Sequences in Satellite Communications," COMSAT Advanced Systems Report, July 20, 1978.

[9] Packet Switching Services and the Data Communications User, Timothy Johnson, Ovum Ltd., 14 Penn Road, London, England, 1976.

[10] Economic Analysis of Integrated DOD voice and Data Networks, Final Report, Network Analysis Corporation, September 1978.

APPENDIX A. CM&M DESCRIPTION

This appendix contains a more detailed description of the command and monitoring module (CMM).

A.1 CMM SOFTWARE

The following section contains a general description of the command and monitor module (CMM) software. This software is described from the viewpoint of SIMP-CMM commands and their effect on the CMM.

All communications from the SIMP to the CMM takes place via a serial link operating at 200 kHz. This link consists of three wires which serve the following functions:

a. transmit data to the SIMP,
b. receive data from the SIMP, and
c. establish a serial clock from the SIMP.

Upon reception of a byte of data from the SIMP, a hardware interrupt is generated within the CMM processor. This causes the CMM to read the data byte, test the data for errors (parity and overrun) and, if no errors are detected, to echo the byte received. The CMM then starts a 20 ms timer. Upon reception of the CMM echo, the SIMP compares this byte with the command byte just transmitted. If the two bytes are identical the SIMP will not send a new command for at least 20 ms. If they are not the same, then the SIMP must retransmit the previous command within 20 ms.

Once the 20 ms interval from the previous command has expired, the CMM assumes that the command is valid and proceeds to process it. All SIMP commands consist of one byte of information. The upper nibble identifies the command function (thus there are 16 possible types of commands) while the lower nibble contains a subfunction which varies from command to command. Table A-1 contains a list of the various commands that the CMM executes and gives their respective formats. These commands will be discussed in greater detail in the following sections.

A.2 SYSTEM SET UP COMMAND (FUNCTION 0)

The system set up commands are used to place the CMM in a predefined initial state. There are two types of set up commands presently implemented in the CMM; these are system RESET and system RESTART. As noted in Table A-1, the function code for this command consists of all zeros, while the subfunction contains a zero in bit D3 for system reset and a 1 in D3 for system restart. The system restart command is also executed automatically on the occurrence of a hardware power on or a hardware reset. The state in which the CMM is placed after the execution of these commands is given in Table A-2.

A.3 PREAMBLE SELECT COMMANDS (FUNCTIONS 1 THROUGH 3)

The preamble select commands allow the SIMP to change:

a. the carrier recovery length select (function code 1),
b. the large station bit timing recovery length select (function code 2), and
c. the small station bit timing recovery length select (function code 3).

Table A-1. CMM Commands

Routine Name	Type	Bit Code (SIMP Oriented)			Function		Device Name	Path	ADDR	Values	Bits	No. of SIMP Words
		2^0 2^1 2^2	2^3	2^4 2^5 2^6 2^7*	CMD Name	(CMD) No.						
Reset	S	0 0 0	0	0 0 0 0	System Reset	#0	Whole System	All Modules		refer to reset and restart documentation		
Restart	S	0 0 0	1	0 0 0 0	System Restart	#0						
CARREC	S	X X X	X	1 0 0 0	Carrier Recovery LNG Sel.	01	XSSIL	.IO 8 Bit Reg	8E00	16		1
LSBITREC	S	X X X	X	0 1 0 0	Large Station Bit Time		XSSIL					
					Recovery LNG Sel.	#02		.IO 8 Bit Reg	8E01	16		1
SSBITREC	S	X X X	X	1 1 0 0	Small Station Bit Time		XSSIL					
					Rec. LNG Sel.	03		.IO 8 Bit Rec	8E02	16		1
SOMSEL	S	X X X	X	0 0 1 0	SOM Sel. (LNG and PTRN)	04	XSSIL	.IO 8 Bit Reg		4 SOM1	4	1
					D3 = 0 SOM PATTERN Sel.			SOM1 4 MSB	8E04	4 SOM2	4	1
					D3 = 1 SOM Length Sel.			SOM2 4 LSB	8E04	4 SOM PTR	4	1
								PTRN 4 LSB	8E05			1
FRQCHN	S	X X X	X	1 0 1 0	NOT USED	05						
DISCTL	S	X X X	X	0 1 1 0	DTC Control	06	DTS	.IO 8 Bit Reg	8E05	16	4	1
					D0 Fast/Slow			4 MSB				
					D1 Free Run							
					D2 Fixed Pkts # active high							
					D3 Single Pkt							
SWIMAX	S	X X X	X	1 1 1 0	Switch Matrix Mode	07	Switch Matrix	.IO 8 Bit Reg	8E06	16	4	1
								4 LSB				
SOMDET	S	X X 0	0	0 0 0 1	SOM Det Threshold	08	RCV SSIC	IF Latch 2^0-	8046	4	2	1
					D0 = 2^0			2^1-	8047			
					D1 = 2^1							
MONSSIC		X X X	X	1 0 0 1	Monitor SSIL PARM	09		IO 8 Bit Reg	8005			
		1 0 0	X		Monitor PRSOMR0		Xmit SSIC	LSB = Preamble		3	3	1
		0 1 0	0		Monitor PRSOMR1		Xmit SSIL	D1, D2 = $2^0 2^1$ SOM#	8E00	5	5	1
		0 0 1	0		SOMERRA		RCV SSIC	5 LSB of	8E04	5	5	1
		0 0 0	1		SOMERRB		RCV SSIC	4 8 Bit Reg	8E04	5	5	1
LBCMDP				0 1 0 1	Command SSIL	10						1
		0 0 0	1/0		XSSIL/RSSIL ECHO 0AAH/055H		XSSIL	IO 8 Bit Reg	8E07			1
		1 0 0	1/0		Reception Control Heard							
					Enable/Disable		RSSIL	2 Strobe Lines				
		1 1 0	X		Reset XSSIL/RSSIL		XSSIL	XSSIL	8004			
		0 0 1	X		XSSIL Interface Enable/Disable		RSSIL	RSSIL	8006			
		1 0 1	1/0		XSSIL-CMM Interface Enable/		XSSIL					
					Disable		RSSIL					
		1 1 1			RSSIL-CMM Interface Enable/							
					Disable							
LBCMDM	S/T	0 0 0	0	1 1 0 1	Monitor AGC	11	RSSIL	1 8 Bit IO Port	8E06	256	8	1
					READ AGC							
	S	X X X	X	0 0 1 1	SIMP Monitor	12	C&M	Serial Link				
		1 1 1	1		ECHO			to SIMP	E409			~40M
		1 1 1	0		Start Monitor							
		1 1 1	0		Continue							
			S/R	1 0 1 1	Latches set = 1, Reset = 0	13		IF Latches				
LATCH	S	0 0 0			Normal/Bypass		XSSIc		8041	2	1	1
		1 0 0			DTS Start/Stop		DTS		8042	2	1	1
		0 1 0			DTS Ready/Clear		DTS		8043	2	1	1
		1 1 0			Normal/Loopback On/Off		RCV SSIC		8044	2	1	1
		0 0 1	0		Single Pkt Strobe		DTS		804d	2	1	1
		1 0 1			Modem Loopback On/Off		IF Subsystem		8049	2	1	
		0 1 1			Data Source DTS/SIMP		Switch Matrix		804A	2	1	1
		1 1 1			Data Source Select		Switch Matrix		804B	2	1	1
					0 SIMP/SIMP							
					1 SIMP/DTS							
					NOT USED	14						
			0	1 1 1 1	SIMP Info/Response LDR	15						
			1	1 1 1 1	SIMP Info/Response Data							
					(not defined or implemented)							

*2^7 = most significant bit.

Table A-2. System Configuration After "Restart" and After "Reset"

Function	System Restart	System Reset
Stack Pointer	Reset	Unchanged
ACIA	Reset	Unchanged
Programmable Timer	Reset	Unchanged
SSIL Monitor Interrupts	Reset	Unchanged
Switch Matrix	Mode 01	Mode 01
Automatic Load SOM	Set	Unchanged
Carrier Recovery Length	30H	Unchanged
LS Bit Time Recovery Length	30H	Unchanged
SS Bit Time Recovery Length	30H	Unchanged
SOM Pattern	0FH	0FH
SOM1 Length	1	1
SOM2 Length	1	1
Normal/$\overline{\text{Bypass}}$	Normal	Normal
Modem Loopback	Off	Unchanged
DTS $\overline{\text{Ready}}$/Clear	01 = clear	clear
DTS $\overline{\text{Start}}$/Stop	01 = stop	Unchanged
RSSIC Loopback $\overline{\text{On}}$/Off	Off	Off
DTS Single Packet Start	Set	Unchanged
DTS Test Mode		Unchanged
DTS Source Select	SIMP/DTS	Unchanged
DTS Mode		
D0 Auto Fast/$\overline{\text{Slow}}$	Set	
D1 Auto Free/Run	Set	
D2 Auto Fixed/No	Reset	
D3 Auto Single/Packet	Reset	
Linkabit Modem	Reset	Reset
All Sources of Interrupt	Reset	Unchanged
XSSIL and RSSIL Reset	Reset	Reset
SSIL Echo Variables		
XSSILSR	0	Unchanged
RSSILSR	0	Unchanged

This is done by setting the correct funtion code in the upper nibble of the SIMP command and using the lower nibble to choose one of the 16 possible values located in the tables in CMM memory. The contents of these tables is listed in Table A-3.

Table A-3. CMM Preamble

Subfunction	LNG Select		
	Carrier Rec.	LS Bit Time	SS Bit Time
0	3CH	3CH	3CH
1	01H	60H	60H
2	28H	36H	36H
3	01H	60H	60H
4	01H	60H	60H
5	01H	60H	60H
6	01H	60H	60H
7	01H	60H	60H
8	01H	60H	60H
9	01H	60H	60H
10	01H	60H	60H
11	01H	60H	60H
12	01H	60H	60H
13	01H	60H	60H
14	01H	60H	60H
15	01H	60H	60H

A.4 SOM LENGTH AND PATTERN SELECT (FUNCTION CODE 04)

The SOM length and PATTERN select command is used to choose one of 8 possible SOM patterns and one of 4 possible SOM lengths for SOM1 and SOM2. This is done by using the subfunction portion of the SIMP command to identify one of 16 possible locations in a SOM select table located in CMM EPROM. Depending upon which location is chosen, the CMM changes SOM1, SOM2 or the SOM pattern. The format for the subfunction bits of the SIMP command is given in Table A-4 while the actual contents of the internal CMM table are given in Table A-5.

The CMM internal table at present contains only one value for the SOM patterns and length as that is all that has been defined. It should be noted that SOM1 and SOM2 length select are 4 bit codes being outputted via an 8-bit register. Thus, the values of SOM2 length are located in the upper nibble of loactions 0CFB8H—0CFBBH, while the values permissible for SOM1 length select are located in the lower nibble of locations 0CFBCH—0CFBFH.

A.5 FUNCTION CODE 05

Not used.

Table A-4. SOM Select Subfunction

D3	1 SOM Length Select 0 SOM Pattern Select
D2	1 SOM1 Length Select 0 SOM2 Length Select
D0-D1	Table Address Identifier

Table A-5. SOM Select Table
(EPROM ADDR 0CFB0H-0CFBFH)

SOM Pattern	SOM2 Length	SOM1 Length
CFB0H 0FH	CFB8H 10H	CFBCH 01H
0FH	10H	01H
0FH	10H	01H
0FH	10H	01H
0FH		
0FH		
0FH		
0FH		

A.6 DATA TEST SET CONTROL (FUNCTION CODE 06)

The data test set control command is used to configure the data test set (DTS). The subfunction bits of this SIMP command are used to set 4 DTS control lines to their correct state. These control lines and the appropriate bit that controls their state is given in Table A-6.

Table A-6. DTS Control

Bit	Function
D0	Fast/$\overline{\text{slow}}$ station
D1	Free run
D2	Fixed number packets
D3	Single Packet

It should be noted that bits D1–D3 are active high signals and that each time this command is executed the state of all four control lines must be specified.

A.7 SWITCH MATRIX CONTROL (FUNCTION CODE 07)

The switch matrix command from the SIMP allows the PSP terminal to be placed in any one of 16 possible modes. The desired mode is specified in binary notation in the subfunction portion of this command. It should be noted that the CMM does not test the subfunction to ensure that the mode is permissible within an appropriate range. It is up to the user to ensure that the subfunction is coded to a correct value. The value actually outputted by the CMM to the switch matrix is in nines complement offset three notation.

A.8 DETECT THRESHOLD (FUNCTION CODE 08)

The SOM detect threshold command allows the SIMP to specify the allowable number of correlation errors when detecting a valid SOM. This number may be anywhere from 2 to 5 errors.

The actual number of errors to be tolerated is specified in the two least significant bits of the command subfunction code. Bits D2 and D3 are presently ignored.

A.9 MONITOR SSIC PARAMETERS

The monitor SSIC command is used to allow the SIMP to request and receive SSIC parameters. At present, the SIMP has the capability of requesting any one of 4 SSIC parameters. These are the following:

a. preamble, SOM1 number
b. preamble, SOM2 number
c. SOM correlation error A, and
d. SOM correlation error B.

The first two parameters have their useful information contained in the three least significant bits of the byte of information transferred to the SIMP. The least significant bit specifies whether the PSP is communicating with a large station (D0 = 1) or a small station (D0 = 0). The next two least significant bits (D1-D2) identify which of four possible SOM numbers is being used.

The last two parameters have their useful information contained in all bits except bit D5. These bytes contain the previous correlation error detected by the CMM. The two most significant bits of these bytes identify whether the correlation error was detected on the AA BB channel (D7 = 0 D6 = 1) or the AB BA channel (D7 = 1 D6 = 0. The five least significant bits give the actual number of errors detected for each channel.

A.10 COMMAND SSIL (FUNCTION CODE 10)

The SIMP has the capability of controlling the SSIL interface by the command SSIL. The functions that this command allows the SIMP to control are:

0) $\overline{\text{XSSIL}}$/RSSIL echo command for 0AAH/055H, respectively
1) reception control packet heard enable/$\overline{\text{disable}}$
2) reception dummy packet and modem and interface T&M words enable/$\overline{\text{disable}}$
3) reset RSSIL or reset XSSIL interface
4) Linkabit interface XSSIL enable/$\overline{\text{disable}}$
5) XSSIL interface enable/$\overline{\text{disable}}$ capabilities of CMM
6) not used
7) RSSIL interface enable/$\overline{\text{disable}}$ capabilities of CMM

The appropriate command is chosen by setting bits D0-D2 of the subfunction to the binary representation of the command number. Bit D3 of the subfunction is set to 1 for active function of each command and is set to zero for the inactive function for the case of command 03; D3 = 1, specifies resetting RSSIL while D3 = 0 specifies resetting XSSIL.

Upon power up, the Linkabit interfaces are active while C&M command function is disabled. Thus, before any other commands are sent to the Linkabit interfaces, a C&M enable command (commands 5 and 7 with D3 = 1) must be executed from the SIMP.

It should also be noted that the dummy packet referred to in command 02, and also the transmit dummy packet which is transmitted automatically when the transmitter identifies a contention condition, has not been defined to date. These packets are 16 words in length and a definition of the contents of these packets including check sums remains to be defined by BBN.

A.11 MONITOR SSIL COMMAND (FUNCTION CODE 11)

The monitor SSIL command at present allows the CMM the ability to read the receive SSIL automatic gain control. The AGC value read will be the value at the time the SSIL processed the command. In order for the SIMP to get this value, it is necessary that the SIMP follow this command with a command for the CMM to dump its registers to the SIMP (Function code 12).

A.12 SIMP MONITOR COMMAND (FUNCTION CODE 12)

The SIMP monitor command is used to allow the SIMP to check the state of the CMM. The first subfunction of this command (subfunction code = 1111B), called ECHO, is used to test the SIMP CMM communications link. Upon reception of this command, the CMM will retransmit the ECHO byte but will take no other action.

The second subfunction of this command, START MONITOR/CONTINUE, is used to give the CMM the ability to dump its internal registers to the SIMP. Upon reception of the first SIMP monitor command with the subfunction bits set to 0111B, the CMM echos the command and sets itself up for a data dump. Upon reception of each succeeding SIMP monitor command, with subfunction bits set to 0111B, the CMM, instead of echoing the command, transmits the next byte of its internal registers. The START MONITOR/CONTINUE command may be terminated by either:

a. allowing it to transmit all its internal registers (30H bytes), or

b. having the SIMP transmit some other command which is not a continue. This causes the CMM to abort its dump of its data registers, echo the command and prepare to execute the command after the specified timeout.

The data variables the CMM transfers to the SIMP, upon reception of a START MONITOR/CONTINUE command, and their meaning is given in Table A-7. It should be noted that unless the particular bits of these registers are specified they are undefined. Also due to the nature of the reception mechanism at the SIMP, these bytes are read bit reversed and complemented. That is to say the most significant bit shown at the SIMP is actually the complement of the least significant bit at the CMM. The bits referred to in this documentation are with reference to the CMM perspective.

A.13 LATCH COMMAND (FUNCTION CODE 13)

The CMM latch command is used to set or reset one of eight control lines of the CMM. The specific control line and its state is identified by the subfunction code of the command. Bits D0—D2 identify which latch is to be controlled while bit D4 is set to the desired state. The possible control functions that this command affects are listed in Table A-8. For all functions, with the exception of the single packet strobe, once the latch is specified in a certain state it remains there until a restart or another command changes its state. The single packet strobe subfunction must have bit D3 = 0. When this control line is specified, the CMM uses D3 = 0 to form a negative going pulse which strobes the data test set. If the command is given with D3 = 1.

a pulse of the wrong polarity will result. It should be noted that the CMM does not record the state of the various control lines. Thus, it is up to the user to set the control lines to their appropriate state if they differ from those defined by a restart of reset command.

Table A-7. Description of CMM Data Registers

Byte	Generic	Description
0	SOMERRSR	SOM correlation error indicates interface the A and B modem data was received in. Bits D0 and D1 contained all desired information D0 = 1 D1 = 0 A data received A channel B data received B channel D0 = 0 D1 = 1 A data received B channel B data received A channel This value indicates where the actual correlation errors were detected
1	SOMERRA	SOM AA/AB correlation error. SOMERRSK (D0–D4) specifies which correlation error is located in this byte. This value is updated with each QPSK SOM detect
2	SOMERRB	SOM BB/BA correlation error. SOMERRSK (D0–D4) specifies which correlation error is located in this byte. This value is updated with each QPSK SOM detect
3	XSSILSR	Storage register containing last byte of data transferred from the XSSIL which was unsolicited by the CMM
4	RSSILSR	Storage register containing last byte of data transferred from the RSSIL which was unsolicited by the CMM

Table A-7. Description of CMM Data Registers
(Continued)

Byte	Generic	Description
5	MINDXLB	Monitor index for whether a monitor command is outstanding for XXIL (MINDLB = 01) or for RSSIL (MINDLB = 02). Register is not used at present
6	SOMDETSR	SOM detect storage register. D0–D1 contains the value of the last SOM correlation value sent to the PSP (see Subsection 1.6)
7 8 9	FRQSNR0 FRQSNR1 FRQSNR2	At present, these three bytes are not used. The lower nibble of these registers will contain the last values sent to the frequency synthesizer
10 (A)	PRSOMR0	Preamble SOM register 0 contains the value of the last preamble SOM number read. D0 = 10 = large/small station preamble. Numbers D1–D2 = SOM number used. This value is updated once a packet for the first SOM in each packet
11 (B)	PRSOMR1	Same as PRSOMR0 only it contains the value for the second SOM, if any, in each packet
12 (C)	CARRECSR	Carrier recovery storage register. D0–D7 contains the value of the last carrier recovery length loaded into the PSP terminal
13 (D)	LSBITRSR	Large station bit timing recovery storage register. D0–D7 contain the value of the last large station bit timing recovery length loaded into the PSP terminal

Table A-7. Description of CMM Data Registers (Continued)

Byte	Generic	Description
14 (E)	SSBITRSR	Small station bit timing recovery storage register D0–D7 contain the value of the last small station bit timing recovery length loaded into the PSP terminal
15 (F)	SIMPBC	Word count of the number of bytes remaining for the CMM to transmit in a START MONITOR/CONTINUE command sequence
16–17 (10H–11H)	SIMPPT- SIMPPT1	High low address pointer for the next byte, the CMM will transmit in a START MONITOR/CONTINUE command sequence
18 (12H)	SIMPFLG	byte indicator as to wheather the CMM is in the middle of a START MONITOR/CONTINUE (SIMPLFG = 1) command sequence or not (SIMPFG = 0)
19 (13H)	MSGWDCTR	Message word counter used in multiword command sequences with SIMP (presently not implemented)
20 (14H)	DATAWCTR	Data word counter used in multiword command sequences with SIMP (presently not implemented)
21 (15H)	ACIADT	AC1A data storage register. D0–D7 of this register contains the last byte of data received from the SIMP
22 (16H)	Not Used	
23 (17H)	SOMLNGSR	Same as SOMSELSR
24 (18H)	SOMPTRSR	SOM pattern storage register. Bit D0–D3 contain the value of the last SOM pattern loaded into the PSP terminal
25 (19H)	SINFOFLG	SIMP information message flag (presently not implemented) used in multibyte commands from the SIMP to the CMM

Table A-7. Description of CMM Data Registers (Continued)

Byte	Generic	Description
26–31 (1AH–1FH)	Not Used	
32 (20H)	RMFLGLB	Reception monitor flag for RSSIL. When this flag equals 01 it means that a read AGC command has been given to the SIMP and that the next interrupr from the RSSIL monitor port will specify the AGC data may be read. If the flag = 0 no such command has been given
33 (21H)	CFLGLB	Command flag for Linkabit command port. When this flag is 01 it specifies that the CMM is in the process of executing a command to Linkabit for the SIMP. If this flag = 0, no such command is being executed. As long as this flag equals 01 no new command will be accepted for SSIL
34 (22H)	SFLGLB	Status flag for Linkabit command port. When this flag = 0, the CMM is able to execute a CMM-Linkabit command. When this flag = 01, it means that the CMM tried to execute a CMM-Linkabit command and the Linkabit interface took longer than 2.2 seconds to respond. This flag is reset to zero as soon as the CMM-Linkabit interface is responsive within 2.2 seconds
35 (23H)	CMDINDX	Command index specifies whether the next command byte to be output to the XSSIL (CMDINDX = 01) or the RSSIL (CMDINDX = 02) command port.

Table A-7. Description of CMM Data Registers (Continued)

Byte	Generic	Description
36 (24H)	SOMSELSR	SOM select storage register (D0–D7). The register contains the most recent SOM length loaded into the PSP. The upper nibble contains SOM1 length select while the lower nibble contains SOM2 length select
37 (25H)- 38 (26H)	CMDPTR CMDPTR1	Pointer to the next byte to be output to the SSIL command port
39 (27H)	CMDSTE	Register containing the number of bytes remaining to be transferred to the SSIL command port
40 (28H)- 41 (29H)	WAITTM	Two byte count used to indicate amount of time it takes the SSIL to respond to a command. When this value is counted down to 0, SFLGLB is set to the index of the Linkabit interface which timed out
42 (2AH)	DTSCTLSR	Data test set control storage register bits D0–D4 contain value of last data test set control command executed
43 (2BH)	SW1MAXSR	Switch matrix storage register. Bits D0–D3 contain value of last mode set in switch matrix
44(2C)–48(30H)	Unused	

Table A-8. CMM Latch Functions

Command Numbers (D0 D1 D3)	Latched Control Line
0	Normal/$\overline{\text{bypass}}$
1	DTS $\overline{\text{Start}}$/stop
2	DTS $\overline{\text{ready}}$/clear
3	Normal/$\overline{\text{loopback}}$
4	Single packet strobe
5	Modem loopback $\overline{\text{on}}$/off
6	Data source DTS/$\overline{\text{SIMP}}$
7	Data source select 0 SIMP/SIMP 1 SIMP/DTS

A.14 FUNCTION CODE 14

Not used.

A.15 MULTIWORD SIMP/CMM COMMAND (FUNCTION CODE 15)

Not defined.

APPENDIX B

LISTING OF TRADEOFF COMPUTER MODEL

A listing of the computer program used in the tradeoff study is given in this Appendix. The mainline program is given in Figure B-1. This program varies certain parameters of the system and calculates cost per 10^6 packets. Data to vary network characteristics is entered as block data and an example is shown in Figure B-8. Subroutine ASI( ), Figure B-2, computes adjacent satellite interference levels. Subroutine BPSHZ( ), Figure B-3, determines channel effeciency in bps/Hz given input C/N. The cost of the earth station, excluding signal processing and switching equipment, is determined by CES3( ), which is listed in Figure B-4. Subroutine CAPSAT( ) determines total spacecraft capacity needed for the three transmission techniques, viz; packet switching, circuit switching, and circuit switching with activity concentration. The subroutine BLOCK( ) determines output channels needed, given input channels, for certain "blocking" or "freeze-out" probabilities. Finally, subroutine ESCOST( ), listed in Figure B-7, calculates the cost of the earth station switching and processing equipment for the three transmission techniques.

```
C      MAIN PROGRAM
C      CALCULATES $/MEGAPACKET
C      AS A FUNCTION OF SEVERAL
C      PARAMETERS AND SWITCHING
C      SCHEMES
       COMMON/COM2/ TCKTS(3),BCKTS(3),ZUES(3),RBPKUU(3),DSI(3),PHAT
       COMMON/COM1/ FU(3),PU(3),UTIL(3),ACT(3),QBEAM,ZES
       DIMENSION DIAM(5,2,3)
       DIMENSION PWR(4),TEMP(3),SPOT(3,2),SATCAP(3)
       DIMENSION TSCOST(4,3),SCCOST(4)
       DIMENSION ESC(3)
       DIMENSION UL(3),DL(3)
       DIMENSION FUP(4),FDN(4),SATSEP(4)
       DIMENSION COCH(2)
       DIMENSION A1(3)
C  E.S. ANT. DIAMETER (METERS)  DIAM(INCREMENT,BEAM SIZE,KBND)
       DATA DIAM/4.,6.,8.,10.,12.,
      2          2.,4.,6.,8.,10.,
      3          4.,6.,8.,10.,12.,
      4          2.,4.,6.,8.,10.,
      5          4.,6.,8.,10.,12.,
      6          1.5,2.,4.,6.,8./
C  E.S. POWER DENSITY (W/MHZ)
       DATA PWR/5.,10.,20.,40./
C  E.S. LNA TEMP F(KBND)
       DATA TEMP/120.,200.,300./
C  SATELITE SPOT SIZE F(CONUS OR SPOT,KBND) CONUS-->7 DEG.
       DATA SPOT/7.,7.,7.,1.,0.75,0.5/
C  UPLINK AND DOWNLINK LOSS F(KBND)
       DATA UL/0.94,0.94,0.32/
       DATA DL/0.99,0.99,0.5/
C  ARRAY OF TRANSPONDER COSTS IN M$/40MHZ/YEAR
       DATA SCCOST/0.5,1.,2.,4./
C  FREQUENCY VALUES RELATIVE TO KBND (GHZ)
       DATA FUP/6.0,8.0,14.,30./ ,FDN/4.0,7.0,11.,20. /
C  SATELLITE SEPARATION SATSET(KBND)
       DATA SATSEP /4.0,4.0,3.0,3.0  /
C COCHANNEL INTERFERENCE FACTORS
       DATA COCH/300.,24./
C  PEAK BIT RATE
       RBIT=60.E+6
       RB=40.
C  EQUIV. SAT. TEMP.
       TSAT=2000.
C  SAT. PWR
       PSAT=0.25
C  PROBABILITY OF BIT ERROR (10**(-IPE))
       IPE=6
C  E.S. ANNUALIZATION COST FACTOR
       AF=0.3
C  FREQUENCY BAND  1=6/4  2=8/7  3=14/11 GHZ
C  RUN FOR 6/4 AND 14/11 GHZ BANDS
       DO 100 J=1,3,2
       KBND=J
```

Figure B-1(a)

```
C  SET E.S. EQUIVALENT TEMPERATURE
      TES=TEMP(KBND)
      FRU=FUP(KBND)
      FD=FDN(KBND)
      TH=SATSEP(KBND)
C  CONUS (7 DEG.) OR SPOT BEAMS
      DO 200 K2=1,2
      CCI=10.**(-0.1*COCH(K2))
      GO TO (5,10),K2
    5 NBEAM=1
      GO TO 15
   10 NBEAM=QBEAM
   15 THETA=SPOT(KBND,K2)
C  VARY E.S. DIAMETER
      DO 300 K=1,5
      D=DIAM(K,K2,KBND)
      L=3
      PES=PWR(L)
      CNUP=(4.14E+4*PES*UL(KBND)*(D**2.))/((THETA**2.)*TSAT)
      CALL ASI(FRU,D,TH,XIU)
      CNUP=1./(1./CNUP+CCI+2.*XIU)
      CNUPDB=10*ALOG10(CNUP)
      TSKT=50.+290.*(1-DL(KBND))
      CNDN=(4.14E+4*PSAT*DL(KBND)*(D**2.))/((THETA**2.)*(TES+TSKT))
      CALL ASI(FD,D,TH,XID)
      CNDN=1./(1./CNDN+CCI+2.*XID)
      CNDNDB=10*ALOG10(CNDN)
C  C/N TOTAL
      CNT=((1./CNUP)+(1./CNDN))**(-1.)
      CNTDB=10*ALOG10(CNT)
      CALL BPSHZ(CNT,IFE,R,MM,PZ)
      RATE=R
      CALL CAPSAT(NBEAM,RBIT,K2,RATE,PZ,SATCAP,PACKDY)
      CALL CES3(KBND,D,PES,RB,TES,A1)
      CALL ESCOST(A1,ESC)
C  CALCULATE $/MEGAPACKET VALUES
      XPHAT=PHAT/1.E+6
      DO 500 M2=1,4
      DO 600 N2=1,3
      TSCOST(M2,N2)=((ZES*ESC(N2)*AF)+SCCOST(M2)*1.E+6*SATCAP(N2))/XPHAT
  600 CONTINUE
  500 CONTINUE
C  OUTPUT VALUES
      A2=A1(1)/1000.
      WRITE(6,66)D,PES,CNUPDB,`NDNDB,CNTDB,R,MM,PZ,THETA
   66 FORMAT(1X,F4.1,1X,F4.1,1,,F5.1,1X,F5.1,1X,F5.1,1X,F5.2,1X,I3,1X,
     1        F5.2,1X,F5.2)
      DO 700 K3=1,3
      ZESC=(ESC(K3)*0.3)/1000.
      WRITE(6,67)K3,ZESC,SATCAP(K3),TSCOST(1,K3),TSCOST(2,K3),
     1              TSCOST(3,K3),TSCOST(4,K3)
   67 FORMAT(62X,I1,1X,2(F8.2,1X),4(F5.2,1X))
  700 CONTINUE
  300 CONTINUE
  200 CONTINUE
  100 CONTINUE
      END
```

Figure B-1(b)

```
      SUBROUTINE ASI(F,D,TH,A)
C-SUBROUTINE 'ASI' COMPUTES ANTENNA GAIN RELATIVE TO MAIN BEAM GAIN,AT ANGLE,
C-TH(DEGREES) OFF MAJOR AXIS OF ANTENNA OF DIAMETER 'D'(METERS)
C- OPERATING AT FREQUENCY'F'(GHZ). DEPENDING ON TH, FINDS RELATIVE GAIN
C- IN 4 REGIONS:
C-       1)TH<10.5/(F*D), IN MAIN LOBE
C        2)10.5/(F*D)<TH<30/(F*D),INTERPOLATE
C        3) TH>30/(F*D) USES 32-25*LOG(TH) RELATIONSHIP
C        4) RELATIVE GAIN NEVER LESS THAN 10 DB BELOW ISOTROPIC
      G=F*D
      GAIN=60.31*G**2
      G10=1./(10.*GAIN)
      Z=30./G
      IF(TH .LT.  Z ) GO TO 1
      X3=TH**2.5
      IF(G .GT.  30.  ) A=26.2/(X3*G**2)
      IF(G .LE.  30.  ) A=787.6/(X3*G**3)
      GO TO 10
    1 Q= 10.5/G
      IF(TH .LE.  Q) GO TO 2
      IF(G .GT.  30. ) X1=0.005315*G**0.5
      IF(G .LE.  30. ) X1=0.159773/G**0.5
      GG=0.5/X1
      B=(Z-TH)/(Z-Q)
      A=X1*GG**B
      GO TO 10
    2 A=1.-0.5*(TH/Q)
   10 IF(A .LT.  G10) A=G10
      RETURN
      END
```

Figure B-2

```
      SUBROUTINE BPSHZ(CN, IPE, R, MM , PZ      )
C-SUBROUTINE 'BPSHZ' COMPUTES BITS/SECPER HERTZ OF BANDWIBTH GIVEN INPUT 'C/N'.
C C/N IS INTERPRETED AS RATIO OF CARRIER POWER TO PRODUCT OF NOISE DENSITY TIMES
C SYMBOL RATE.  COHERENT PSK IS ASSUMED WITH M PHASES WHERE M=2, 4, 8,  OR 16.
C-TABLES ARE STORED OF: EB/NO AS A FUNCTION OF M AND P(E); CODE RATE, R; AND
C-IMPLEMENTATION AND IMPAIRMENTS LOSS AS A FUNCTION OF M.  SUBROUTINE FINDS
C-MAXIMUM VALUE OF PZ(BITS/SEC/HZ) GIVEN C/N AND P(E)=10-**IPE.  ALSO RETURNS
C- 'M' AND CODE RATE 'R'.
C         EBNO(M, IPE)-ENERGY PER BIT TO NOISE DENSITY RATIO
C         FOR MODULATION TYPE "M" AT BIT ERROR RATE 10**-IPE
C         EBNO IS NOT IN DB
C         RCOD(K)- RATE OF CODE"K"
C         CODES ARE STORED IN ORDER OF DESCENDING RATE; I. E. , RCOD(1), THROUGH
C         RCOD(8) ARE: 8/9, 7/8, 4/5, 3/4, 2/3, 1/2, 1/3, AND1/7,  RESPECTIVELY
C         GCOD(K, IPE)- CODING GAIN OF CODE "K" AT DECODED PB(E)=10**-IPE
      DIMENSION EBNO(4, 8), GCOD(8, 8), RCOD(8)
      DIMENSION XMI(4)
      DATA XMI/1. 58, 1. 78, 2. 00, 2. 24  /
      DATA RCOD/0. 889, 0. 875, 0. 8, 0. 75, 0. 667, 0. 5, 0. 333, 0. 143 /
      DATA EBNO/0. 813, 0. 813, 2. 510, 7. 079,
     2          2. 692, 2. 692, 6. 310, 18. 62,
     3          4. 786, 4. 786, 10. 47, 31. 62,
     4          6. 918, 6. 918, 15. 14, 43. 65,
     5          9. 120, 9. 120, 20. 00, 57. 54,
     6          11. 22, 11. 22, 24. 55, 69. 18,
     7          13. 49, 13. 49, 29. 51, 85. 11,
     8          15. 85, 15. 85, 35. 48, 95. 50 /
      DATA GCOD/0. 50, 0. 50, 0. 50, 1. 20, 0. 50, 0. 50, 0. 50, 0. 50,
     2          0. 50, 0. 50, 0. 50, 1. 40, 0. 50, 1. 70, 0. 50, 0. 50,
     3          1. 00, 1. 10, 1. 30, 1. 91, 2. 15, 2. 40, 2. 80, 3. 20,
     4          1. 45, 1. 58, 1. 86, 2. 29, 2. 60, 2. 88, 3. 31, 3. 80,
     5          1. 60, 1. 74, 2. 00, 2. 57, 2. 90, 3. 24, 3. 40, 3. 90,
     6          1. 70, 1. 80, 2. 40, 2. 75, 3. 00, 3. 47, 3. 50, 3. 95,
     7          1. 80, 2. 20, 2. 70, 2. 88, 3. 30, 3. 63, 3. 70, 4. 00,
     8          2. 00, 2. 50, 3. 00, 3. 20, 3. 40, 3. 80, 3. 90, 4. 10 /
      BMARG=1. 26
      R=1.
      MM=1
   10 MM=MM+1
      CN1=MM*R*EBNO(MM, IPE)*XMI(MM)
      IF(CN1 . GT.  CN) GO TO 20
      IF(MM . LT.  4) GO TO 10
      PZ=MM*R/BMARG
      GO TO 30
   20 CONTINUE
      BIG=(MM/BMARG)*(CN/CN1)
      DO 40 I=1, 8
      CN2=MM*RCOD(I)*EBNO(MM, IPE)*XMI(MM)/GCOD(I, IPE)
      IF(CN2 . GT.  CN) GO TO 50
      BBB=MM*RCOD(I)/BMARG
```

Figure B-3(a)

```
      IF(BBB .LT. BIG) GO TO 40
      BIG=BBB
      R=RCOD(I)
      GO TO 40
50    BIG=(MM*RCOD(I)/BMARG)*(CN/CN2)
      R=RCOD(I)
40    CONTINUE
      IF(MM .EQ. 2) GO TO 60
      CCC=(MM-1)/BMARG
      IF(CCC .LT. BIG) GO TO 60
      BIG = CCC
      MM=MM-1
      R=1
60    CONTINUE
      PZ=BIG
30    CONTINUE
      RETURN
      END
```

Figure B-3(b)

```
      SUBROUTINE CES3(KBND, D, PES, RB, TES, A1)
C SUBROUTINE COMPUTES THE COST OF AN EARTH STATION GIVEN THE
C FOLLOWING INPUT PARAMETERS:
C PALMER 2-26-79-- COST INCLUDES ANT+HPA+LNA+MODEM
C     KBND-DENOTES THE FREQ. BAND WHICH IS CODED AS:
C          1=6/4 GHZ
C          2=8/7 GHZ
C          3=14/11 GHZ
C          4=30/20 GHZ
C     D- ANTENNA DIAMETER
C     PES- POWER DENSITY(WATT/MBPS) RADIATED FROM THE EARTH STATION
C     RB- BIT RATE FROM/TO THE EARTH STATION, IN MBPS
C     TES- LNA NOISE TEMPERATURE
C     A1- COST OF ANT+HPA+LNA+MODEM IN DOLLARS
C COMPUTE THE COST OF THE ANTENNA
      DIMENSION A1(3)
      RB=RB*4.
      DO 100 K=1, 2
      IF(K. EQ. 1) GO TO 7
      RB=RB/4.
    7 GO TO (10, 11, 12, 13), KBND
   10 IF(D . GE.  10. ) CA=0. 222E+6*(D/10. )**2
      IF(D . LT.  10. ) CA=0. 150E+6*(D/10. )**3 +3. E+3
      GO TO 14
   11 CA=0.
      GO TO 14
   12 IF(D . GE.  5. ) CA=0. 0625E+6*(D/5. )**2
      IF(D . LT.  5. ) CA=0. 045E+6*(D/5. )**3 +3. 5E+3
      GO TO 14
   13 CA=0. 0
   14 CONTINUE
C COMPUTE THE COST OF THE H. P. A.
      P=PES*RB
      GO TO(20, 21, 22, 23),  KBND
C 6 GHZ
   20 IF(P . GE.  3000. ) CP=2. E+5*(P/3000. )**0. 5
      IF(P . LT.  3000. ) CP=1. 2E+5*(P/3000. )**0. 5 +2. E+3
      GO TO 24
C 8 GHZ
   21 IF(P . GE.  3000. ) CP=2. 2E+5*(P/3000. )**0. 5
      IF(P . LT.  3000. ) CP=1. 22E+5*(P/3000. )**0. 5+2. 2E+3
      GO TO 24
C 14 GHZ
   22 IF(P . GE.  2000. ) CP=2. 0E+5*(P/2000. )**0. 5
      IF(P . LT.  2000. ) CP=1. 3E+5*(P/2000. )**0. 5 +3. E+3
      GO TO 24
C 30 GHZ
   23 IF(P . GE.  1000. ) CP=1. 8E+5*(P/1000. )**0. 5
      IF( P . LT.  1000. ) CP=1. 2E+5*(P/1000. )**0. 5+3. 5E+3
   24 CONTINUE
```

Figure B-4(a)

```
C COMPUTE THE COST OF THE LNA
      GO TO(30,31,32,33), KBND
   30 IF(TES .GE. 120.) CT=30.E+3*(120./TES)**1.9
      IF(TES .LT. 120.) CT=30.E+3*(120./TES)**0.67
      GO TO 34
   31 CT=0.0
      GO TO 34
   32 IF(TES .GE. 500.) CT=10.E+3*(500./TES)**4.75
      IF(TES .LT. 500.) CT=10.E+3*(500./TES)**1.62
      GO TO 34
   33 CT=0.0
   34 CONTINUE
C COST C1 IS COST OF ANT+LNA+HPA
      C1=CA+CP+CT
C COST C2 IS COST OF QPSK MODEM
      C2=2.E+4*(RB)**0.2
      A1(K)=C1+C2
  100 CONTINUE
      A1(3)=A1(2)
      RETURN
      END
```

Figure B-4(b)

```
      SUBROUTINE CAPSAT(NBEAM,RBIT,K2,RATE,PZ,SATCAP,PACKDY)
C THE TOTAL NATWORK TRAFFIC IS 'PHAT' PACKETS PER YEAR.  A PACKET
C BY DEFINITION HAS 'PACK' INFORMATION BITS.  THE NETWORK HAS 'NUSER' USERS
      COMMON/COM1/ FU(3),PU(3),UTIL(3),ACT(3),QBEAM,ZES
      COMMON/COM2/ TCKTS(3),BCKTS(3),ZUES(3),RBPKUU(3),DSI(3),PHAT
      DIMENSION SATCAP(3)
      DIMENSION ZUST(3)
      DIMENSION PAKCON(2)
      DATA PAKCON/0.1,0.5/
      ZBEAM=NBEAM
      NES=ZES
      BTRAN=40.E+6
      PHAT=8.E+12
      NUSER=1000
      ZUSER=NUSER
      PACK=500.
      TFRAME=1.
      YS=365.*24.*3600.
      ZBURST=ZBEAM
      NPBK=2
      NFREZ=6
      PRE=100.
C USERS ARE IDENTIFIED BY 3 CATEGORIES: IU=1,HIGH RATE; =2, MEDIUM RATE;
C =3,LOW RATE.  OF THE TOTAL POPULATION OF USERS, A FRACTION FU(IU)
C ARE IN EACH CATEGORY.  THE FRACTION OF THE TOTAL TRAFFIC BY EACH
C USER CATEGORY IS GIVEN BY PU(IU).  THE AVERAGE UTILIZATION FOR EACH
C USER CLASS IS GIVEN BY UTIL(IU) HOURS/DAY.  THE ACTIVITY FACTOR
C BY USER CLASS IS ACT(IU).
C*****************************************************************************
C POSTULATE A NETWORK THAT IS DIVIDED INTO REGIONS.  EACH REGION
C HAS ONE OR MORE EARTH STATIONS AND AN  EQUAL SHARE OF THE USERS IN
C  EACH CLASS.
C DETERMINE THE NUMBER OF USERS BY CLASS PER EARTH STATION,NUES(IU),
C AND THE PEAK INFORMATION BIT RATE PER USER-TO-USER CONNECTION PER
C EARTH STATION,RBPKUU(IU)
C*****************************************************************************
      DO 1 K=1,3
      DSI(K)=1.
      ZUST(K)=ZUSER*FU(K)
      ZUES(K)=ZUSER*FU(K)/ZES
      RBPKUU(K)=PHAT*PU(K)*PACK/(YS*ZUST(K)*(ZUST(K)-1 )*UTIL(K)*
     1 ACT(K)  )
      ROUT=RBPKUU(K)
    1 CONTINUE
```

Figure B-5(a)

```
C*******************************************************************************
C NOW COMPUTE THE SATELLITE CAPACITY NEEDED PER EARTH STATION FOR THREE
C TRANSMISSION TECHNIQUES:
C       ITRAN=1  PACKET SWITCHING
C       ITRAN=2  CIRCUIT SWITCHING
C       ITRAN =3 CIRCUIT SWITCHING WITH ACTIVITY CONCENTRATION
        DO 2 ITRAN=1,3
        GO TO (23,21,22), ITRAN
C COMPUTE THE SATELLITE CHANNEL CAPACITY REQUIRED FOR CIRCUIT SWITCHING
C COMPUTE THE CAPACITY PER EARTH STATION
C*******************************************************************************
   21 CAP=0.
      DO 3 IT=1,3
      ZCK1=ZUES(IT)*(ZUST(IT)/ZBEAM-1.)
      CALL BLOCK(ZCK1,UTIL(IT),NPBK,SATC1)
      ZCK2=ZUES(IT)*ZUST(IT)/ZBEAM
      CALL BLOCK(ZCK2,UTIL(IT),NPBK,SATC2)
      SATCH=SATC1+(ZBEAM-1.)*SATC2
      TCKTS(IT)=ZCK1+(ZBEAM-1.)*ZCK2
      BCKTS(IT)=SATCH
      CAP=CAP+SATCH*RBPKUU(IT)
    3 CONTINUE
      CAP=CAP/(PZ*BTRAN)*(1.+ZBEAM*PRE/(CAP*TFRAME))
C  NOW MULTIPLY BY THE NUMBER OF EARTH STATIONS
      SATCAP(ITRAN)=CAP*NES
      GO TO 2
C*******************************************************************************
C COMPUTE THE SATELLITE CAPACITY REQUIRED FOR CIRCUIT SWITCHING WITH
C ACTIVITY CONCENTRATION
   22 CAP=0.
      SMCONT=0.
      DO 4 IT=1,3
      ZCK1=ZUES(IT)*(ZUST(IT)/ZBEAM-1.)
      CALL BLOCK(ZCK1,UTIL(IT),NPBK,SATC1)
      CALL BLOCK(SATC1,ACT(IT),NFREZ,CONS1)
      ZCK2=ZUES(IT)*ZUST(IT)/ZBEAM
      CALL BLOCK(ZCK2,UTIL(IT),NPBK,SATC2)
      CALL BLOCK(SATC2,ACT(IT),NFREZ,CONS2)
      CONSC=CONS1+(ZBEAM-1.)*CONS2
      CAP=CAP+CONSC*RBPKUU(IT)
C ADD OVERHEAD FOR CONCENTRATION CONTROL
      CON1=SATC1
      CON2=SATC2
      IF(CONS1 .EQ. SATC1) CON1=0.
      IF(CONS2 .EQ. SATC2) CON2=0.
      SUM=CON1+CON2
      IF(SUM.EQ.0.) DSI(IT)=0.
      SMCONT=SMCONT+CON1+CON2*(ZBEAM-1.)
    4 CONTINUE
      CAP=(CAP+(ZBEAM*PRE+SMCONT)/TFRAME)/(PZ*BTRAN)
```

Figure B-5(b)

```
C  NOW MULTIPLY BY THE NUMBER OF EARTH STATIONS
      SATCAP(ITRAN)=CAP*NES
      GO TO 2
C*****************************************************************************
C COMPUTE THE SATELLITE CAPACITY REQUIRED FOR PACKET SWITCHING
   23 CONTINUE
C     HEADER-HEADER BITS /PACKET
C     FDELAY- OVERSIZING FACTOR TO AVOID DELAY
C     PAKCON(L)- FACTOR TO ACCOUNT FOR CONTROL; L=1,CONUS,=2,SPOT
      FDELAY=0.1
      HEADER=275.
      G1=(PACK+PRE+HEADER)/PACK
      G2=1.+FDELAY
      G3=1.+PAKCON(K2)
      FPAK=G1*G2*G3
      SATCAP(ITRAN)=PACK*PHAT*FPAK/(YS*PZ*BTRAN)
      PACKDY=0.27
    2 CONTINUE
      RETURN
      END
```

Figure B-5(c)

```
      SUBROUTINE BLOCK(XN,U,NPBK,SATCH)
C  SUBROUTINE PROVIDES APPROXIMATE NUMBER OF SATELLITE CHANNELS (SATCH)
C  REQUIRED TO PROVIDE A CERTAIN "BLOCKING" PROBABILITY GIVEN XN INPUT
C CIRCUITS WITH UTILIZATION "U".  NPBK IS CODED AS
C FOLLOWS;
C NPBK=1,0.05;  =2,0.01;  =3,0.005;  =4,1.E-3;  =5,5.E-4;  =6,1.E-5;=7,3.E-5
      DIMENSION XK(7)
      DATA XK/1.65,2.36,2.62,3.1,3.3,3.7,4.0  /
      SATCH=XN
      IF(XN.LE.10.) GO TO 5
      XMEAN=XN*U
      SIGMA=SQRT(XMEAN*(1.-U))
      SATCH=XMEAN+XK(NPBK)*SIGMA
    5 RETURN
      END
```

Figure B-6

```
      SUBROUTINE ESCOST(A1,ESC)
C     CALCULATES FULL E.S. COST GIVEN THE FOLLOWING PARAMETERS:
C       ITRAN- 1=PACKET  2=CIRCUIT  3=CONCENTRATED CIRCUIT
C       ZUES- DISTRIBUTION OF E.S. USERS(#H,#M,#L)
C       A1- COST OF ANT+HPA+LNA+MODEM
C       ESC(ITRAN)- TOTAL E.S. COST IN DOLLARS RETURNED
      COMMON/COM2/ TCKTS(3),BCKTS(3),ZUES(3),RBPKUU(3),DSI(3),PHAT
      DIMENSION CBBN(3),CSTATN(3),CSIMPN(3)
      DIMENSION ESC(3)
      DIMENSION A1(3)
      DATA CBBN/10.,20.,100./
      DATA CSTATN/100.,200.,1000./
      DATA CSIMPN/1000.,5000.,20000./
      NH=ZUES(1)
      NM=ZUES(2)
      NL=ZUES(3)
      XH=TCKTS(1)
      XM=TCKTS(2)
      XL=TCKTS(3)
      YH=BCKTS(1)
      YM=BCKTS(2)
      YL=BCKTS(3)
      CSTFIX=50000.
      IF(DSI(1).EQ.0.) YH=0.
      IF(DSI(2).EQ.0.) YM=0.
      IF(DSI(3).EQ.0.) YL=0.
      TOT=YH+YM+YL
      IF(TOT.EQ.0) CSTFIX=0
      CBBFIX=20000.
      CPSP=200000.
      CSIMPF=100000.
      CGTWN=500.
      CGTWF=20000.
      CTDMA=300000.
      DO 100 ITRAN=1,3
      GO TO(10,20,15),ITRAN
C     PACKET SWITCHING COSTS
   10 CONTINUE
      CGATE=CGTWF+(CGTWN*(NL+NM+NH))
      CSIMP=CSIMPF+(CSIMPN(1)*NL)+(CSIMPN(2)*NM)+(CSIMPN(3)*NH)
       ESC(ITRAN)=A1(ITRAN)+CGATE+CSIMP+CPSP
      GO TO 100
C     CIRCUIT SWITCHING COSTS
   15 CONTINUE
C     WITH STAT. MUX
      CBB=CBBFIX+(CBBN(1)*XL)+(CBBN(2)*XM)+(CBBN(3)*XH)
      CBC=CSTFIX+(CSTATN(1)*YL)+(CSTATN(2)*YM)+(CSTATN(3)*YH)
      ESC(ITRAN)=A1(ITRAN)+CBB+CBC+CTDMA
      GO TO 100
C     WITHOUT STAT. MUX
   20 CBB=CBBFIX+(CBBN(1)*XL)+(CBBN(2)*XM)+(CBBN(3)*XH)
   30 ESC(ITRAN)=A1(ITRAN)+CBB+CTDMA
  100 CONTINUE
      RETURN
      END
```

Figure B-7

```
      BLOCK DATA
      COMMON/COM1/ FU(3),PU(3),UTIL(3),ACT(3),QBEAM,ZES
C  SET EARTH STATION AND USER SCENARIO
C RUN B2
      DATA PU  /0.6,0.3,0.1  /
      DATA  FU  /0.04,0.16,0.8  /
      DATA  ACT / 0.25,0.10,0.05  /
      DATA UTIL /1.0,0.5,0.125  /
      DATA ZES/40./
      DATA QBEAM/40./
      END
/*
```

Figure B-8