NRL Report 8404

Analysis of Alternative Satellite Channel Management Systems

M. E. Melich, C. E. Landwehr, and P. J. Crepeau

Navy Command and Control Systems Lead Laboratory
Communications Sciences Division

October 16, 1980

NAVAL RESEARCH LABORATORY
Washington, D.C.

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

UNCLASSIFIED

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER NRL Report 8404	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) ANALYSIS OF ALTERNATIVE SATELLITE CHANNEL MANAGEMENT SYSTEMS		5. TYPE OF REPORT & PERIOD COVERED Final Report 6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Michael E. Melich Carl E. Landwehr Paul J. Crepeau		8. CONTRACT OR GRANT NUMBER(s)
9. PERFORMING ORGANIZATION NAME AND ADDRESS Department of the Navy Naval Research Laboratory Washington, DC 20375		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS NRL Problem 75-0217-0-0 65866N, X0741CC
11. CONTROLLING OFFICE NAME AND ADDRESS Chief of Naval Operations (OP-940) Washington, DC 20350		12. REPORT DATE October 16, 1980 13. NUMBER OF PAGES 61
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office) Commander Naval Electronic Systems Command (PME-108) Washington, DC 20362		15. SECURITY CLASS. (of this report) UNCLASSIFIED 15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

A

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

Approved for public release; distribution unlimited.

18. SUPPLEMENTARY NOTES

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

Communication network protocols	SSIXS
Satellite communications	ASWIXS
CPODA	UHF-DAMA
CUDIXS	Network simulation
TADIXS	Navy message demand

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

The Fleet Satellite (FLTSAT) Communications System provides U.S. Navy units with mobile communications services. The sharing of the satellite transmission capacity among several units is achieved using adaptive network management systems, which grant transmission intervals to users based on the amount of communications traffic each has. The FLTSAT space segment is scheduled to be replaced by the Leasat satellites, which will be leased by the Navy. The Leasat will have fewer transponders for Navy use than does FLTSAT, but the overall transmission capacity available to the

DD FORM 1473 1 JAN 73 EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102-014-6601

 SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

20. Abstract (Continued)

Navy is to be increased by more efficient use of the transponders. The ground segment terminal equipment is being modified, under the UHF-DAMA project, to achieve this greater efficiency.

This report presents the results of a Monte Carlo computer simulation analysis of the performance of different network management systems. The systems considered were the existing FLTSAT Information Exchange Systems (IXS's), the IXS's operating through semiautomatic UHF-DAMA, and a reservation protocol based on CPODA. The time delay in delivering messages of a given quality was the performance measure used. A model of message service demand was developed using historical operational data and extrapolations based on expected usage by new fleet systems.

CONTENTS

EXECUTIVE SUMMARY

BACKGROUND

The U.S. Navy has been using ultrahigh frequency (UHF) satellites for long-haul communications since the late 1960's. Since 1976, coverage of the Mediterranean Sea and the Pacific, Indian, and Atlantic Oceans has been provided by leased services using Gapfiller satellites. Gapfiller was originally intended as a limited capacity system providing interim service until the 4 FLTSATCOM satellite constellation was in orbit.

The first FLTSATCOM satellite was launched in 1978, covers the continental United States and portions of the eastern Pacific and western Atlantic, and has an expected life of 5 years. Congress has directed the Navy to buy just five FLTSATCOM satellites. In 1983, when the FLTSATCOM satellites are expected to begin failing, the Navy is to obtain replacements by leasing service, using the so-called Leasat satellites.

Gapfiller satellites have two 25 kHz and one 500 kHz radio frequency (rf) transponders. Today, the Navy uses one 25 kHz and portions of the 500 kHz transponder on each spacecraft for fleet teletype broadcasts, submarine and surface ship general purpose message services, secure voice networks, and various dedicated 100 word per minute (wpm) services. FLTSATCOM satellites provide ten 25 kHz channels, one of which is jam resistant for a protected fleet broadcast. Leasat is to have one jam resistant fleet broadcast, one 500 kHz, six 25 kHz, and five 5 kHz rf channels.

The rf resources planned for Navy use in each of the Leasat spacecraft will be three 25 kHz channels, the jam-resistant fleet broadcast channel, and two 2400 bit per second (bps) accesses in the 500 kHz channels. The Army and Air Force will make use of the remainder of the spacecraft's rf capabilities. Therefore, the demand for service from the Navy's users of nine of the existing FLTSATCOM channels must somehow be consolidated into the reduced number of rf channels assigned on Leasat.

The UHF demand assigned multiple access (DAMA) project has the task of designing and installing those ground segment terminal modifications that will allow more efficient use of the capacity of the leased spacecraft transponders. The increased number of circuits which will be provided will allow the existing surface ship and landbased users of the UHF satcom to continue to receive the service they currently receive. The dominant service is for message exchange and is provided *via* the various network channel management systems called the Information Exchange Systems (IXS), which are described in Appendix A. Also in UHF-DAMA, it is expected that additional telephone-like services will become available. Some users, such as submarines with SSIXS and aircraft with ASWIXS, will probably continue to use dedicated rf transponders because of the low-gain antennas which they use. That is, they will not be UHF DAMA subscribers.

PURPOSE OF THIS STUDY

The purpose of this study is to examine the consequences of, develop options for, and suggest action for dealing with a reduction in the number of UHF satellite channels available to the Navy during the 1980's.

PRINCIPAL RECOMMENDATIONS

OPNAV should:

- Establish a policy on the post-1987 use of UHF satellite communications to support Navy mobile platforms.

- Define a standard Navy communications demand model which can be used to evaluate the performance expected from proposed satellite communications system changes. (A partial example is provided in this study.)

- Define standard Navy telecommunications performance measures in terms of which naval communications services will be expressed, *e.g.*, blocking probability for telephone-like circuits or time delay for computer messages.

- Perform trade-off studies to show the merit in buying terminal modifications, such as UHF-DAMA, as opposed to suboptimal utilization of more rf transponders in Leasat.

- Cease investment in satellite network management protocol engineering development, *i.e.*, additional IXSs, automatic DAMA, until it can be shown either analytically or in a test program that such a diversity of network management schemes performs significantly better than existing protocols, such as CPODA (contention-based, priority oriented, demand assigned)-like ones.

- Establish a common test program to verify the comparative performance of semiautomatic UHF DAMA, a reservation protocol, such as CPODA, and a simple fixed time division multiplexing of existing polling protocols.

- Direct that the tests consider the following factors: responsiveness—overall network time delays; transparency—ability to support a mixture of digital voice, data, and messages; flexibility—levels of utilization achievable in the face of fluctuating demand.

- Encourage NAVELEX and DARPA to emphasize the communications protocol performance portion of the mobile access terminal (MAT) experiment.

- Use the results of such testing in the definition of Navy EHF Satcom Program (NESP) or any follow-on satcom system which the Navy may invest in after the expiration of the first Leasat contract.

- Limit the investment in automatic UHF DAMA unless it has been shown that over the life of UHF satcom in the Navy that buying additional rf channels is more expensive or that a new kind of "telephone-like" service is required.

SUMMARY OF THE STUDY, RESULTS, AND CONCLUSIONS

Utilization of a 25 kHz Channel at Higher Data Rates

The current 2400 bps operation of CUDIXS, SSIXS, and secure voice does not fully utilize the theoretically available transfer capacity of the three assigned and dedicated 25 kHz channels nor will the other IXSs. Shipboard UHF DAMA TechEval results have shown that higher signaling and consequently higher data rates can be achieved. Realizing these data rates requires changing the shipboard radio, its modulator/demodulator, and the addition of error protection encoding.

The amount of money to be invested for increased utilization of the 25 kHz channels should have been established in a trade-off between the cost of shipboard terminal equipment versus the cost of transponders over the expected life of the UHF satcom system. To perform such a trade-off, a policy which stated the expected life of Navy UHF satcom was required. The UHF DAMA project has proceeded to that stage of development where such a policy statement and its consequences are irrelevant for the current Leasat contract. For the post-1987 period when the current Leasat contract ends, there is a need for the Navy to enunciate its intentions to use or not use UHF satcom for its mobile unit communications.

Recommendation

Formulate a Navy policy on UHF satcom use post-1987. Based on this policy select an investment strategy for UHF terminal improvements and installation.

Allocation of Higher Data Rates

Assuming that some investment to increase the utilization of the 25 kHz channels is justified, the allocation of the, say, 9600 bps capacity must be defined. The present CUDIXS, for example, allocates a portion of the 2400 bps capacity to each ship according to the amount of traffic that ship has queued for transmission. (In principle, CUDIXS could equally well operate at 9600 bps.) Thus, CUDIXS divides the available transmission capacity so that a ship with many messages gets more capacity than does one with fewer. That is, it is a form of demand assigned multiple access network management for general service message traffic. But, unlike UHF DAMA, it has not assumed that each ship has a standard set of data rates which are multiples of 75 bps up to a maximum of 4800 bps. The UHF DAMA is designed with telephone-like users in mind; *i.e.*, the system is designed to reduce the blocked call probability (frequency of a "busy signal"). Message communication as opposed to voice communication has, to date, dominated long-haul Navy communications.

Recommendation

The apparent shift in emphasis from message and data to voice-like usage signified by the UHF DAMA approach to the division of increased capacity among users should be reflected in formal Navy requirements. If no such shift is intended, then the performance of other network management schemes should be formally investigated and a selection for post-1987 satellite systems made based upon lowest overall cost. For the near term UHF DAMA, as presently developed, is the only programmatically available way to increase the utilization of Leasat rf resources.

Performance of Various Network Management Protocols

The analysis in this study assumed the data rate available in a 25 kHz channel is increased and that the objective is to provide the widest set of users with the shortest possible delay times. The average delay time for three different priorities of messages of differing lengths was estimated in a 30-node network of CUDIXS subscribers and 19-node network of TACINTEL subscribers. The networks were assumed to be operating with a crisis load similar to that experienced by the SIXTHFLT during the 1973 Yom Kippur War augmented by a large number of short, computer-to-computer messages.

For the highest priority traffic the CPODA protocol provided equal or shorter time delays than the UHF DAMA or the existing IXS and in the case of the lower priority traffic did much better—delays as much as ten times shorter. Thus, this analysis demonstrates that there exists an alternative satellite channel management scheme which provides faster message delivery than either UHF DAMA or the existing IXS. However, the value of having faster service has not been determined in terms of the ability of any satcom user to do a job at all or do it better. For example, the timeliness requirement for

TACINTEL could not be justified from an examination of the overall response time needs of systems which exchange the BE-3 messages. In that particular case either a much faster exchange is needed or a slower one is tolerable.

Recommendation

Alternative network management subsystems, such as CPODA in the mobile access terminal experiment, compared to those currently being proposed for Leasat, such as DAMA with existing IXSs, have sufficiently attractive performance improvement to warrant full-scale testing. However, the time it would take to implement such alternative channel management schemes would preclude their being available in time for Leasat. Thus, the purpose served by such testing should be directed at any follow-on satcom system which the Navy may develop or use.

ANALYSIS OF ALTERNATIVE SATELLITE CHANNEL MANAGEMENT SYSTEMS

PURPOSE OF THIS STUDY

It is the purpose of this study to examine the consequences of, develop options for, and suggest action for dealing with a reduction in the number of ultrahigh frequency (UHF) satellite channels available to the Navy during the 1980's.

TECHNICAL APPROACHES AND ISSUES

General Consolidation Options

The consolidation of telecommunications implies that some aspect or some subsystem or subsystems of the telecommunications system will be shared. The issues that arise in selecting what consolidation scheme to use depend upon answers to the following three questions:

- What will be shared?
- Who will do the sharing?
- What affect will the sharing have on the users of the telecommunications system?

It is entirely possible for consolidation to occur outside of the boundary of the telecommunications system. For example, a consolidation of user groups could accomplish reduction in the apparent number of channels that the telecommunications systems would have to provide. A consolidation of the special intelligence (SI) and the general service (genser) communications services aboard a ship would result in a relatively natural merger of like services. It would involve having SI and genser messages handled by a common message center. Thus, when viewed by the telecommunications system only one transmission path would be used by the SI/genser message users. This type of merger has been proposed, in part, and is being implemented at CINCUSNAVEUR, London. Mergers of user groups usually result in the users assuming responsibility for tasks previously done by the telecomunications system, *e.g.*, the introduction of optical character readers caused all Navy message writers to become accurate formatters of typed messages.

Other consolidation schemes involve the sharing of one or several of the major subsystems of the telecommunications system. The three major subsystems of interest in this study were:

- User's Services or Injection Subsystem, which consists of those facilities necessary to provide narrative message, voice, data transmission, facsimile, video, or television transmission services to the user.

- Network Management Subsystem, which takes the traffic being generated by the various user services subsystems and attempts to allocate transmission capacity in an equitable manner in response to established priorities.

- Transmission Subsystem, which provides for the delivery of bits from point-to-point with a certain degree of accuracy. It generally contains the modulation and demodulation functions and the appropriate bit level error protection coders, the transmitters, and the receivers necessary to signal through a particular medium.

Manuscript submitted March 13, 1980.

The major consideration in selecting the kind of sharing to implement should be the type and quality of service that the users need or expect. In today's FLTSATCOM system, the dedication of a radio-frequency (rf) channel has created the impression among the users that such dedication is necessary for timely message delivery. This is particularly true for users of the TACINTEL system, who are very much concerned with timely delivery of very short messages and are not particularly concerned with the efficient utilization of the rf channel. The suggestion that the TACINTEL user community and the genser user community be merged would, based upon these communities' historical experience, be considered unacceptable, since the time delays for TACINTEL users' messages might well turn out too long. The demands of elaborate security classification systems are argued to be best met with separate communications centers. (Similarly, a consolidation which merged voice traffic and data or message traffic could result in degraded performance for each of the users if the traditional way of providing such services were employed.) This study has only cursorily considered consolidation alternatives that require the total merging of user communities.

The second class of consolidation would result in the sharing of the network management subsystem. For example, one of the easiest ways to accomplish a consolidation is to increase the data rate from 2400 bits per second (bps) to 4800 bps over the transmission subsystem and then to divide each second's worth of capacity evenly between the two 2400 bps users. Some such time division multiplexing method of sharing is the approach proposed in all Leasat era consolidation schemes. The major issue that arises is how responsive to changing demands the multiplexing or the sharing scheme will be. At issue here is:

- whether allocation of a fixed share of the total capacity to a particular network for periods of tens of minutes will be acceptable; or
- whether there should be an ability to adjust allocated capacity between networks on a more timely basis in response to a rapidly shifting demand, as occurs in interactive computer networks.

For example, aboard a particular ship which is equipped with TACINTEL and CUDIXS the allocation of 2400 bps to each of them could well result in the CUDIXS not fully using the capacity it was given while on the same ship the TACINTEL experiences backlogs. A consolidation which does not recognize the opportunity to shift or balance traffic loads from one user group's information exchange system (IXS) or network to another may result in reduced performance, generally reflected by an increase in the average delay time.

The timely, accurate delivery of his messages is the telecommunications system's performance measure of interest to the user. The efficiency with which the resources of the telecommunications system are exploited is of little interest to the user (so long as he is served, of course). In other words, a user prefers to have timely and accurate delivery and is not particularly concerned that this may involve the dedication of an rf communications channel which consequently remains idle much of the time. In the commercial world, of course, the cost of this idle resource is visible to the user since the user (management) sees his money going to pay the telephone bill. The operators, fiscal planners, and designers of the communications system, however, are concerned with the efficient use of communication system resources. As a result, there is a natural desire for the designer and operator to strive first for high utilization. This may conflict with the desire by each user to have the fastest, most accurate communications.

This study addresses the communications system from the standpoint of the user and attempts to answer the question:

"What will be the effect of various consolidation options on the average time delays that the user will experience?"

The consolidation option to be pursued should be selected on the basis of economic costs as well as the time delays and utilization. The major technical issue in this study is which of the various alternative ways of using an increased rf channel data rate is most effective for the user. *Implementation costs and utilization have not been a primary concern of this study.* Thus, this study is not a complete analysis of the decision problem which Navy planners face.

Consolidation Options Studied

The average time delays were analyzed for in the following three telecommunications system configurations.

- Current System. The current CUDIXS and TACINTEL message services operating in dedicated 2400 bps channels.

- Fixed Time Division Multiplex. The dedicated allocation of time slots in the proposed semiautomatic UHF DAMA (demand assigned multiple access) to the existing CUDIXS and TACINTEL networks which requires modification to the existing shipboard satcom terminal.

- Reservation Protocol. The sharing of a single rf channel for all CUDIXS and TACINTEL subscribers with allocation dynamically established by a contention-based, priority oriented, demand assignment (CPODA)-like protocol, which would require modification of the shipboard satcom terminal.

The implementation of either UHF DAMA or a reservation protocol involves hardware and software modifications of the shipboard equipment. The cost of UHF DAMA through fiscal year 1985 is estimated to be over $50 million excluding R&D and installation costs. The cost of a reservation protocol has not been estimated, but the major equipment changes have been developed for other projects and appear to be competitive with or cheaper than the comparable UHF DAMA costs. However, they have not gone through the operational testing procedures which lead to certification for service use that UHF DAMA is just now completing. A less elaborate fixed time division multiplexer than UHF DAMA should be significantly less expensive and significantly less capable of fully utilizing the satellite transponder, but no cost data have been developed for such an alternative approach.

GENERAL CONCLUSIONS AND OBSERVATIONS

Utilization of a 25 kHz Channel at Higher Data Rates

The current 2400 bps operation of CUDIXS, SSIXS and secure voice does not fully utilize the transfer capacity (Shannon Capacity) of the three assigned and dedicated 25 kHz channels nor will the other IXSs—AWIXS, TADIXS. Laboratory tests have demonstrated that 9600 bps data rates should be achievable in a noisy shipboard environment. Shipboard UHF DAMA TechEval results have shown that higher signaling and consequently higher data rates can be achieved. Realizing these data rates requires changing the shipboard radio, its modulator/demodulator, and the addition of error protection encoding. The degree of sophistication to be incorporated in these transmission subsystem changes should have been determined in a trade-off analysis of the lease cost of a 25 kHz channel versus the cost of the terminal equipment changes.

For example, if the cost of a leased channel for all satellites for four years is $12.5 million and the cost of the terminal modifications is $50 million, then it is cheaper to lease four underutilized channels and use existing terminal equipment than to make the modifications. Beyond the first leasing period the cost of the leased channels should drop and also the post-1987 Navy satcom may not employ the same class of transmission subsystem. There is no indication that this economic analysis was done for the Leasat including Navy terminal costs nor is there evidence that the post-1987 Navy satcom system has been defined to include or not include UHF.

Allocation of Higher Data Rates

Assuming that some investment to increase the utilization of the 25 kHz channels had been justified in trade-off analysis of terminal versus transponder costs, the operational allocation of a portion of the capacity to each ship terminal must be defined. The present CUDIXS, for example, allocates, over a time of a few tens of seconds, a portion of the 2400 bps transmission capacity to each ship according to the amount of traffic the ship has queued for transmission. (In principle, CUDIXS could equally well operate at 9600 bps.)

Thus, CUDIXS divides the available transmission capacity so that a ship with many messages gets more capacity than does one with fewer. It is a form of demand assigned multiple access network management for specialized message traffic. But unlike UHF DAMA it has not assumed that each ship has a standard set of data rates which are multiples of 75 bps up to a maximum of 4800 bps. And unlike UHF DAMA, the CUDIXS/NAVMACS does collect information on the queues of messages waiting to be transmitted and makes its allocation decisions according to a computer algorithm. Semiautomatic UHF DAMA works much more like an older style telephone exchange in that it does not attempt to find out how many "customers" are lined up outside the telephone booths trying to make a call. If a customer wants to use an unused circuit he must ask for a controller at the Naval Communication Station to connect a "new instrument" for him. This manual operation generally takes minutes depending on how busy the central controller is. If this customer were prescient enough to arrange to have the "new instrument" already connected, then the semiautomatic DAMA would automatically respond when he tries to dial whomever he wishes to call.

The speed with which a CUDIXS-like network management subsystem responds to changes in demand is a fraction of second to seconds. Semiautomatic UHF DAMA takes minutes to tens of minutes to adapt, because it depends on manual recognition of load imbalances. The choice of which way to use the increased capacity ultimately depends then on how constant the demand at each ship is. If a ship has highly erratic or fluctuating demand, then dedication of a fixed capacity transmission segment is going to be either wasteful or inadequate and average delay times will be longer than necessary. If, however, the demand is relatively constant such as in a telephone conversation, then dedicating a fixed segment is best. (Note that the notion of what is a fluctuation depends upon the scale used in making the measurement. Thus, if the demand statistics used to characterize a particular service are in tens of minutes, then the second-by-second fluctuations are masked.)

What is needed is a network management subsystem which can adapt with ease to either a constant or a rapidly varying demand. The class of reservation algorithms, such as the contention-based, priority oriented, demand assignment (CPODA), hold promise of providing just such network management capability. They are designed to collect loading data which are used to allocate communication capacity so that imbalanced demand is rapidly accommodated.

The planned concept of operation of semiautomatic UHF DAMA in the Leasat era envisions the assignment of fixed capacity segments to the current and planned users of the FLTSATCOM channels by a DAMA controller at an area communications station. The assumption is that fluctuations in the demand from a CUDIXS or TACINTEL subscriber group can be adequately accommodated by a manual control system. It will use demand statistics collected over a teletype-speed order wire from shipboard communications people and IXS network controllers to detect loading imbalances.

The design of UHF DAMA is more appropriate to a demand that is heavily dominated by voice-like communications—a demand which is much closer to that historically associated with the Army in its wire-based field communications. Message communication as opposed to voice communication has, to date, dominated long-haul Navy communications. The apparent shift in emphasis signified by the UHF DAMA approach to the division of increased capacity among users should be reflected in formal

Navy requirements. Under channel management schemes which are being tested by DARPA this shift would not be of much consequence, if they perform as expected. However, for the scheme being implemented in UHF DAMA there is an implied shift in Navy communications requirements.

Performance of Various Network Management Protocols

The analysis in this study assumed the data rate available in a 25 kHz channel is increased and that the objective is to provide the widest set of users with the shortest possible delay times. The average delay time for three different priorities of messages of differing length distributions was estimated in a 30-node network of CUDIXS subscribers and 19-node network of TACINTEL subscribers. The networks were assumed to be operating with a crisis load similar to that experienced by the SIXTHFLT during the 1973 Yom Kippur War [1] augmented by a large number of short, computer-to-computer messages.

For the highest priority traffic the CPODA protocol provided equal or shorter time delays than the other systems, and in the case of the lower priority traffic did much better—delays as much as ten times shorter. (See Tables 1 through 8.) Thus, this analysis demonstrates that there exists an alternative satellite channel management scheme which provides faster message delivery than either UHF DAMA or the existing IXS. However, the value of having faster service has not been determined in terms of either the ability of any satcom user to do his job at all or to do it better. For instance it has been suggested that the computer-to-computer exchange use of TACINTEL requires more timely response than will be provided by a 2400 bps circuit. Yet investigation of the consequences of the faster service for this particular use failed to show that a better job could be done even if messages got delivered ten times faster.

Table 1—Performance of TACINTEL and CUDIXS Operated as Separate Nets with Dedicated DAMA Slots on a Single FLTSAT Channel (zero error rate)

Polling		Mean Message Delay (s)		Utilization	
		CUDIXS	TACINTEL	CUDIXS	TACINTEL
4800 bps					
	Priority 1	8.28	4.36		
	Priority 2	20.1		0.28	0.26
	Priority 3	35.7			
2400 bps					
	Priority 1	12.8	5.35		
	Priority 2	34.3		0.47	0.45
	Priority 3	117.			
CPODA (Merged CUDIXS and TACINTEL Demand) (Averages from Table 3)					
9600 bps					
	Priority 1	1.50			
	Priority 2	2.88		0.26	
	Priority 3	3.18			
4800 bps	(See Table 2)				

Table 2—Baseline Environment with 4.8 kbps Information Transfer Rate (9.6 ksps signaling rate, half-rate code)

Seed Set	Mean Message Delay (s)		Utilization	
	CPODA	Polling	CPODA	Polling
1				
Priority 1	6.12/6.74[a]	7.11		
Priority 2	7.44/9.75[a]	20.4	0.47/0.51[a]	0.36
Priority 3	7.12/16.4[a]	36.5		
2				
Priority 1	6.83	7.48		
Priority 2	12.7	23.3	0.53	0.38
Priority 3	20.7	179.		
3				
Priority 1	8.10	7.20		
Priority 2	8.33	25.0	0.50	0.37
Priority 3	16.0	123.		

[a]Measured after 2 simulated hours

Table 3—Baseline Environment Studies
(9.6 kbps information transfer rate)

Seed Set	Mean Message Delay (s)		Utilization	
	CPODA	Polling	CPODA	Polling
1				
Priority 1	1.48	6.08		
Priority 2	2.71	15.8	0.25	0.20
Priority 3	3.04	23.1		
2				
Priority 1	1.51	6.08		
Priority 2	3.03	17.0	0.27	0.21
Priority 3	3.39	92.6		
3				
Priority 1	1.52	6.09		
Priority 2	2.90	18.4	0.26	0.20
Priority 3	3.12	75.3		
Means for Seed Sets 1-3				
Priority 1	1.50	6.08		
Priority 2	2.88	17.1	0.26	0.20
Priority 3	3.18	63.7		

Table 4—Baseline Environment with Doubled CUDIXS Load

Priority	Mean Message Delay (s)		Utilization	
	CPODA	Polling	CPODA	Polling
1	1.75	6.35		
2	3.06	66.2	0.34	25
3	3.83	499.		

Table 5—Baseline Environment with Doubled TACINTEL Load

Seed Set	Mean Message Delay (s)		Utilization	
	CPODA	Polling	CPODA	Polling
1				
Priority 1	2.02	5.98		
Priority 2	3.80	15.5	.38	.23
Priority 3	3.40	24.8		
2				
Priority 1	1.98	6.23		
Priority 2	3.53	16.8	.40	.25
Priority 3	3.91	89.4		

Table 6—Parameters for Error Environment

Probability of modem synchronization	=	0.999
Probability of crypto synchronization	=	0.999
Probability of error in header	=	0.0001
Probability of error in data block	=	0.01
Mean time to loss of reservation synchronization (CPODA only)	=	600s
Mean time to regain reservation synchronization (CPODA only)	=	2.0 s

Table 7—Baseline Environment with Errors Introduced

Seed Set	Mean Message Delay (s)		Utilization	
	CPODA	Polling	CPODA	Polling
1				
Priority 1	1.79	6.43		
Priority 2	3.77	15.8	0.26	0.20
Priority 3	3.87	25.4		
2				
Priority 1	1.86	6.61		
Priority 2	4.53	17.3	0.27	0.21
Priority 3	4.94	91.8		

Table 8—Performance of TACINTEL and CUDIXS Operated as Separate Nets with Dedicated DAMA Slots on a Single FLTSAT Channel (Zero Error Rate)

Polling	Mean Message Delay (s)		Utilization	
	CUDIXS	TACINTEL	CUDIXS	TACINTEL
4800 bps				
Priority 1	8.28	4.36		
Priority 2	20.1		0.28	0.26
Priority 3	35.7			
2400 bps				
Priority 1	12.8	5.35		
Priority 2	34.3		0.47	0.45
Priority 3	117.			
CPODA (merged CUDIXS and TACINTEL Demand) (Averages from Table 3)				
9600 bps				
Priority 1	1.50			
Priority 2	2.88		0.26	
Priority 3	3.18			
4800 bps (see Table 2)				

The results of simulation studies upon which these estimates are made should be considered indicative of relative performance, rather than predictions of what will be realized when implemented. Such results are significant when the differences between system time delays are as large as several hundred percent. A few percentage points difference in such analyses should not be considered significant enough to warrant taking action. The results of this analysis are significant enough to warrant further testing of the configurations analyzed, if the increased speed of delivery and flexibility is judged worthwhile.

Operation of CUDIXS and TACINTEL with UHF DAMA

A set of simulation studies was performed to provide a rough indication of the performance that might be expected using the UHF DAMA modem and the IXSs. A likely way of operating CUDIXS and TACINTEL with DAMA would be to dedicate different DAMA time slots to each of the IXS networks. Each network would then operate approximately as it does now, with a Network Control Station (NCS) and a polling scheme.

Assuming that the DAMA equipment on ships operates at a signaling rate of 19.2 ksps, the maximum information rate for a CUDIXS and a TACINTEL net operating on the same FLTSAT channel would be 4.8 kbps for each net, if a half-rate code is used.

To assess the performance of the independent TACINTEL and CUDIXS nets in this mode, the CUDIXS (including urgent) and TACINTEL loads were separated, so that one net includes 30 CUDIXS stations and the other includes the 19 TACINTEL users. Each of these was tested over a dedicated 9.6 ksps and 4.8 ksps channel, using the polling protocol only. (Note that this protocol is not an exact model for either CUDIXS or TACINTEL protocols, but is quite similar to them.)

The tests with 4.8-kbps channels, then, represent the best performance that can be expected of the DAMA system with the current IXS polling techniques. Table 1 displays the delays observed and contrasts them with the average delays in the baseline CPODA studies (which represent CPODA performance under comparable conditions). As the table shows, the delays are substantially higher in the DAMA system with IXS's applique.

The DAMA system with each net operating at 2.4 kbps may be compared with CPODA in tests at 4.8 kbps (Table 2). The TACINTEL delays under DAMA are marginally lower than under CPODA in this case, but all other traffic again receives substanially better service under CPODA.

MODEL OF THE USER MESSAGE DEMAND

Rationale for Selected Demand Model

The characteristics of the user demand on which the analysis of this study is based were extrapolated from the 1973 Yom Kippur War crisis experience of SIXTHFLT [1]. Since that time various capabilities have been introduced, such as Outboard which provides a long-range detection and identification capability for certain surface combatants. Their communications impact has been estimated and included in the demand characterization.

The SIXTHFLT during the Yom Kippur crisis had 45 to 50 ships. The network simulated in this report uses 29 ships and one shore station, but 19 of these ships are assumed to have two kinds of non-voice, high-speed duplex communications systems aboard—TACINTEL and CUDIXS. The 6700 messages per day volume of traffic for the 30-node network simulated is approximately the same as the highest traffic day during the crisis. Over and above this number approximately 55 000 TACINTEL 80-character messages per day were simulated in the network.

In view of the Navy's current shipbuilding program and the way the units are deployed, the mid-1980's should not see a great change in the message service demand. The largest change in demand may occur with the fleet-wide introduction of narrowband secure voice equipment. The actual policy for the operational allocation of the satellite transmission capacity, when secure voice is widely available, is unknown at this time. The Navy's use of an Air Force radio relay aircraft stationed over the Gulf of Tonkin during the Vietnam War would suggest that tactical employment of the secure voice service will be extensive and complex.

Parameters for the Demand Model

The basic demand environment to be modeled includes 30 nodes, distributed as follows:

1 major shore station, TACINTEL and CUDIXS
5 major ships, TACINTEL and CUDIXS

10 medium ships, CUDIXS
14 small ships, TACINTEL and CUDIXS.

A network of this size is large enough to be representative, but small enough to remain within the computational time and space constraints of the simulator when run on a PDP-10 computer. The traffic generated by each platform consists of CUDIXS traffic of two types, and (for the 19 major ships and smaller ships) TACINTEL traffic. The more routine SI traffic is assumed to be part of the traffic handled by CUDIXS. This makes the delivery of short, sensor-related messages the only load on TACINTEL and thus provides a good estimate of how responsive TACINTEL might be.

The amount of CUDIXS traffic generated by a platform depends on the platform type, as indicated below. For convenience, we assume that all messages from ships travel to shore, and messages from shore are uniformly distributed among all ships. Notice that this assumption has little impact on the utilization of the satellite channel, since the broadcast nature of the channel allows all addresses of a message to receive it from a single broadcast. This statement does not apply to the FLTSATCOM fleet broadcast, which uses 15 individual teletype channels in much the same way as the high-frequency multichannel fleet broadcast.

CUDIXS Traffic Generation

CUDIXS messages have been separated into two general categories to distinguish the highest priority messages (Flash and Flash Override) from the large volume of Immediate, Priority, and Routine traffic. These highest priority messages generally account for a small percentage of the traffic, but they have the most stringent delay requirements. Messages of this type are assumed to range from half a page (600 characters, including headers) to a page (1200 characters) in length, and to arrive randomly, at a rate of 6 per hour per the entire network. These messages are assumed equally likely to originate from any platform. Thus, the following parameters are used to model the generation of CUDIXS "urgent" messages at each of the 30 nodes:

message priority:	priority 1 (highest)
message length:	uniformly distributed from 345 characters (short header plus 150 characters text) to 1200 characters (header plus text of one page, approximately 200 words)
message inter-arrival times:	negative exponential distribution with mean of 18 000 s (=.2/h)

Non-urgent CUDIXS traffic is lumped into two priority categories for the purposes of this simulation: Immediate and Priority traffic are both considered to be priority 2, and routine traffic is priority 3. The volume of non-urgent CUDIXS traffic is assumed to be a function of the platform type, but the length and priority distributions of the generated traffic are assumed common. The following parameters are used to model the generation of non-urgent CUDIXS traffic:

message priority:	priority 2 (80% of messages generated) priority 3 (20% of messages generated)
message length:	from 345 characters to 2400 characters (2 pages), according to the following tabular distribution (see Table 9).

Table 9—Cumulative Distribution Function

No. Of Characters, C	Probability of Message Length Less Than or Equal to C
345	0.0
600	0.3
1200	0.6
2400	1.0

message inter-arrival times: negative exponential distribution with the following means according to platform types:
Major shore station: 18 s (=200/h)
Major ship: 600 s (=6/h)
Medium ship: 1800 s (=2/h).
Small ship: 1800 s (=2/h)

TACINTEL Traffic Generation

Traffic patterns in the TACINTEL system are substantially different from those in the CUDIXS system. Traffic flow security is a requirement for TACINTEL: operationally, this constraint leads to the operation of the TACINTEL polling scheme in a mode that closely resembles fixed time division multiple access (TDMA). Once a station is granted a slot allocation within the TACINTEL cycle, it must continue to transmit enough bits to fill that slot in each cycle regardless of whether it has useful traffic to be transmitted or not.

Within the present TACINTEL protocol, if all possible stations in a net are operating and each has a minimum length message to send in each frame, the network can accept approximately one 80-character message per 30 second cycle from each user. Consequently, we will model the TACINTEL traffic demand on the basis of this demand structure. It is quite possible that traffic flow security constraints might be met in some more efficient way in a consolidated network, but if the consolidated net can handle the maximum traffic in the existing system, it is clearly sufficient. The following parameters, then, govern the generation of TACINTEL messages on TACINTEL-equipped ships:

message priority: 1 (all TACINTEL traffic is considered urgent)
message length: 480 bits (80 characters, fixed length for all messages)

message inter-arrival times: fixed at one message per 30 s

ANALYTICAL APPROACH AND SIMULATION DESCRIPTIONS

In any simulation, some model of the system to be simulated must be defined. This model must be at a more abstract level than the system under study (otherwise we have in fact constructed the system that was to be simulated). The following paragraphs describe how the models for CPODA and the polling protocol used in the simulations differ from actual implementations of the protocols. A general polling protocol is used to model the operation of CUDIXS and TACINTEL.

CPODA Simulation Model

CPODA was chosen as representative of recently developed packet satellite protocols. The protocol was developed under DARPA sponsorship by BBN, Linkabit, and COMSAT and is described in

Appendix A. The Mobile Access Terminal (MAT) project plans to use this protocol to share a FLTSATCOM channel among several ships and a shore station.

CPODA Frame Structure

CPODA divides time into fixed length frames, as shown in Fig. A-8 of Appendix A. Frames are divided into subframes which are in turn divided into packets. Each frame includes an information subframe (possibly null), a leader timing subframe containing just one packet (frequently called the "leader packet"), and a contention subframe. The simulation uses the identical frame structure, except that the leader packet precedes the information subframe instead of following it. This is more convenient for the simulator, since, during the leader packet, the simulator can schedule the stations which are to next transmit information packets during that frame's information subframe. The scheduling is based on the reservations contained in the contention packets just received in the preceding frame. The unscheduled remainder of the current frame is then automatically devoted to contention. (Notice that, if the leader timing subframe is short relative to the frame length, the actual and simulated protocols are equivalent.) This change should affect delays by at most the length of a leader slot.

The algorithm used to allocate traffic to the information subframe is FIFO (First-In-First-Out) by priority; each frame is filled completely with information packets subject to the constraint that there be at least two contention packets and one leader packet per frame.

Reservations and Acknowledgments

Reservations are handled in the simulator essentially as described in Appendix A. Each contention packet may contain up to two reservations and up to four acknowledgments. In the proposed MAT system such packets may contain two reservations or several acknowledgments. The actual number of reservations and acknowledgments that can fit into a contention packet depends on precise data formats and the packet size; the discrepancy between the simulator and the system should not have any significant effect. Piggybacking of reservations and acknowledgments is managed in the simulator as it is in the CPODA design: the equivalent of one contention packet is appended to each information packet within the information subframe.

The protocol uses an ARQ (Automatic Repeat Request) mechanism; in the simulation, one acknowledgment is required for each data block transmitted. A data block corresponds roughly to 80 characters of text. An information packet may contain from one to eight data blocks. When a data block is received successfully, an acknowledgment for it is created and queued for transmission. If the sender of a data block does not receive an acknowledgment for the block within a specified time, it schedules the block for retransmission.

Contention

The contention behavior is essentially identical to that planned for the MAT system. Each time a message is generated at a node or a message arrives over the channel (requiring an acknowledgment), the node must decide whether or not to access the contention subframe to make the reservation (or send the acknowledgment). If the node already has reservations pending in the queue, it can piggyback the new reservation (or acknowledgment) on its next transmission in the information subframe. In the simulator, this choice is made on the basis of the number of acknowledgments and data blocks to be sent, the number of reservations pending for this node in the queue, and the expected waiting time until the first outstanding reservation for the node will be served.

When a contention packet is transmitted, the sending node sets up an "echo" timeout. If the echo of the packet is not received before the timeout occurs, the acknowledgments and the reservations from the packet are requeued and the contention threshold is halved, to reduce the probability of contention packet transmissions. If the echo is received successfully, the reservations are entered into the global reservation queue, the contention threshold is adjusted upward (the distance between the threshold and 1 is halved) to increase the probability of contention packet transmission, and the echo timeout is canceled.

Subscriber Start-up

It is in this part of the protocol that the simulated CPODA is most abstracted from CPODA implementation. To include the synchronization states in detail would require keeping track of the predictions made by each node for each information slot transmission. The complexity and computational overhead that would be incurred would overwhelm the present simulator and would not in any case seem to justify the minor increase in precision.

The approach adopted to model the achievement and loss of reservation sync is to consider the loss of sync to be a random event. The time until the next loss of sync is modeled as a random variable with a negative exponential distribution. The mean of this variable is a simulation parameter, the mean time until loss of reservation synchronization. The time to regain reservation sync is modeled similarly. While a station is out of sync, it does not transmit any messages or acknowledgments. All stations are assumed to start in reservation synchronization, and the initial acquisition state is omitted.

Polling Protocol Simulation Model

The polling protocol model is based on the TADIXS protocols as described in Appendix A, but extended to allow three priorities of traffic and up to 30 nodes. TADIXS was chosen as a basis for the polling protocol analysis because, of the several polling protocols designed for FLTSATCOM systems, it was designed to achieve the lowest delays and thus represents at least as good (possibly better) baseline performance as could be delivered by either the TACINTEL or CUDIXS.

Polling Protocol Cycle Structure

The polling cycle, as simulated, starts with a Network Control Station (NCS) transmission including a header, a Sequence Order List (SOL), and a data transmission. Subscribers then transmit in the order specified in the SOL. Each subscriber transmits a header and up to twelve 500-bit data blocks; message blocks are sent FIFO by priority.

The times for users to begin their transmissions are not sent as part of the SOL; each subscriber listens to the header transmitted by its predecessor in the SOL to determine how many blocks its predecessor will transmit. In this way, each subscriber can schedule its transmission to arrive at the satellite immediately following the termination of its predecessor's transmission. The goal of this strategy is to allow each subscriber to use only as much of the 12-block allotment as it needs and to allow the following user to avoid idle periods in the satellite.

There is, however, a penalty of a round trip delay time if, for example, a station has no traffic at all to send. That station's successor must wait until it hears the header indicating "no data blocks" before it can begin its own transmission.

The total cycle length is thus variable, with a minimum length of N (the number of subscribers) round trip delays. To restrict the maximum cycle length, so that high-priority traffic receives adequate services, the protocol also defines a threshold time. Any stations that begin transmitting after the current cycle length has exceeded this threshold may transmit only priority 1 traffic. To keep from blocking indefinitely the lower priority traffic from stations occurring at the end of the SOL, the order of transmission specified by the SOL is rotated at the beginning of each cycle.

Reservations and Acknowledgments

Since each subscriber has a slot in every cycle, reservations are unnecessary. Acknowledgments for all of the traffic in the last cycle are included in the headers transmitted by each subscriber.

Subscriber Start-up

The technique generally used in FLTSATCOM polling schemes to allow new subscribers to enter a net is contention-based. A random access time slot (RATS) is scheduled periodically in the SOL, and when a new subscriber wishes to enter the net, it must broadcast its identification in the RATS. The NCS then includes the new subscriber in the next SOL. If two new subscribers access the same RATS, some contention resolution scheme is required.

In practice, subscribers generally remain members of polling networks for relatively long periods, as compared to CPODA, but short compared to semiautomatic UHF DAMA, so that random access slots are not frequently employed. If a subscriber has no traffic to send in such a net and if it wishes to receive traffic, it must have an allocation in the polling cycle in order to transmit acknowledgments. Consequently, the polling protocol simulator does not attempt to model the subscriber start-up process. The SOL always includes all 30 users, and (since all of the users are active) no RATS slot is included in the net cycle.

Protocol Simulation Parameters

Both the CPODA and polling protocols include parameters that may be adjusted to attempt to optimize performance. The parameters chosen here for each protocol were selected to try to provide the best achievable combination of throughput and delay characteristics within the demand environment and for the hardware configuration of a consolidated Leasat system. Since the determination of "optimal" parameters for either protocol would be a study in itself, the choices here should be viewed as reasonable guesses based on the assumed channel characteristics and the demand model. Where the two protocols have similar parameters, comparable values have been used in both cases. Based on a feasible selection of hardware, the analysis assumes that both protocols are to be operated at 19.2 ksps signaling rate, half-rate encoded to allow a 9.6 kbps information transfer rate.

CPODA Parameters

The fundamental parameters for CPODA are the frame length, the length of a data block, the maximum number of data blocks per information packet, and the length of a contention packet. Table 10 lists initial assumptions and corresponding justifications for these and for several other protocol parameter values.

Polling Protocol Parameters

The most important parameters in the polling algorithm are those that determine the maximum length of user transmissions, since they determine the maximum cycle time. In the polling algorithm simulated, three priorities of data are recognized, and after the current polling cycle (frame) has exceeded a certain threshold length, only traffic of the highest priority may be transmitted by the remaining stations in the cycle. To compensate for this discrimination against stations appearing late in the polling list, the starting point of the polling list is rotated from cycle to cycle. This mechanism should be less crucial when the protocol is operated at 19.2 ksps instead of at 4.8 ksps for which it was designed. As in CPODA, error protection is provided by applying a half rate code to all information bits transmitted. This type of error protection coding could be considered part of the function of the transmission subsystem, not the network management subsystem. Parameters used in the simulations are documented in Table 11.

Table 10—Parameters for CPODA Simulations

Parameter and Value	Justification
frame length = 0.5s	Initial simulator test runs and discussions with BBN CPODA implementers. SATNET presently operates at 64 kbps (uncoded) with a 0.25 s frame.
data block = 1000 bits	Allows shortest messages to fit into one block without too much waste. Longest messages will require fewer than 15 data blocks; most will require fewer than 8.
maximum data blocks per packet = 8	Allows most messages to fit in one maximum length packet; longest messages require two maximum length packets. Prevents excessive number of acknowledgments.
contention packet length = 100 bits	Allows two reservations. BBN implementers report that a SATNET reservation requires about 48 bits. We assume a contention packet may contain up to 2 reservations and up to 4 acknowledgments. Note that these are information bits, not encoded bits.
leader packet length = 100 bits	Length of timing packet sent once per frame. Assumed same length as contention packet.
added overhead per information packet (beyond contention packet length) = 180 bits	An information packet includes space for a contention packet (to piggy back acknowledgments and reservations), some addressing and link control overhead for the data blocks on the packet, and the data blocks themselves. This figure is based on the link control and addressing overhead in the MAT design.
#guard bits separating user transmissions = 20 bits	MAT design value for 19.2 ksps.
modem synchronization preamble = 208 bits	MAT design value.
crypto synchronization preamble = 64 bits (after compression)	MAT design value. Note MAT takes KG-36 preamble of 841 bits and compresses it prior to transmission.
modem turnaround time = 0.0 s	Assume operation in full duplex (using Trident modifications) as in MAT.
bit per character = 6	Typical for current Navy system.

Table 11—Parameters for Polling Protocol Simulations

Parameter and Value	Justification
data block length = 500 bits	Corresponds roughly to short TACINTEL message.
maximum data blocks per user transmission = 12	Allows most messages to fit in one user transmission.
threshold time = 30 s	Reduces user transmission times after 30 s passes in frame. Value is large enough so that it was never reached in simulation experiments. Use of 5 s threshold showed no significant change in delays.
length of header (per slot) = 100 bits	Comparable to CPODA
sequence order list length = 100 bits	Comparable to previous polling algorithm designs. Net control station transmission includes both SOL and standard header.
overhead per data block = 100 bits	From CPODA estimate and previous polling algorithm designs.
No. of guard bits = 20	As in CPODA.
Modern synchronization = 208 bits	As in CPODA.
Crypto synchronization = 64 bits	As in CPODA.
Modern turnaround time = 0.0	As in CPODA.
No. of bits per character = 6	As in CPODA.

PERFORMANCE ANALYSIS—LIMITATIONS, SCOPE, AND RESULTS

The goal of these studies has been to demonstrate the feasibility of operating CUDIXS and TACINTEL on a shared FLTSAT channel and to assess the relative performance of CPODA and polling protocols in such an environment. More than 25 separate simulation experiments have been performed, most representing an hour of simulated activity on a FLTSAT channel. This section discusses the limitations of the simulation techniques used, the particular simulation experiments conducted, and the results obtained.

Limitations

The basic techniques used in Monte Carlo computer simulation impose some constraints on the results that can be obtained. Any simulator experiment conducted for a fixed length of time corresponds to a single realization of the simulated system's behavior (under the conditions assumed by the simulator) for that time period. Thus, to get accurate predictions for such measures as mean delay times in the real system, not only must the simulation assumptions correspond to real world conditions but also multiple replications of the same experiment (with different random number streams) are required.

Predictions that are relative in nature ("Mean delays are less in situation A than in situation B") can be made with greater confidence than predictions of absolute performance ("the mean delay in situation A will be 6.83 seconds"). To make such relative predictions it is necessary to test the two situations against a common environment. In the cases at hand, a common pattern of message arrivals can be assured in two separate runs by employing identical sequences of random numbers (by using identical "seeds"), for the message generator processes.

Transient behavior in a simulator can also cause difficulties. Particularly if the period simulated is relatively short or if the system is heavily loaded, transient effects in a given run may lead to atypical values for queue lengths and delays. To detect behavior in the simulator that is not representative of steady state conditions, measures for queue lengths and delays can be recorded periodically and plotted as function of simulated time so that trends may be observed. Some of the queue lengths and delay measures in these simulations were recorded in this fashion.

Further complications arise when saturation phenomena are present (*e.g.*, if the arrival rate in the system exceeds the rate at which the system can provide service). It can be difficult to distinguish between a saturated condition and a transient buildup of demand.

Simulation experiments are also subject to the limitations of memory space and computing time that apply to any computer program. In the present instance, memory became a critical resource when the simulator attempted to collect complete statistics on all 30 nodes in the scenario. Consequently, some measures of marginal interest were eliminated and many successful experiments were carried out with the revised simulator. The constraint on memory space still had an effect in a few cases, where (perhaps transient) backlogs developed in message queues and memory was exhausted. In such cases it is difficult to know whether the observed behavior is due to a transient peak in the simulated message traffic, to the saturation of the system simulated, to an instability in the protocol, or simply to a bug in the simulator. Fortunately, these problems were rare and did not appear to indicate any instabilities.

The difficulties cited above can also make investigation of failure modes of a protocol difficult and expensive. If an allocation technique has an instability, for example, it may or may not appear in a given, finite-length experiment. If it does appear, it may or may not be distinguishable from a transient.

As a consequence of these general limitations of Monte Carlo simulation, the results presented below are best interpreted as a feasibility demonstration, *not a guarantee of performance.* They demonstrate that, under the assumptions described for the demand model, the protocol model, and the physical environment, it is possible for the system to behave as the simulator indicates. In several cases, multiple replications with different seed sets yield somewhat greater confidence that observed values in the simulator are stable and represent a reasonable expectation for the actual system performance. More often, a single common seed set (and thus a single pattern of message arrivals) has been employed to assess the relative performance of two different modes of system operations.

Performance Measures

The fundamental measures reported below are mean message delay and system utilization. Message delay is measured from the time of file for a message at the source node until the last bit of the message is received at the destination node. Utilization is the fraction of time that the channel is busy, and thus includes overhead bits as well as information bits. The number of information bits transferred in comparable cases is generally identical or nearly so; differences in utilization indicate differences in channel overhead. The number of information bits transferred per unit time (excluding all overhead) is the channel throughput. Utilization should not be considered a measure of residual capacity in itself, since the behavior of most protocols limits achievable utilization to a value less than 1.0. Exactly what the achievable utilization for a given protocol is will depend on traffic patterns, noise conditions, *etc.* This study does not investigate the achievable utilizations.

Experimental Baseline Assumptions

A number of factors have been held fixed throughout most or all of the experiments. These will be described here and will be noted in the following sections only when they differ from the baseline.

- All runs are made for 1 h of simulated time. Preliminary experiments run for various time intervals indicated that after half an hour only minor changes were noted in the principal delay statistics. Longer studies would require more CPU time at a rate of 1 CPU minute (DEC KI-10 processor) per roughly 3 to 10 min of simulated time, depending on the traffic load.

- An error-free environment is assumed. Except in those runs investigating the behavior of the system under errors, it is assumed that all traffic is received error-free, that modems and cryptos always sync correctly, and that (in CPODA) reservation synchronization is never lost. Collisions, of course, can still occur on slots transmitted in the CPODA contention subframe. The assumption of an error-free environment is probably reasonable for the Leasat system when operating in good weather and in the absence of scintillation, and it provides a common base for comparison.

- Progation delays are assumed to be 0.28 s (round trip) for all stations. This assumption corresponds to operation with the satellite approximately on the horizon. Best case condition (satellite overhead) could lead to roundtrip delays of 0.25 s.

- The channel information transmission rate is 9.6 kbps achieved by using half rate coding with a signaling rate of 19.2 ksps.

- Parameters for protocol overheads, frame formats, and message arrival rates are fixed. The determination of these parameters was described above. Except where noted, the same random number seeds are used for the message arrival stream in all cases.

Baseline Experiments

The first set of simulation experiments investigates system performance for the baseline environment, as it has been defined above: a 30-node network of 29 ships and one shore station handling CUDIXS (including urgent) traffic (all nodes), and TACINTEL traffic (19 nodes). Because of the special interest in this case, the experiment was run three times with three separate seed sets. Table 3 records the results of each of the three tests and also the mean delays and utilizations averaged over the three cases. The delays in both the CPODA and polled protocol tests would be tolerable as absolute performance values in an operational system. The relative performance between the two protocol clearly favors CPODA; the average delay for each class of traffic under CPODA is less than one quarter of the corresponding case under polling.

Loading Experiments

To investigate behavior under heavier loads than those postulated for the baseline environment, several studies were conducted with increased CUDIXS or TACINTEL traffic loading. Table 4 reports the effect of doubling the arrival rates for all CUDIXS message generators, thereby increasing CUDIXS message traffic to about 14 000 messages per day. The results obtained for the baseline message stream show CPODA providing much better service than the polling protocol. In fact, the average message backlog (per minute) under polling was monotonically increasing throughout the duration of the 1 h test run, despite the relatively low utilization.

The TACINTEL loading was next doubled from the baseline so that each TACINTEL message generator produced a message once every 15 s instead of every 30 s, thereby increasing the number of 80-character messages to about 110 000 per day. (CUDIXS generators were operated as specified in the baseline.) Two different seed sets were used to investigate this case; again, CPODA performs with significantly lower delays than polling. Table 5 details these results. That priority 2 messages had a larger mean delay than priority 3 messages under CPODA with seed set 1, appears to be an artifact of the particular arrival stream.

An alternative to increasing the traffic in the system is to decrease the transmission rate of the channel. To investigate system behavior in this case, experiments were conducted assuming an information rate of 4.8 kbps, achieved by utilizing a half-rate code and a signaling rate of 9.6 kbps. To maintain the same number of bits per frame, the CPODA frame length was increased to 1 s from .5 s. Table 2 displays the results obtained.

For priority 1 traffic, the performance of the two protocols is roughly comparable, although CPODA still appears to provide better service to lower priority traffic than does polling. However, the CPODA experiments indicated that the system might be approaching saturation in this mode. The CPODA experiments for seed set 2 had a sizable backlog at the end of an hour (about 50 messages) so the delays reported for lower priority traffic in that case are probably underestimated. The experiment of seed set 1 with CPODA was also conducted for a period of 2 simulated hours to examine its stability, and the average priority 3 message delays measured over 2 h showed significant increases over those observed after 1 h.

Error Studies

Time and manpower limitations did not allow a complete model of the error environment to be constructed and studied with the simulator. Nevertheless, a simple model of how errors might occur in the system under normal operations was constructed and used to determine simulator parameters.

To successfully receive a transmission, a receiver must: (a) establish modem synchronization, (b) establish crypto synchronization, and (c) receive the header and data portions of the packet error-free (after Forward Error Correction is performed). If detected (but uncorrected) errors occur in a received packet, or if modem or crypto sync is not established, the packet will have to be retransmitted by the sender. In addition, CPODA requires that a node establish reservation synchronization before it can actually transmit data. Once established, reservation synchronization may be lost if a packet containing a reservation is missed or received incorrectly.

The parameters for the simulated error environment are shown in Table 6. The modem and crypto sync probabilities are based loosely on published data for the devices operating in an unstressed environment. Failure to obtain crypto or modem sync much more often than once in a thousand attempts in normal operation indicates a poor synchronization design.

Rates of detected errors in header and data blocks were determined based on the assumption of an error rate of 10^{-5} after the FEC has been applied. In the absence of RFI and scintillation, the channel will probably have a considerably lower error rate for bits delivered; this value is chosen as a reasonable upper bound for normal operations.

Data blocks are on the order of 1000 bits long, so, at a bit error rate of 10^{-5}, roughly one block in 100 is expected to contain a detected error. Headers are roughly 100 bits long, so one in a thousand is expected to contain a detected error. Additional error coding on the header block (as is customarily included) might lower this rate to one erroneous header in ten thousand.

Estimates for the mean time to loss of reservation synchronization are based on the error rate for header transmissions and the number of reservations transmitted per hour as observed in the simulator. Reservation sync will only be lost by a node if that node fails to receive a header containing a reservation. The probability of receiving an arbitrary header correctly is the product of the probabilities of successful modem sync, crypto sync, and error-free header receipt.

For the baseline environment, the number of reservations transmitted in an hour was roughly 2700. Assuming that each reservation represents a separate trial, this figure (together with the modem, crypto, and header error probabilities) indicates that each platform would lose reservation synchronization approximately 6 times per hour, or once every 10 min.

The time to regain reservation synchronization is complex to estimate. As a minimum, the out-of-sync station would need to hear correctly one reservation and the corresponding transmission. (Or, the out-of-sync node could observe a leader packet followed by any empty information subframe.) The reservation queue lengths observed in the zero error case for the situation at hand were generally short (average length less than one); consequently, an estimate of 2 s (4 half-second frames) was chosen as a reasonable value for the mean time to regain scheduling sync in both cases.

Experiments with the baseline demand plus errors were conducted for two seed sets; the results are shown in Table 7. The relative performance of CPODA and polling still favors CPODA substantially.

Comparison of Tables 3 and 7 shows that CPODA delays increase by a somewhat larger percentage when errors are introduced than do the delays under polling. This discrepancy may be due in part to the absence of Negative Acknowledgment (NAK) in CPODA. If a CPODA subscriber realizes it has received a packet with an error, it merely throws it away and waits for the sender to retransmit it. The sender does not retransmit the packet until a timeout period elapses, however, and in the baseline environment this timeout was set to about 30 s. Decreasing this period could lead to smaller increases in delay with increased numbers of errors.

Operation Under DAMA

A final set of studies was performed to provide a rough comparison with an environment using the UHF DAMA modem. A likely way of operating CUDIXS and TACINTEL under DAMA would be to dedicate different DAMA time slots to each of the networks. Each network would then operate approximately as it does now, with an NCS and a polling scheme.

To assess the performance of the independent TACINTEL and CUDIXS nets in this mode, the CUDIXS (including urgent) and TACINTEL loads were separated, so that one net includes 30 CUDIXS stations and the other includes the 19 TACINTEL users. Each of these was tested over a dedicated 9.6 ksps (= 4.8 kbps) and 4.8 ksps (=2.4 kbps) channel, using the polling protocol only. (Note that this protocol is not an exact model for either CUDIXS or TACINTEL protocols, but is quite similar to them.)

The tests with 4.8 kbps channels, then, represent the best performance that can be expected of the DAMA system with the current polling techniques. Table 3 displays the delays observed and contrasts them with the average delays in the corresponding baseline CPODA studies. Use of two DAMA channels, each operating at 4.8 kbps, corresponds to CPODA operation on a single 9.6 kbps channel

(Table 3), and use of two DAMA channels, each operating at 2.4 kbps, corresponds to CPODA operation on a single 4.8 kbps channel (Table 2). As Table 3 shows, the delays are substantially higher in the DAMA system with IXS applique. In the 2.4 kbps case, delays for TACINTEL under DAMA are marginally lower than under CPODA, but all other traffic receive substantially better service under CPODA.

REFERENCE

1. T. Oberlin, Genser Message Information in Sixth Fleet During the Yom Kippur War, CRC 334, Center for Naval Analyses, 1976.

Appendix A
NETWORK MANAGEMENT PROTOCOLS

Described in this appendix are presently defined FLTSATCOM protocols, the contention-based, priority oriented, demand assigned (CPODA) protocol developed for the DARPA packet satellite studies, and the multiplexing scheme available on the demand assigned multiple access (DAMA) modem developed as part of the TRI-TAC program. This appendix is restricted to the consideration of time division multiple access (TDMA) protocols; frequency and code division multiple access protocols have been omitted because the use of either of these techniques would require substantial alteration in shipboard equipment.

TERMINOLOGY

Since some of the terms used here are used in different ways by other authors, definitions used in this study are listed below:

Protocol:	Refers both to an agreement about how data transmitted over a channel are to be interpreted (*i.e.*, the precise data formats) and also to an agreement among users as to how access to the channel will be controlled (when and how each user may transmit data).
Channel:	An rf communications path. The characteristics of primary interest to us are its transmission rate, error rate, and availability. We may also be interested in whether it is simplex or half/full duplex.

Terms used in describing protocols include:

Fixed Assignment:	A static control policy. Each user is assigned a fixed time period within a frame, and time not required by one user cannot be reassigned to another.
Polling:	A policy in which a master station requests traffic from subscriber stations. Traffic is only transmitted in response to a poll.
Reservation:	A packet of control information transmitted by a subscriber requesting the allocation of some channel time for the transmission of its traffic.
Contention:	A protocol that allows two or more subscribers to transmit simultaneously (without violating the protocol) resulting in a collision is said to include contention.

PROTOCOL DESCRIPTIONS

The Fleet Broadcast, Secure Voice, and SSIXS protocols, because of their limited applicability to a consolidated system, are described only briefly. CUDIXS, TACINTEL, and TADIXS protocols are similar and are described together. Next the contention-based ASWIXS (formerly the TSCIXS) and CPODA protocols are reviewed, followed by an overview of the planned DAMA operating modes. The descriptions of FLTSATCOM protocols are derived from Refs. A1 through A8; CPODA is described in Refs. A9 through A11; and the DAMA control procedures are discussed in Ref. A12.

Fleet Broadcast, Secure Voice, and SSIXS

The FLTSATCOM channel for the Fleet Broadcast is a simplex channel operated in a fixed assignment mode. Fifteen subchannels are defined, each operating at a rate of 75 bps in a fixed time division multiplexed (FTDM) mode. A 16th subchannel is used for frame synchronization. Except to the extent that a consolidated system might use the Fleet Broadcast as a strictly shore-to-ship transmission path, this subsystem is not important to the protocol aspects of this study. FTDM (or FTDMA) protocols are best suited to environments where users' traffic demands are both high and stable.

The FLTSATCOM secure voice subsystem transmits encoded, digitized voice signals at 2.4 kbps. Operation is half-duplex (push-to-talk). Two channels on each satellite are dedicated to secure voice, and the channel designated as backup for Fleet Broadcast may be used for secure voice when the broadcast does not require it. Two access methods are defined for the voice channels:

1. If a voice channel is not busy, a user may contact another user directly or *via* a voice net controller.

2. If the channels are busy, the user may send an operator-to-operator text message to the voice controller *via* the CUDIXS network.

The first of these methods corresponds loosely to a so-called carrier-sense multiple access algorithm—a "contention" protocol: the second is a type of reservation scheme, using a separate channel (CUDIXS) for the transmission of reservations to the central controller. Both modes are implemented as manual procedures and are thus unlikely to be applicable in a pooled system with short time delay requirements.

The SSIXS network moves traffic between shore stations and submarines on a single FLTSATCOM channel. The shore stations broadcast messages to submarines at scheduled times during the day. At other times, any submarine may use the channel on a radom access basis: when the ship wishes to send traffic, it does so without regard to other ships. Thus, collisions may occur. This protocol relies on submarines to transmit data relatively infrequently so that collisions will be rare. In a consolidated system, this assumption would be unlikely to hold; thus the current SSIXS protocol is not likely to be directly applicable to a consolidated network scheme.

CUDIXS, TACINTEL, AND TADIXS

This section outlines the common features of these three IXS protocols; the following section supplies details on the distinctive aspects of each of the systems. The users of these systems are called subscribers. A distinguished user, called the Network Control Station (NCS), maintains control of the system as a whole. The time during which a subscriber's transmitter is active is called a Transmission Unit (TU). The information sent during a TU has fixed and varying length parts. The data to be transmitted by the system (*i.e.*, the information that the end user is interested in) are sent in the varying part of the TU. These data are considered to be a sequence of messages or portions of messages. Messages are the information units of interest to the end user.

The communication system subdivides messages into blocks for transmission purposes. (In some cases, there is an intermediate level between a message and a block called a Message Unit (MU) or a

Logic Unit (LU).) The varying part of a TU may contain one or more blocks from one or more messages, depending on the length of the traffic present in the system and on instructions received from the NCS. The NCS provides such instructions in its TU.

An important part of the NCS TU is the Sequence Order List (SOL), which defines the order in which subscribers may transmit. A network cycle consists of a sequence of TUs, the first being the NCS's, followed by TUs from the active subscribers in the order defined by the SOL. A typical network cycle (from TACINTEL) is shown in Fig. A-1. NCS and Subscriber TUs (from CUDIXS) are shown in Figs. A-2 and A-3.

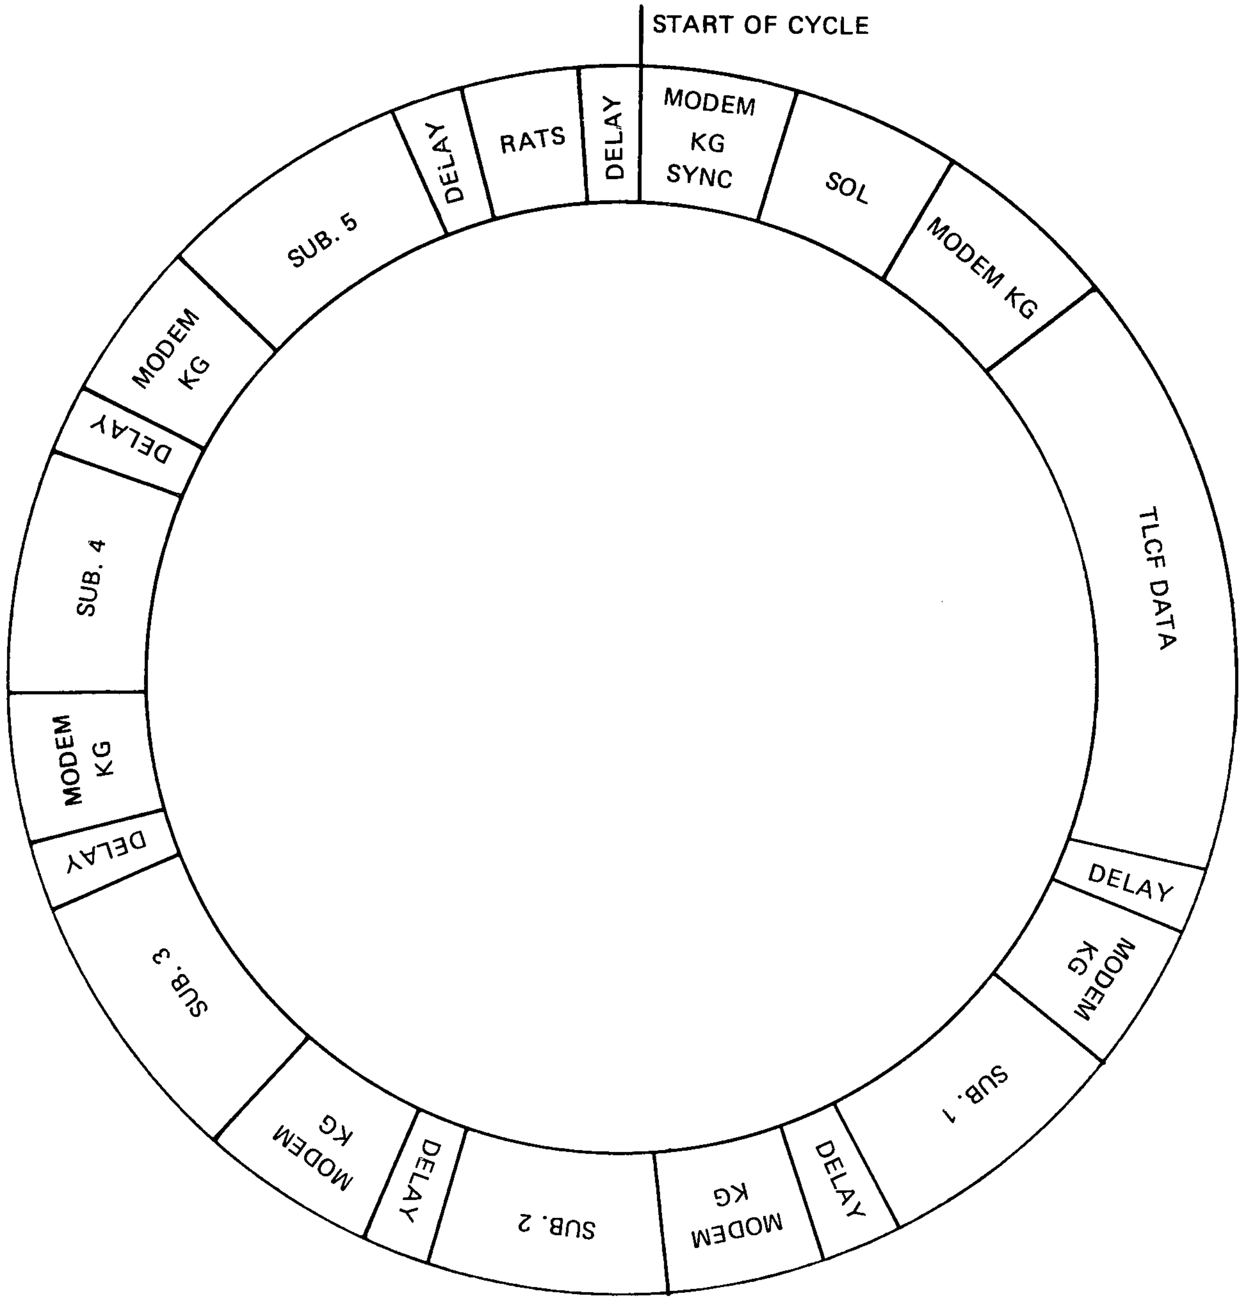

Fig. A-1—TACINTEL net cycle

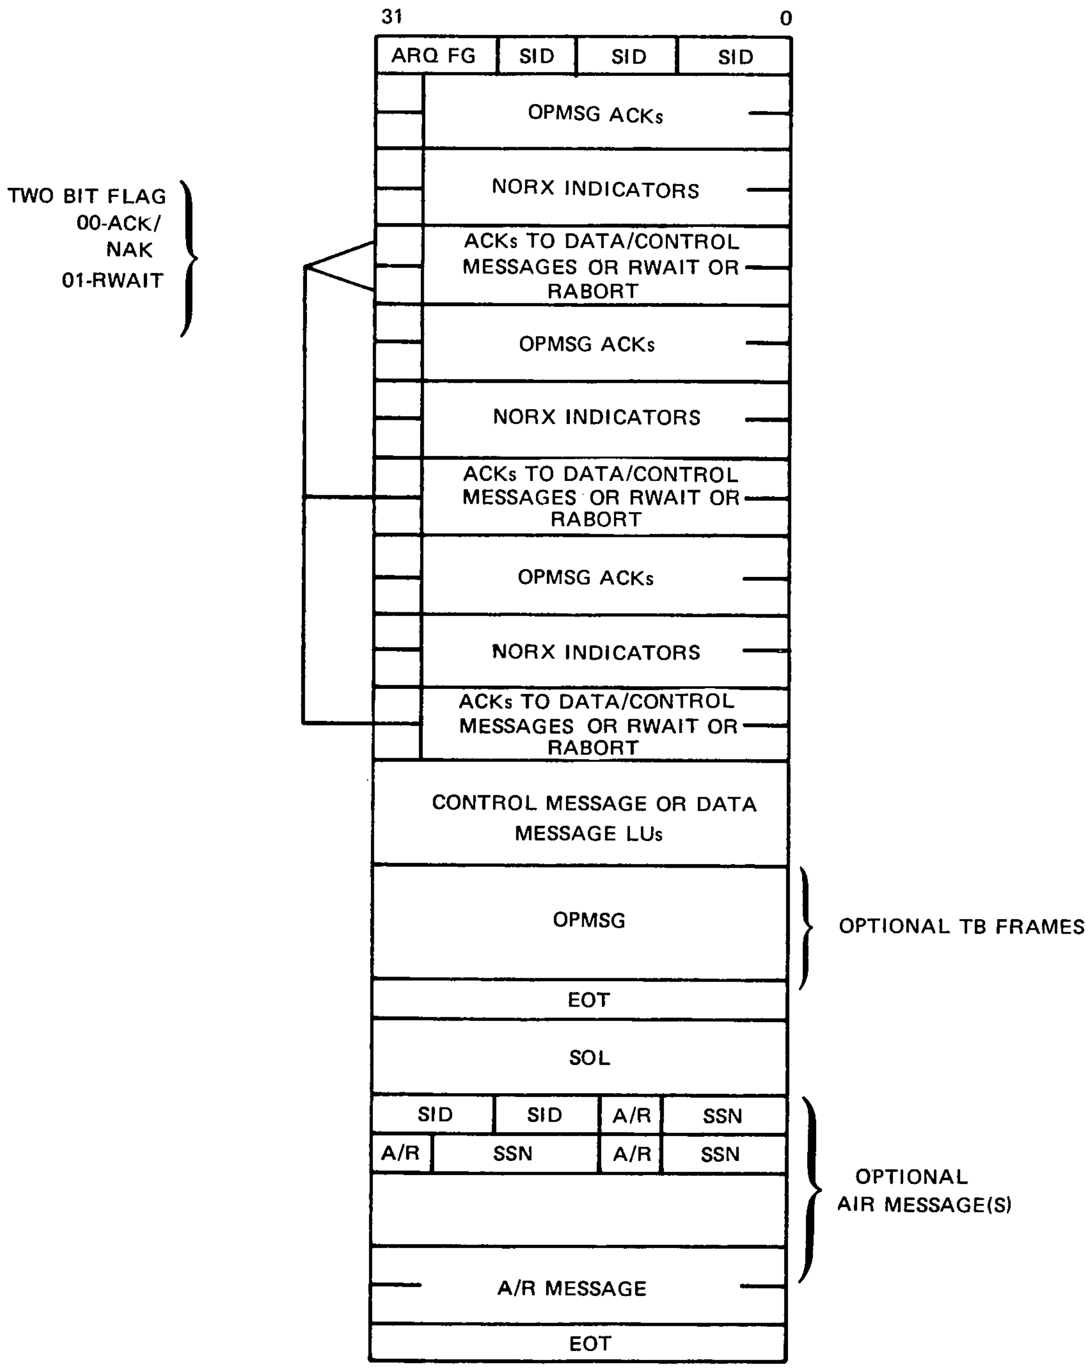

Fig. A-2—CUDIXS NCS transmission unit

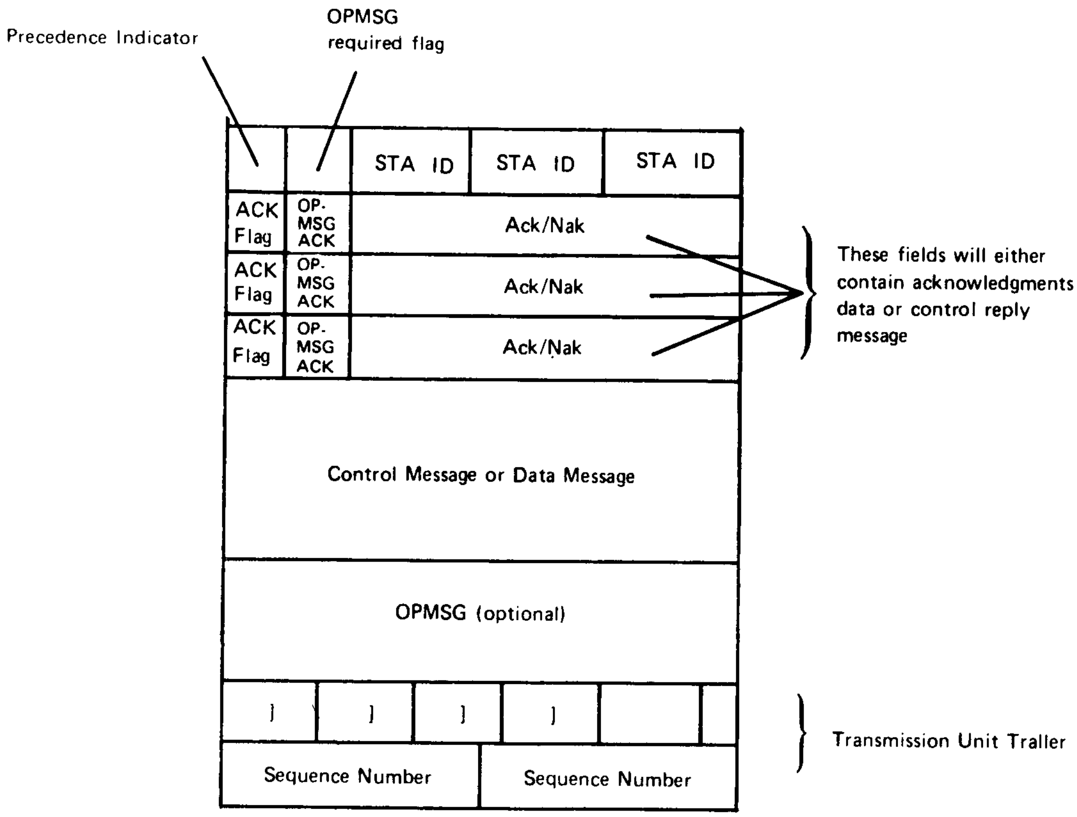

Fig. A-3—CUDIXS subscriber transmission unit

If not all subscribers are active, one or more Random Access Time Slots (RATS) are included, usually at the end of the cycle. During these periods, subscribers desiring to become active may transmit required identifying information (usually the fixed part of a standard subscriber TU) so that the NCS will include them in the next SOL. If more than one user transmits in a RATS (and the NCS detects this occurrence), the NCS transmits a special (ALL CALL) SOL that includes all possible system subscribers. In the following cycle, the SOL is reduced to include only those stations that responded to the "ALL CALL."

Network Control Station (NCS)

The NCS acts as central controller for the system, defining which subscribers may transmit, when they may transmit, how much they may transmit, and what priorities of traffic are allowed. Backup for the NCS is usually provided by defining a subscriber station to assume that role if the NCS should fail. Moving to backup mode may require manual intervention; complete information on this point has not been located.

Network operation necessarily begins with an NCS TU. The initial SOL may include subscribers known to desire to communicate, or it might include only a RATS, or it might be an ALL CALL. This choice is apparently up to the operator at the NCS. In some cases, the NCS transmits an INITIATE control command. Presumably, network shutdown occurs by the transmission of control messages among participants, and the NCS ceases to transmit. In some cases, the NCS operator can transmit a TERMINATE control command to subscribers. Subscriber stations then refuse to accept any new traffic and transmit only the messages already queued.

Information in the NCS TU for CUDIXS, TADIXS, and TACINTEL includes:

- NCS station identifier
- SOL substructure, that is, the list of subscribers to transmit in next cycle and related control information
- Error control coding and redundancy for SOL substructure
- Acknowledgment/Negative Acknowledgment (ACK/NAK) information concerning data received in the last cycle (sent redundantly)
- Address information and data blocks to be delivered to subscribers, with error-control

More error control is provided on the control fields than on the data to be transmitted. Each TU (including the NCS TU) begins with a preamble to allow a guard space between different unit's transmissions, and to allow modem and crypto synchronization.

Subscriber Station

Subscriber stations collect traffic to be transmitted from local sources (including other automated systems) and transmit this traffic according to precedence and priority rules of the system in response to control from the NCS. A subscriber station also receives traffic from the network, acknowledges the traffic, and provides for its local distribution. Subscriber stations may provide a variety of ancillary functions as well, such as journaling, and local distribution of messages. Some subscribers may be designated as backup stations for the NCS.

Subscriber operation presumably begins when the subscriber's receiver is turned on and tuned to the network operating frequency. The station then may listen for the NCS TU. If the station is already included in the SOL, it may then begin transmitting (and receiving) traffic. Otherwise, it must wait for a RATS and transmit its identifying information as required. A subscriber may cease operation in one of several ways. Normally, the subscriber station will send a control message indicating that it is terminating operation. The NCS may then eliminate the station from the SOL. If a station simply stops transmitting in its assigned position (*e.g.*, because of equipment malfunction), the NCS may leave the station in the SOL for a fixed number of cycles and then remove it. Emissions Control (EMCON) operation is generally provided for by allowing the subscriber to send a control message indicating it is entering EMCON mode. In this mode, the station will not be expected to acknowledge traffic sent to it. Traffic may be sent to the station redundantly to ensure its receipt.

A subscriber TU generally includes the following items:

- Sending station identifier
- ACK/NAK information concerning data received in the previous cycle
- Control and status information for the NCS (operator messages, queue lengths, machine status, *etc.*)
- Address information for blocks to be transmitted

- Actual data blocks to be delivered
- Error control information for all of the above

As in the NCS TU, more error protection is provided for the control information (ACK/NAK, addresses, *etc.*) than for the data itself. Also as before, each subscriber TU begins with a guard space and bits for modem and crypto synchronization.

CUDIXS

CUDIXS provides for the largest number of subscribers of the three systems: up to 60 subscribers are allowed per CUDIXS net. Subscribers are designated as either "primary" or "special." ("Primary" derives from the old designation of HF circuits as primary ship-shore circuits.) CUDIXS primary users send information in only one direction—from ship to shore. Responses are generally obtained *via* the Fleet Broadcast, which acts as a simplex shore-ship circuit. CUDIXS special users may also receive traffic from shore (*via* data blocks transmitted by NCS). CUDIXS allows up to 10 of the 60 users in a net to be designated as special users. We are not aware of any differences in equipment required for special users as opposed to primary users. CUDIXS is an operational system and is installed on more than 125 ships.

A CUDIXS primary user generally observes the following mode of operation:

a. The ship accumulates traffic to be sent over CUDIXS.

b. A shipboard WSC-3 radio is tuned to the CUDIXS net (it may previously have been operating on the FLTSAT secure voice channel, for example).

c. The ship listens for an SOL with a RATS

d. The ship transmits a request for service into the RATS, including:

   - ship identifier (SID)
   - highest precedence traffic queued
   - number of data blocks the subscriber wishes to transmit in next cycle

e. The ship listens for an SOL including its SID. It transmits its traffic as scheduled, and includes with the traffic updated values for the highest precedence message in the queues and for the number of data blocks it wishes to send in the next cycle.

f. When the ship's traffic queue is empty (after several net cycles, perhaps), the subscriber terminates transmission to the net. The WCS-3 may be diverted for use by the other systems.

A CUDIXS special subscriber operates similarly, except that it will generally remain active in the net even when it has no traffic to send, since it may receive traffic from shore.

The CUDIX NCS constructs the SOL in response to the subscriber data on queue size and precedence. The algorithm for constructing SOL allotments places subscribers with the highest precedence traffic earliest in the SOL. Users with the same precedence traffic are ordered first come—first served (FCFS). The number of RATS per three net cycles is set by the CUDIXS operator. The placement of the RATS into the three cycles is random.

CUDIXS NCS Transmission Unit

The CUDIXS NCS TU includes specialized acknowledgment fields for different types of messages and also fields where nonreceipt of a TU (as opposed to receipt of a message in error) can be indicated. These are termed NORX indicators. The NCS TU consists of the following fields:

a. ARQ (Automatic Repeat Request) Flag: If on, this TU contains messages retransmitted in response to user NAKs received in the previous cycle. It is unclear whether this means that the entire TU from a previous cycle is being repeated or only some of the messages. (1 bit, 8 times redundant = 8 bits)

b. NCS Station ID (8 bits, 3 times redundant = 24 bits)

c. Acknowledgment Data:

   - OPMSG ACKS to ACK/NAK operator messages. (1 bit per possible subscriber = 60 bits)
   - NORX indicators to denote whether a TU was received from a subscriber in the last frame. (1 bit per possible subscriber = 60 bits)
   - ACKs to data/control messages. (One bit per possible user to acknowledge receipt of a data or control message = 60 bits)

   This field may also be used to indicate that the shore station must terminate receiving (RABORT) or that each subscriber must retransmit its previous transmission. (Total length = 60 × 3 = 180 bits, 3 times redundant = 480 bits)

d. Control or Data Message Logic Units (LUs): Control messages are implicitly addressed to all subscribers. Typical messages include INITIATE, TERMINATE, TRANSMIT WAIT, TRANSMIT ABORT, and NULL (sent when there are no control or data messages to be sent). Data message LUs contain message blocks addressed to special users. The shore TU may contain up to 24 LUs, with each LU containing from 1 to 5 data blocks. (Total length: minimum with redundancy = 64 bits; maximum including addressing and redundancy = 74,448 bits)

e. Operator-Operator Message: This is an optional section. If used, it is a message to be transmitted to the operator at another station. The station address of the destination is included. This is the field sometimes employed to coordinate use of the FLTSAT Secure Voice channels. (Total length: minimum = 0, maximum = 18 × 32 bits = 576 bits)

f. Sequence Order List: This is the list of subscribers in the order in which they are to transmit. Each subscriber station ID is preceded by an indication of the number of data blocks that station may transmit: 0, 3, 6, or 10 blocks will be allowed. (Length per entry = 6 bits per ID, 3 times redundant = 18 bits, plus 3 bits control x 4 times redundant = 30 bits)

g. End of Transmission (EOT) Marker: Provides timing for stations to compute their transmission starting time. (Length = 32 bits)

CUDIXS Subscriber Transmission Unit

Each active subscriber computes its transmission time by searching the SOL until it finds its station identifier. On the basis of the number of data blocks each of its predecessors in the SOL is allowed to transmit, the subscriber can determine the absolute time at which it must send its traffic. Thus, it makes no difference whether or not a preceding station actually utilizes its assigned time or not; the subscriber must wait until its turn. Of course, the NCS can revise the SOL each cycle to provide appropriate service to busy stations.

A subscriber transmission includes the following fields:

a. Precedence Indicator: Indicates highest precedence message held in queue at this station. (Length = 4 bits)

b. OPMSG Required: If set to one, indicates this station has an operator-to-operator message to send. The NCS may then provide an SOL allocation for transmission of the message. (Length = 4 bits)

c. Station ID: (8 bits × 3 times redundant = 24 bits)

d. Acknowledgment Field: Contains 24 bits, 1 for each of the up to 24 LUs the NCS sent in the previous cycle. This field may alternatively be used to send a control message to the NCS, indicating that the subscriber was unable to receive the previous NCS TU (*e.g.*, no crypto sync or no buffer space).

e. Control Message or Data Message LUs: The possible control messages (sent if the SOL assigned this station 0 data blocks) are similar to those that the NCS may send. Ordinarily, this field will include the traffic that the ship wishes to send to the shore. It is formatted as sequence of data blocks, each with a longitudinal parity check. Partially filled data blocks at the end of a message are padded out with ETXs. For retransmission purposes, the entire ship TU is considered a single unit. (Length of Message LU = minimum (3 blocks) 576 x 3 = 1728, maximum = 5760 bits)

f. Data Memory LUs Requested: Number of data blocks (0, 3, 6, or 10) requested to be allocated (by the NCS) in the next cycle.

g. Operator Message: Optional field, if present, same format as NCS operator-operator message. (May only be present if so indicated in SOL.)

h. End of Transmission: 32 bits, fixed format.

TACINTEL

Up to 24 subscribers are supported in a single TACINTEL network: one TACINTEL Link Control Facility (TLCF) that functions as the NCS, and up to 23 subscriber stations. One subscriber may be designated as the secondary TLCF; this station acts as NCS in the event that the shore-based TLCF fails. Operations are degraded when the secondary TLCF is in operation if it is a shipboard subscriber (as is intended). Unlike CUDIXS, TACINTEL allows shore/ship and ship/shore communication to all of its subscribers. In addition, some ship/ship communication is possible. Messages transmitted directly between ships are restricted to two short formats, BE-3 and short OPSCOMM. BE-3 messages have a maximum length of one data block (80 characters), and short OPSCOMM messages are at most two data blocks (160 characters). To facilitate ship/ship communication over half-duplex equipment, TACINTEL includes a 280-ms guard time between the end of one subscriber's transmission and the beginning of the next. In addition, each transmission (at 2.4 kbps) includes 115 ms for modem sync and 353 ms for crypto sync as in CUDIXS. Although the system is defined to be capable of operation at 1.2, 2.4, 4.8, or 9.6 kbps, no parameters are defined for speeds other than 2.4 kbps. The capabilities of the ON-143 connecting group are cited as the reason for the given set of transmission rates.

TACINTEL is quite similar to CUDIXS in its general operation. The NCS sends an SOL from which each subscriber directly determines the time at which it may transmit. Since TACINTEL is intended specifically for SI message traffic, traffic flow security is a concern. The requirement is apparently met by requiring each subscriber to transmit during the entire period assigned to it whether or not it has actual data to be sent. Mechanisms are included for altering the transmission time allotted different subscribers, but apparently these would not be invoked frequently. The assignment of subscriber ordering and transmission length per subscriber is apparently done manually by an operator at the TLCF. Unlike CUDIXS, a high-priority TACINTEL subscriber may be assigned two positions in the cycle for transmission of its data. A high-priority subscriber transmits half of its assigned Message Units per cycle in each of its two assigned positions. After receipt of a high-priority message locally, such a high-priority subscriber need wait only about a quarter cycle (on the average) before it may transmit the message; ordinary users must wait an average of a half cycle. Maximum cycle length in TACINTEL is approximately 480 s (at 2.4 kbps transmission rate).

TADIXS

TADIXS was designed to support a relatively small net of users transmitting digital data with requirements for low delays. This requirement was imposed by the planned communication between a shore-based Fleet Command Center (FCC) and shipboard Tactical Flag Command Centers (TFCCs). The TADIXS design is intended to accommodate up to 11 subscribers, representing 10 TFCCs and 1 FCC. The FCC acts as NCS and transmits the SOL to the other members of the net. TADIXS is more flexible than either CUDIXS or TACINTEL in its information transfer capabilities; traffic can pass in either direction between any two nodes in the network.

Two traffic priorities, normal and alert, are included in the design. To meet the delay requirements on alert messages and the accuracy requirements for digital data, the design includes some features different from those in CUDIXS and TACINTEL. For example, the transmission rate called for is 4.8 kbps as opposed to 2.4 kbps in the previously described systems. In addition, special Forward Error Correcting (FEC) codes and an interleaver are included in order to combat shipboard Radio Frequency Interference (RFI) on the downlink. Finally, the time base used by a subscriber to calculate its transmission time is different. In CUDIXS and TACINTEL, each subscriber examines the SOL to determine how much time will be required by each of its predecessor's transmissions. The subscriber then schedules the start of its own transmission as a fixed displacement from the end of the NCS transmission. If a user should not require the entire allotment provided to it by the SOL, the unneeded time is simply wasted. (In TACINTEL, the unneeded time is padded by repeating the message already sent and provides traffic flow security.) In TADIXS, the subscriber observes its position in the SOL and notices which subscriber is its predecessor. It then listens for the predecessor's transmission. The first part of this transmission indicates the actual length of the data to be sent (*i.e.*, the number of data blocks). Based on this length and the knowledge of the propagation delay, the subscriber can then initiate its own transmission so that there is a minimum of dead time on the channel.

Since the design is for a half-duplex system, if the final data block in the predecessor's transmission is addressed to the station occurring next in the SOL, that station cannot, of course, avoid introducing a dead time of one round-trip propagation delay into the channel. (That is, the addressed station must remain in receive mode until the final block is received; only then can the modem be turned around from receive mode to send mode. This turn around time may be up to 40 ms in a typical shipboard environment; the WSC-3 by itself introduces a 7-ms delay.

Each subscriber may transmit up to three data blocks of high-precedence data in its TU. If the subscriber has fewer than three high-precedence data blocks to be transmitted, it may send normal-precedence data in the remaining lock (under the following conditions):

1. The subscriber may not transmit more than the currently defined maximum number of low-precedence blocks per TU. The value of this number is included in the SOL and may be 0, 1, 2, or 3.

2. If, at the time the subscriber begins transmitting, more than a certain number of seconds has passed since this cycle was initiated, the subscriber may only transmit high-precedence blocks. The value of this number of seconds, called the threshold, is also included in the SOL.

ASWIXS

The protocol for ASWIXS differs substantially from the polling techniques used in CUDIXS, TACINTEL, and TADIXS, but there are still some points of similarity. There is an ASWIXS Net Control Station (NCS) that broadcasts a Net Control Block (NCB). The NCB defines the current operating mode of the network and may designate which subscriber is to transmit. There is a cycle (defined as the time between NCB transmissions), but it typically includes only a single subscriber transmission. Two priorities of traffic are supported. Example operating patterns for ASWIXS are shown in Figs. A-4 through A-7.

Up to 60 subscribers may be served by a single ASWIXS net; the NCS will normally be a shore based ASW Operations Center (ASWOC), formerly VP Tactical Support Center, and the other subscribers will be P-3C aircraft. The network is designed to carry digitized secure voice as well as digital data. Traffic from one subscriber can be addressed to one or more of the other subscribers. In place of ARQ error control on data traffic, each data transmission is repeated three times. This technique facilitates EMCON operations; any user may receive traffic passively.

Different equipment is utilized in ASWIXS than in the previously described systems, but the data transmission rate is still 2.4 kbps, the modulation is still DPSK, and operation is half-duplex. The computer system implementing the protocol is microprocessor-based and is called the Memory Storage Buffer (MSB) or the ON-143(V6). KG-36 crypto units ashore, KG-35 on the aircraft, and CV-3333 voice digitizers are employed. The existing ARC-143B transceivers are proposed to be utilized by all subscribers. The system is presently in a development phase; equipment procurement has not been approved as yet.

Cycle Structure for ASWIXS

There are four modes defined for ASWIXS and two priority levels. The modes are (1) Idle, (2) Teletype, (3) Data Link, and (4) Voice. Operations in Teletype and Data Link modes are virtually identical. The priority levels are Flash (high) and Immediate (low). The precise structure of a net cycle depends on the current mode and priority, but net cycles can be categorized into the following three groups:

Idle Cycles

In idle mode, in the NCB is followed by a period in which subscribers may transmit requests for service. The subscriber requests service by transmitting its station identifier and the desired mode into either the Precedence Request Time Slot (PRTS) (for Flash priority) or one of the following General Access Time Slots (GATS) (for Flash or Immediate priority). The GATS utilized is chosen randomly by the subscriber.

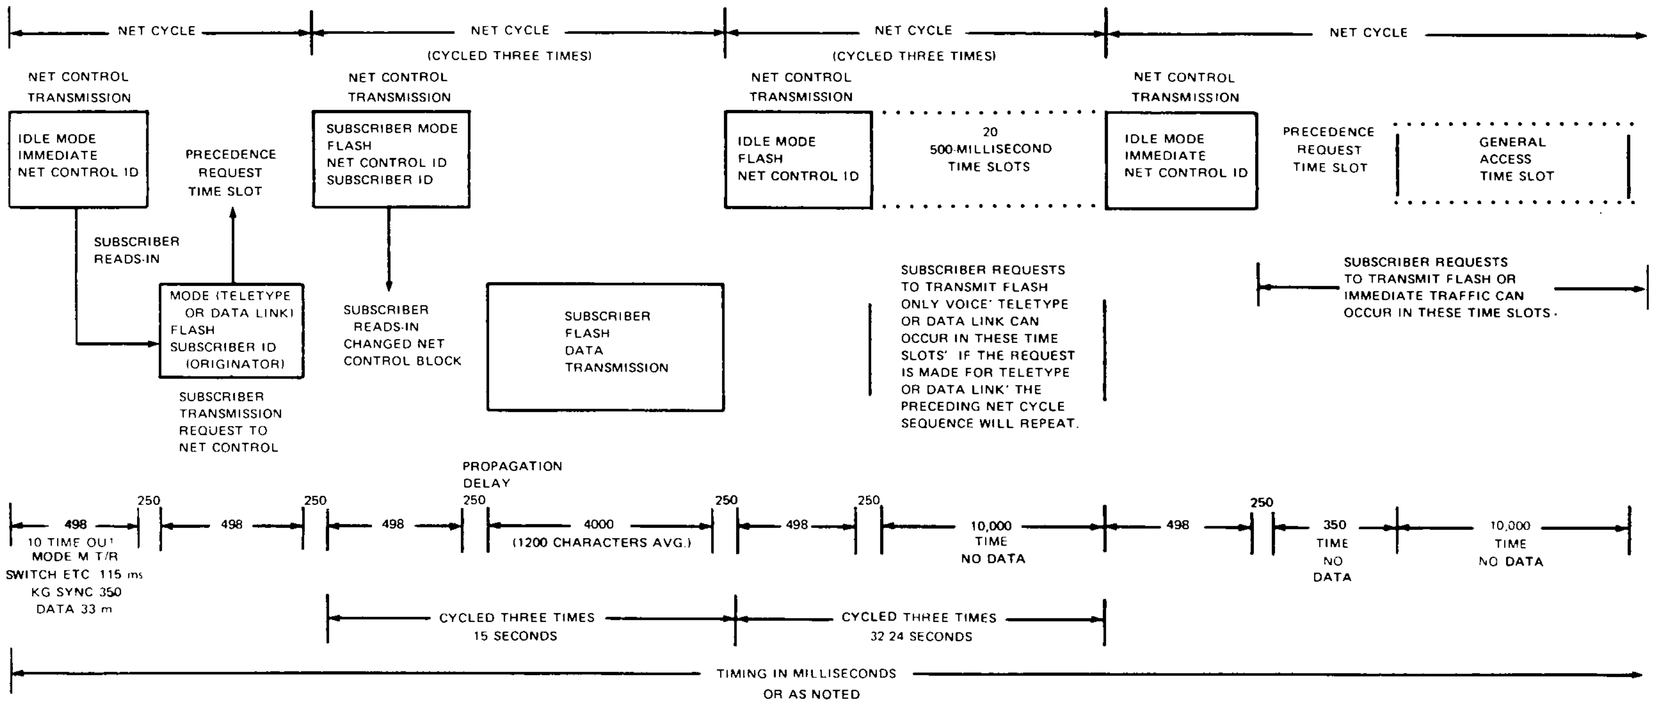

Fig. A-4—ASWIXS—flash data transmission

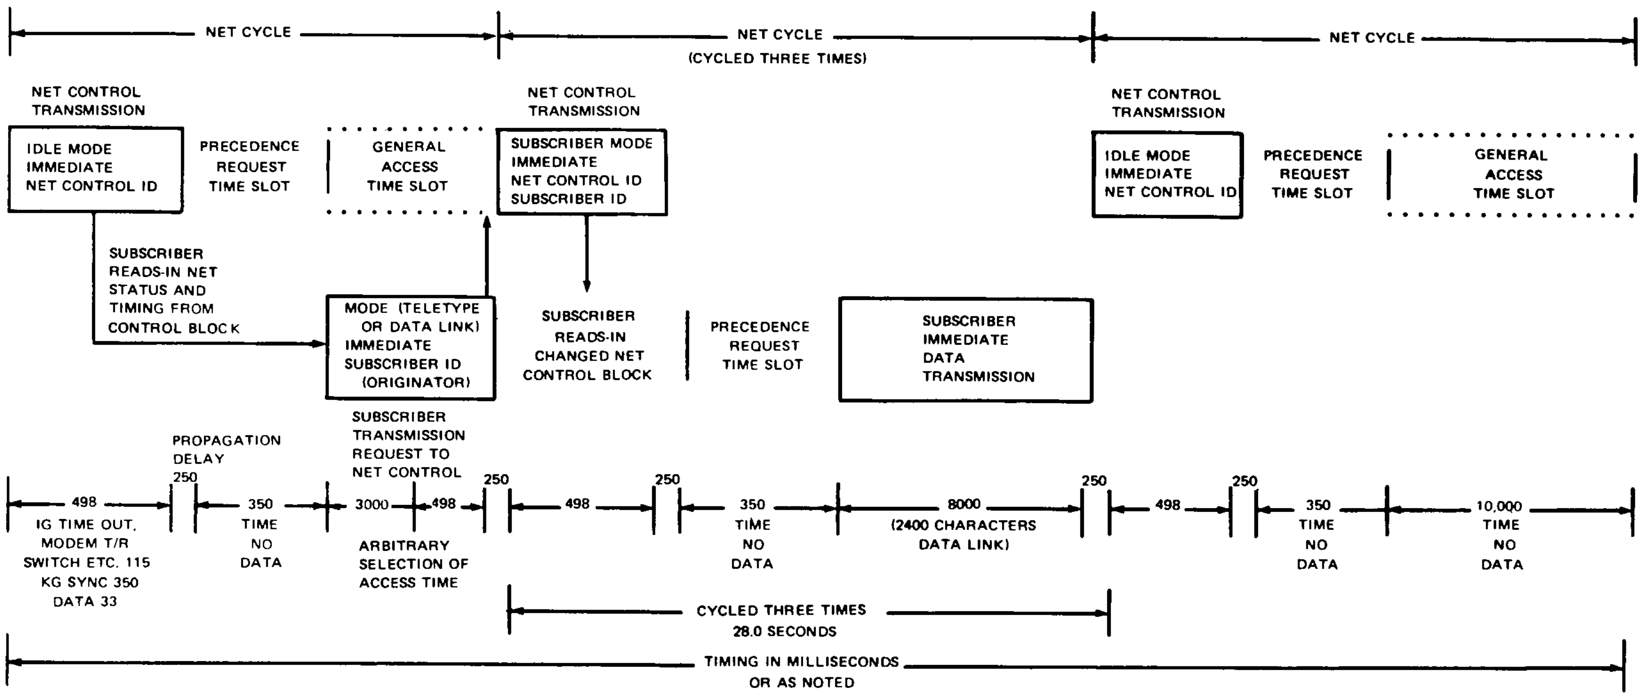

Fig. A-5—ASWIXS—intermediate data transmission

UNCLASSIFIED

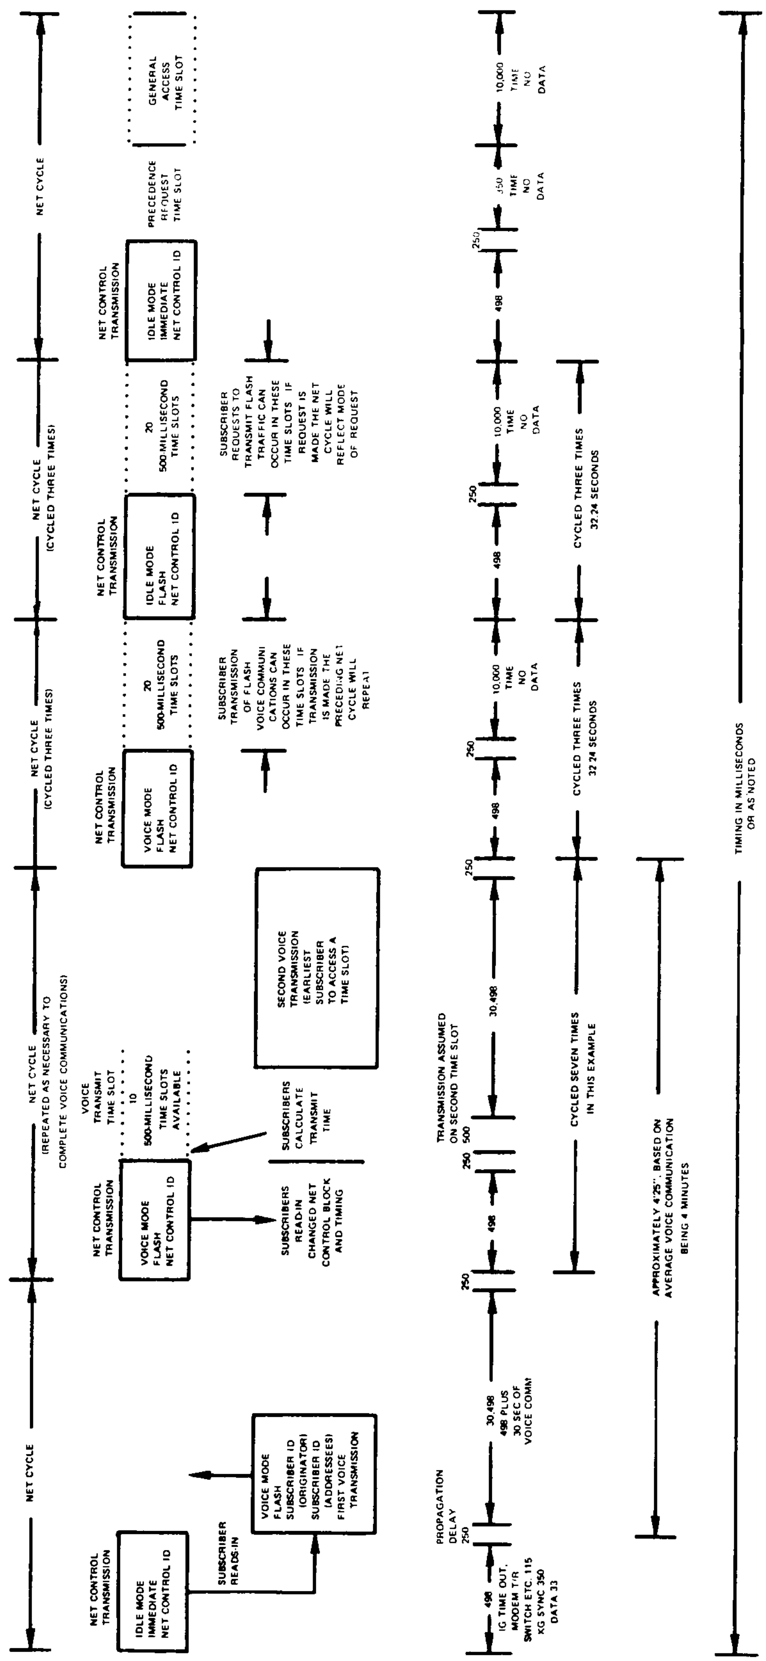

Fig. A-6—ASWIXS—flash voice transmission

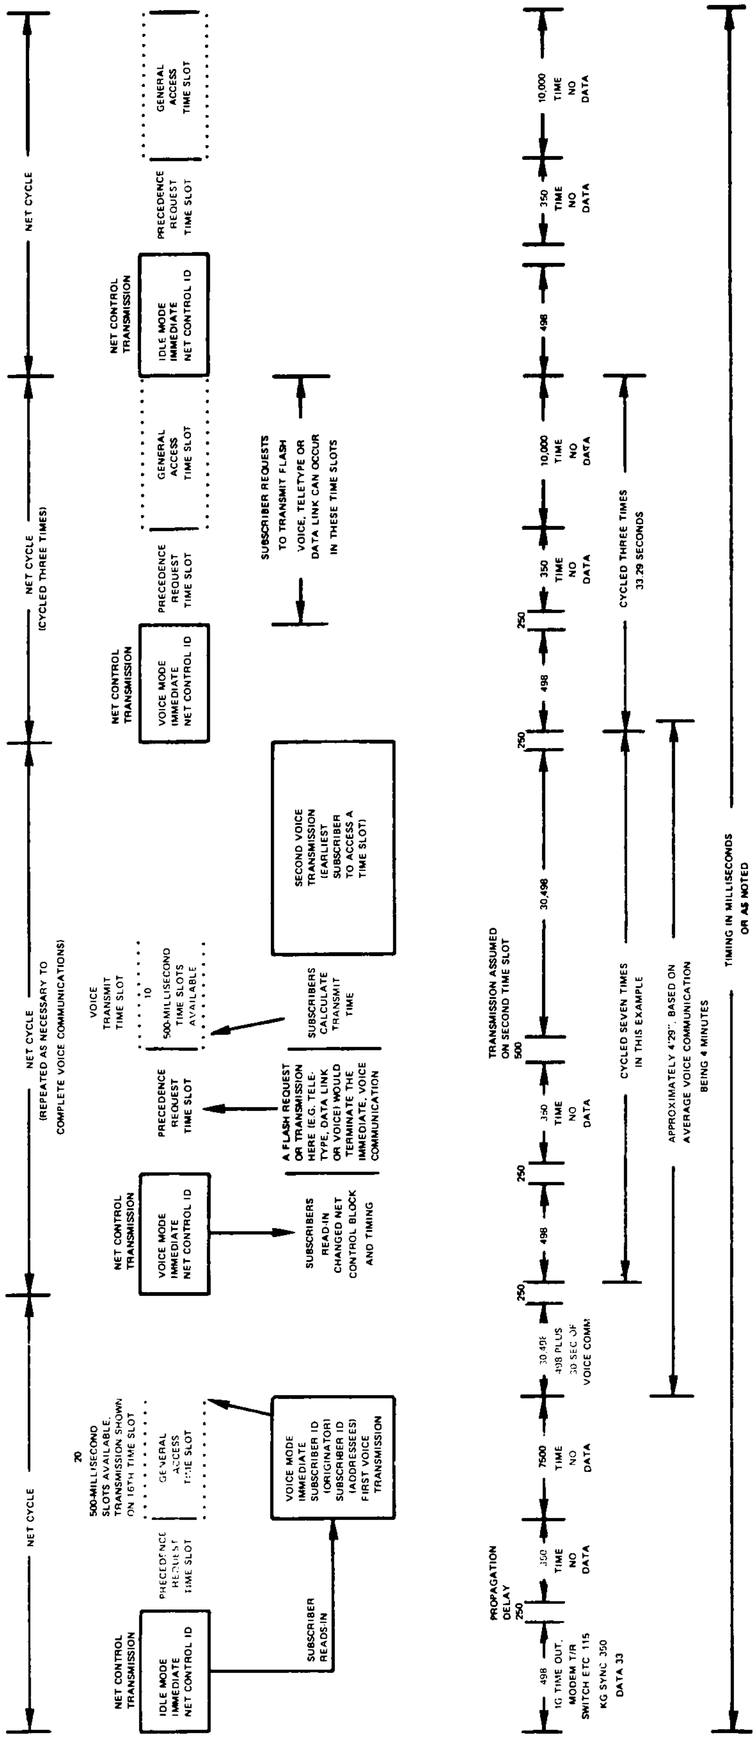

Fig. A-7—ASWIXS—flash data transmission

Data Transmission Cycles

Once a user request is heard by NCS, it constructs a new NCB indicating the mode and priority requested and the subscriber that is to transmit. The subscriber then transmits the data to be sent. The same NCB is next repeated by the NCS, followed by the same subscriber response. This pattern is repeated a third time before the NCS again sends an NCB indicating a return to Idle mode.

Voice Transmission Cycles

Since repeating each voice transmission three times would result only in confusion, a different strategy is adopted in this case. A request for voice service is transmitted similarly to a request for data service, except that the subscriber also transmits the voice data (up to 30 s in length) immediately following the request data (the subscriber does not wait for the NCS to respond).

Following the initial voice transmission, the NCS transmits an NCB indicating voice mode and the appropriate priority level, but without designating a subscriber. Each subscriber wishing to talk then generates a random number to determine which of the 10 voice transmit time slots following the NCB it will use. The subscriber that accesses the earliest of these slots then sends its voice transmission of up to 30 s, and so on.

The priority of the network is defined independently of the mode; the priority and mode together determine the state of the protocol (*i.e.*, the frame structure to be used). A description of the ways that Flash priority may be entered when the network is operating at the immediate level in each mode is given below.

Idle Mode

A subscriber may transmit a request to send Flash traffic (voice, data, or teletype) in either the PRTS or any of the following GATS. The following NCB will indicate Flash priority and the requested mode. No PRTS or GATS will be in the following cycles until transmission of Flash data is completed, since no higher priority traffic can preempt Flash traffic.

Data or Teletype Mode

An Immediate priority frame of either of these types contains a PRTS immediately after the NCB and before the user transmission. A subscriber with Flash traffic may use the PRTS to transmit its request. The Immedite priority user and the NCS will respond by sending an NCB indicating Flash priority and the requested mode.

Voice Mode

Similar to the previous case, each voice mode NCB at Immediate priority is followed directly by a PRTS. The subscriber with Flash traffic transmits its request in the PRTS, the previous Immediate level) subscriber suppresses its transmission, and the NCS responds appropriately.

Following a Flash transmission, the NCS will place the network at Flash level in Idle mode. This frame contains no GATS but has 20 PRTSs to allow subscribers with queued Flash traffic (suppressed during the just completed Flash transmission) to so indicate. Again, each such user must compute a random number from 1 to 20 and transmit its request in the corresponding slot. The user that transmits first in the sequence of 20 slots is the next one served, and its transmission terminates the sequence of request slots. The request slots are 500 ms in length to allow subscribers to monitor the previous slot before transmitting into the next one. If no subscriber broadcasts a request in the Idle mode, Flash level cycle, the NCS returns to Idle mode, Immediate level.

CPODA—Distributed, Contention-Based, Priority-Oriented, Demand Assignment

This protocol was developed and used by BBN and Linkabit in DARPA-supported packet satellite experiments. It is presently planned for use in the Mobile Access Terminal (MAT) project, in which computer terminals on several ships will communicate *via* satellite to shore-based computer facilities. CPODA will be used to coordinate the transmission from the shore and the various ships over the satellite link.

Since this protocol is not part of a presently operational Navy system, precise formats (determining the numbers of bits in various identification and acknowledgment fields) have not been specified. This section therefore describes the general structure of the protocol as it is proposed for use in the MAT project.

Frame Structure

In CPODA, time is divided into fixed length *frames*. Once per frame a designated leader station sends a timing packet to allow all subscribers to maintain synchrony. The leader function may be assumed by any subscriber; all subscribers in the system execute an identical algorithm for scheduling the channel.

Each frame is divided into two subframes: the information subframe and the contention subframe. The portion of the frame devoted to contention varies with the traffic loading. If there is no traffic to be delivered, the entire frame (except for the leader packet) is devoted to contention. As loading increases, the contention subframe is gradually reduced to a minimum size. Figure A-8 displays the CPODA frame structure.

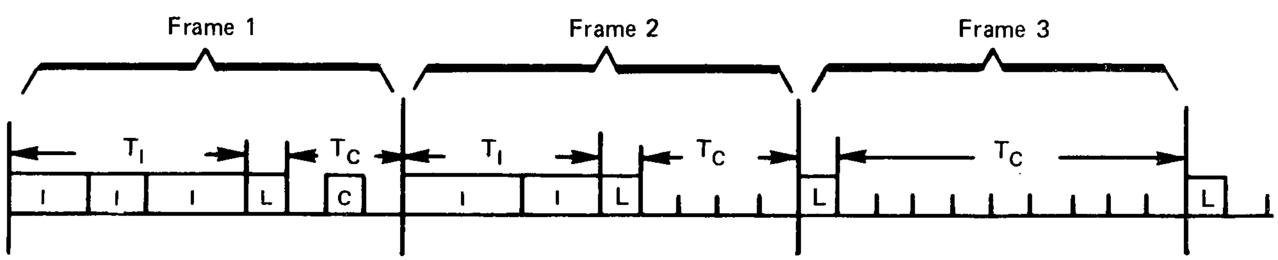

L = Leader packet
I = Information packet
C = Contention packet
T_I = Information subframe
T_C = Contention subframe

Fig. A-8—CPODA frame structure

Reservations

CPODA is a reservation-based protocol and is similar to ASWIXS in this respect. Recall that to send data in ASWIXS, the subscriber first transmits a header giving its identity and traffic precedence in a contention slot. The NCS then responds to the request to grant the channel to the subscriber. There is no queueing of requests. In CPODA, the subscriber transmits a reservation indicating its identity and its traffic characteristics. The reservation is queued by all network participants, and when the traffic corresponding to preexisting and higher priority reservations has been transmitted, the subscriber transmits (in an information subframe) the data corresponding to its reservation.

Notice that in distributed CPODA (unlike ASWIXS) there is no explicit granting of the channel; all subscribers listen to all reservations. Each user maintains a copy of the current reservation queue and uses this queue to determine which subscriber is currently allowed to transmit. Even when a given subscriber is not actively transmitting traffic, it can passively listen to the reservations and transmission on the channel and predict which subscriber will transmit next.

Reservations may also be transmitted by "piggybacking." Up to two reservations may be included with each data transmission in the information subframe. This provision allows users that have already submitted reservations to avoid using the contention subframe. Thus, as the traffic level increases and the length of the contention subframe decreases, fewer users need to use the contention subframe.

Although not planned for implementation in the MAT network, Ref. 9 also describes a "stream" reservation. This type of reservation requests a stream of reservations by specifying a period, priority, and tolerance of variation. When a stream reservation is received by a station, the station introduces a new reservation of the given priority into its queue periodically thereafter (without a new reservation being broadcast over the channel). The reservation is terminated either by an explicit message or a period of inactivity. The purpose of stream reservations is to support stream data (*e.g.*, voice or periodic data transmissions).

Message acknowledgments, when required, may be sent either by piggybacking them on a data transmission or by sending them in the contention subframe. The former method is preferred if the subscriber already has a reservation in the queue. An acknowledgment generally contains a frame number (frame numbers are provided by the leader) and a packet number within that frame to indicate which packet is being acknowledged. Thus the acknowledgment does not have a format that depends on the number of active or potential users in the network.

Contention and Priorities

Contention problems are handled as follows: the station transmitting a packet in the contention subframe monitors the channel to hear it echoed. If the echo is heard correctly, the subscriber assumes that the reservation (or acknowledgment) has been heard by the other net members. Otherwise, the subscriber queues the packet for retransmission.

The choice of exactly when in the contention subframe to transmit (or retransmit) a packet is made as follows: at the start of each contention subframe, a subscriber with a contention packet to send generates a psuedorandom number. If the number is below a certain threshold, the packet is transmitted in a randomly selected slot in the current contention subframe. Otherwise, the subscriber waits until the next contention subframe and repeats the procedure. To maintain stability, the threshold value is lowered slightly each time a packet retransmission is required and is raised slightly each time a packet is transmitted successfully.

Priorities enter the system *via* the algorithms for entering reservations into the queue and for determining which reservation is currently at the head of the line. As long as all subscribers employ the same algorithm, arbitrary priority structures can be accommodated.

Subscriber Start-up

A new subscriber entering the system must listen to reservations and transmissions until it is able to make accurate predictions about which subscriber has the next turn in the information subframe.

This ability corresponds to building an up-to-date reservation queue. Three states are defined for subscribers in this respect: initial acquisition, out-of-sync, and in-sync. A subscriber just turning on its receiver is in initial acquisition state until it makes a certain number of correct predictions. It then enters out-of-sync state until it passes a threshold number of correct predictions without an error. At this time, it enters in-sync state and may begin transmitting reservations and traffic.

An in-sync station that suffers from down link noise may occasionally miss reservations and therefore make erroneous predictions. If an in-sync station makes more than a threshold number of bad predictions, it reenters out-of-sync state until its queue is up-to-date.

If the network is idle, a new subscriber will achieve reservation sync almost immediately, but if a large backlog of reservations exists, the subscriber may have to wait several cycles before it can begin transmitting reservations. This characteristic suggests that subscribers would maintain the passive listening mode even when there is no traffic to be transmitted. (Of course, subscriber wishing to receive traffic must listen anyway.)

DAMA

The UHF-DAMA modem is designed to allow burst data rates of up to 32 ksps on a single FLTSATCOM channel. A single channel can then be shared among several FLTSATCOM user networks, each operating at its current rate of 2.4 kbps or at other rates. A DAMA frame is 1.38 s in length and may be formatted in over 500 different ways (see Appendix B).

The choice of frame format provides the "demand assignment" aspect of the scheme: a user or group of users desiring to initiate, say, a CUDIXS net, indicates this by sending signals in an orderwire portion of the frame. An operator at a central site recognizes the request and chooses a new frame format to allow the new set to be initiated. The operator can notify the members of the new net *via* a return channel control orderwire. The net members then set their equipment to use the designated portion of the DAMA frame. The DAMA program also includes plans to automate the manual process just described, so that allocation of frame formats would not require operator intervention.

Once a frame format has been established, DAMA operation is on a fixed TDMA basis. The bandwidth is demand-assigned in that a channel need not be allocated to a net that is not operating, but during the operation of a net, that net will have a fixed dedicated portion of the DAMA frame.

REFERENCES

A1. "Fleet Satellite Communication System Description," prepared by NOSC, San Diego, for NAVELEX PME-106, Washington, D.C., 5 Jan. 1978.

A2. "The Fleet Satellite Communications Program Information Exchange Systems CUDIXS and SSIXS System Description (U)," Navy Space Project Office (PME-106), NAVELEX, Washington, D.C. (Confidential).

A3. Private communication, 6 June 1978, with CDR McKenzie, Deputy Director, Naval Telecommunications System Integration Center, NAVCOMMU, Cheltenham, Md.

A4. "Baseline Type A System Specification for the FLTSATCOM TACINTEL, Rev. A (U)," prepared for Navy Space Project Office (PME-106) by Computer Sciences Corp. System Div., Dec. 1976 (Confidential).

A5. "Baseline Performance Specification for FLTSATCOM TACINTEL (U)," prepared for NAVELEX PME-106 by NELC, 18 May 1975 (Secret).

A6. "TADIXS Interim Report Network Design (U)," prepared for NAVELEX PME-106 by Computer Sciences Corp., 15 Sep. 1976 (Confidential)

A7. "TADIXS System Engineering Study (U)," prepared for NAVELEX PME-106 by Computer Sciences Corp., Sep. 1977 (Secret).

A8. "Equipment and Netting Operational Description," Appendix I of Type A Specification, Tacintel Support Center Information Exchange System, ESCS 207-1A, prepared for NAVELEX PME-106 by NELC, 31 Jan. 1977.

A9. I. Jacobs *et al.*, "CPODA—A Demand Assignment Protocol for SATNET," in *Proc. Data Comm. Symp.*, ACM/IEEE, Sept. 1977, pp. 2-5—2-9.

A10. "Mobile Access Terminal (MAT), Technical Proposal," Bolt, Beranek, and Newman, Inc., Apr. 1978.

A11. "Channel Scheduling for the CPODA Protocol," Hsu, Nai-Ting and Lin-Nan Lee, 1978 Int. Conf. on Comm., Conf. Record Vol. 3, pp. 43.3.1—43.3.5, IEEE Cat. No. 78CH135D-8CSCB.

A12. "UHF DAMA Critical Design Review," Motorola, Inc., Government Elec. Div., 14-16 Feb. 1978, Vols. 1 and 2.

Appendix B
BRIEF DESCRIPTION OF THE UHF DEMAND ASSIGNED MULTIPLE ACCESS (UHF–DAMA) SYSTEM

The ultrahigh frequency demand assigned multiple access (UHF DAMA) system is being developed in order to increase the available throughput capacity of UHF SATCOM 25 kHz channels. It provides for time division multiple-access (TDMA) operation within these channels along with the capability for the flexible reassignment of channel capacity among users by employing demand-assignment techniques. User requests for TDMA slot assignments are made during a return-order-wire time slot, and a master control station (MCS) makes the assignment of the slot during an order-wire time slot after an appropriate delay for processing the request. In its early development phase, both requests and assignments were made manually by the user and the MCS operator (manual operation). In semiautomatic DAMA, the user request/assignment function will be performed automatically with the MCS still being operated manually (semiautomatic operation). Finally, in its projected full development, UHF DAMA will have all user and MCS request/assignment functions performed automatically by computers and microprocessors, without manual intervention (fully automatic system).*

The principal element of the UHF DAMA system is the subscriber unit (SU). As shown in Fig. B-1, a subscriber unit interfaces with up to four users at baseband and provides orderwire capabilities. The radio transceiver interface is 70 MHz (IF). The subscriber unit accepts I/O data rates of 75, 300, 600, 1200, 2400, and 4800 bps. Convolutional coding with Viterbi decoding EDAC is performed with code rates of 1/2 or 3/4. The transmit/receive burst rates are 2.4 ksps, 9.6 ksps, and 19.2 ksps using BPSK modulation. Also, 32 ksps is available with QPSK modulation.

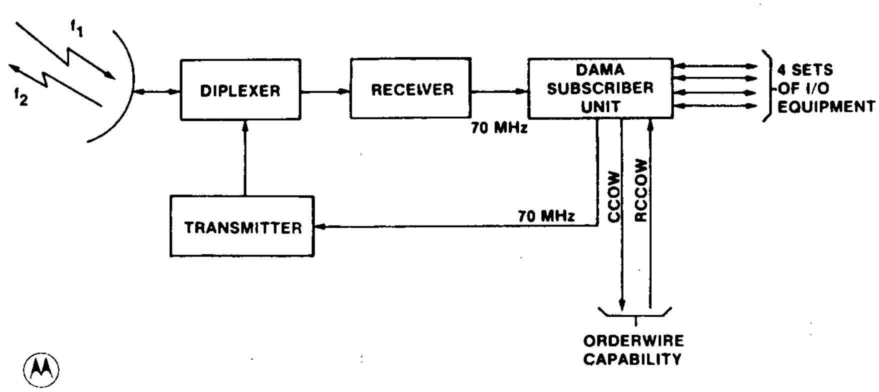

Fig. B-1—Typical user configuration

The basic DAMA signaling format structure is shown in Fig. B-2. Three groups of data time slots (designated A, B, and C) are shown. Each of these slots can be formatted by the MCS in eight possible ways as shown in Fig. B-3. A maximum of 18 users may be simultaneously utilizing the 25 kHz channel. (Note: In Fig. B-3, the three numbers in each box refer (from top to bottom) to the baseband I/O data rate, the T/R burst rate, and the code rate.)

*An important issue concerning DAMA is the turn-around time for request/assignment. For the manual system (with human operators at each end) this time will be approximately 10 min. For the semiautomatic system the turn-around time will be reduced to approximately 1.5 min, and for the full automatic system the response time wil be reduced to approximately 4 s.

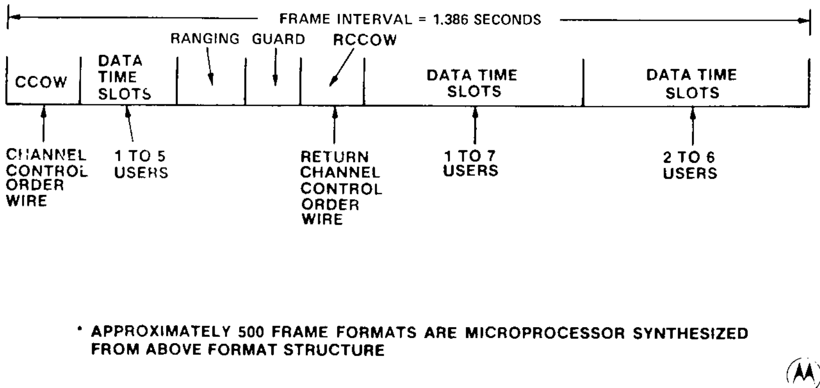

Fig. B-2—Basic DAMA format structure

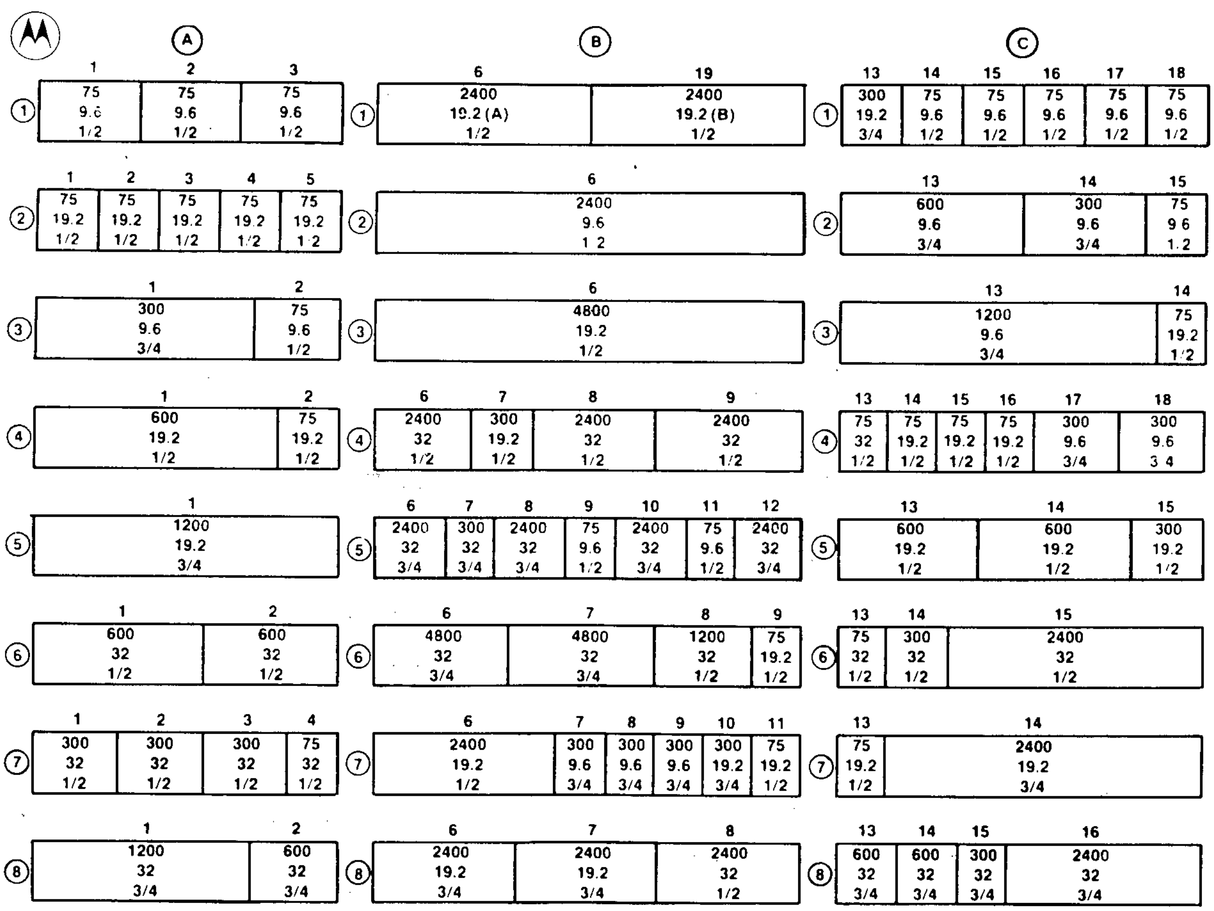

Fig. B-3—DAMA quick reference chart

On a given platform, up to four users may simultaneously utilize the DAMA subscriber unit. These users may be transmitting either data or digitized voice at the given input rates (up to 4.8 kbps). A typical example showing three platforms is illustrated in Fig. B-4.

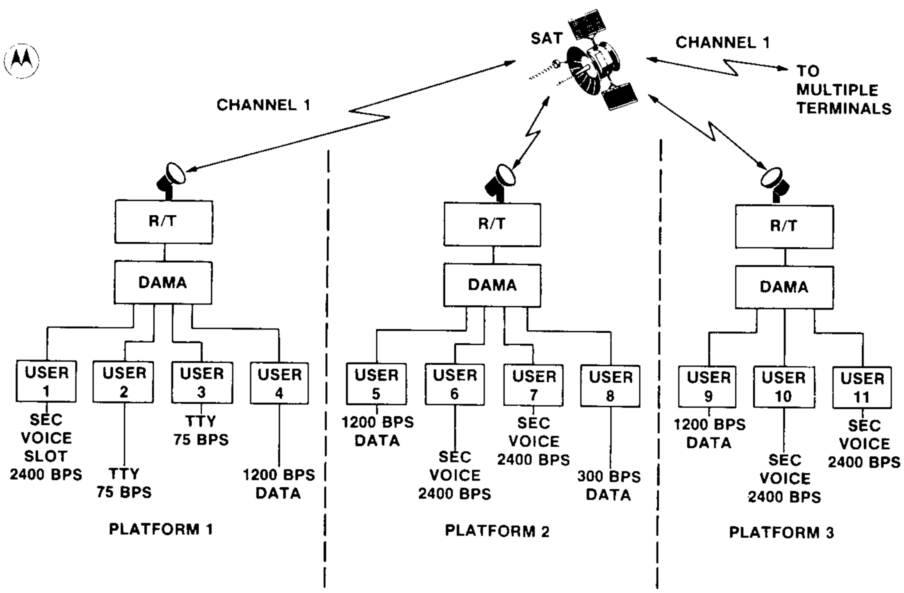

Fig. B-4—DAMA communications example—11 simultaneous users

Some consideration has been given to the problem incorporating TACINTEL and CUDIXS/NAVMACS subscribers into the UHF DAMA channels. A block diagram for an IXS configuration is shown in Fig. B-5. The DAMA subscriber unit has an ON-143 as its I/O interface and the 70 MHz DAMA output interfaces with the WSC-3, bypassing the OM-43 modem in the WSC-3 radio.

There are several difficult issues concerning the incorporation of TACINTEL and CUDIXS/NAVMACS into UHF DAMA. Since the IXSs are slow TDMA systems with multiple subscribers, there is a need to make their timing compatible with that of UHF DAMA, which is also a TDMA system. These interface (timing) problems are currently being given careful consideration. One possible solution involves the dedication of an individual DAMA channel to an individual IXS system (CUDIXS/NAVMACS or TACINTEL), in fact, this is the assumption that is used in describing the performance of DAMA-IXSs in this report. This would avoid the turn-around time delay involved in establishing and reestablishing a physical DAMA channel to an IXS user. However it detracts from the flexibility of DAMA because the loss of time slots to dedicated IXS channels reduces the capacity available for demand-assignment usage.

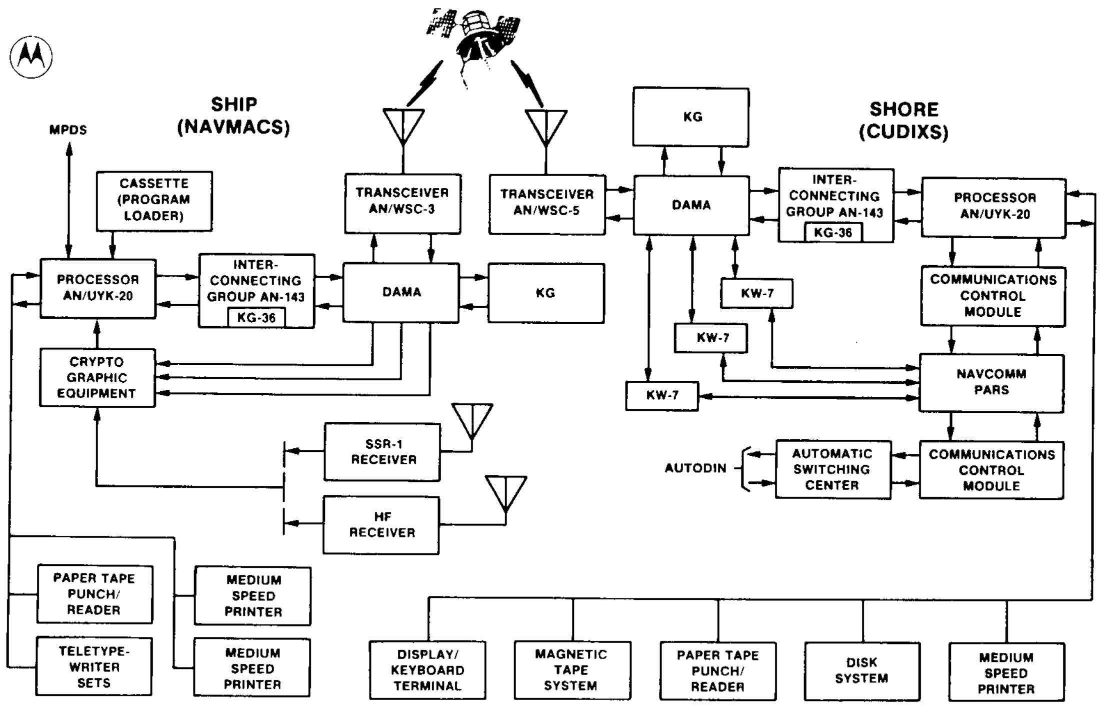

Fig. B-5—Demand assigned message traffic system (DAMIXS)

Appendix C
ERROR CORRECTION CODING FOR NAVY UHF SATCOM

With the use of increased data rates in UHF SATCOM it becomes necessary to employ error correction coding in order to maintain adequate margin in the link power budgets. In this connection there are two general situations of interest:

(1) For advantaged terminals (surface ships), there is a severe radio frequency interference (RFI) problem which can produce bursts of errors at all data rates.

(2) For disadvantaged terminals (submarines and aircraft), although RFI is seldom a problem, there is an inadequate margin in the link power budget due to insufficient received signal power at the higher data rates. Here again, coding gain is important in order to restore a usable link margin for reliable communications.

A recent study has addressed both aspects of the UHF SATCOM coding problem [C1,C2]. A breadboard model of a coder-decoder (codec) was built and comprehensively tested using the shipboard RFI simulator at the Naval Postgraduate School. The results of these tests serve to indicate the performance which is attainable in a combined Gaussian noise and shipboard RFI environment.

A convolutional coder was built having selectable rates of 1/2, 2/3, and 3/4, with constraint lengths k=9, 10, 11, respectively. (The 1/2 rate, k=9 coder is of principal interest in the present study.) These constraint lengths are longer than those commonly employed in Viterbi decoders, but it was found that the increased constraint lengths provide substantially better performance in the presence of RFI, and because of the relatively low operating speeds of interest, the longer codes are feasible. The principal results of the Monterey tests are given as follows:

- For disadvantaged users where RFI is not a factor, but link margin is a major concern, rate 1/2 coding can provide as much as 5 dB coding gain over uncoded operation in a Gaussian noise background at an error rate of 10^{-4}.
- Coding provides protection against bursts of erasures caused by periodic pulsed RFI blanking of theAN/WSC-3, and approximately 4 dB coding gain is still achievable.
- Even for the highest currently envisioned transmission rate of 19.2 ksps and the largest pulse width of 200 μs, for which erasure bursts are five symbols long, the use of interleaving does not provide significant performance improvement over the performance of coding along. That is, the code alone is powerful enough to provide essentially all the RFI protection. This conclusion applies even in the case of two independent RFI sources of 200 μs pulse width and 5% duty cycle. Thus, unless channel rates substantially in excess of 19.2 ksps or more severe RFI parameters are envisioned, use of pseudorandom interleaving is not warranted.
- In addition to protecting against erasure bursts resulting from high-power power pulsed RFI sources, coding can provide substantial performance enhancement over uncoded operation in a very dense RFI environment, such as it typically encountered aboard surface ships.

The overall conclusion from this study is that coding is an effective approach to alleviating the two most severe interference problems over UHF-SATCOM links—RFI and thermal noise. The Harris study cited here provides experimental evidence for this conclusion.

REFERENCES

C1 Coding for Navy UHF Satellite Communications (Phase I), Final Report under contract N00039-76-C-0384, Harris Corporation, 1 April 1977.

C2 Coding for Navy UHF Satellite Communications (Phase II), Final report under contract N00039-77-C-0259, Harris Corporation, 11 Aug 1978.

Appendix D
THE AN/WSC-3 TRANSCEIVER AND ITS MODIFICATIONS

The WSC-3 UHF modular radio is used in all FLTSATCOM subsystems except Fleet Broadcast and ASWIXS. It is an all-solid-state radio in the 225 to 400 MHz band with 7000 channels spaced at 25 kHz increments. The WSC-3 in different configurations can be used for either SATCOM or line-of-sight (LOS) applications. In the SATCOM mode, four modulation choices are available (PSK, FSK, FM, and AM) and the transmit and receive frequencies are offset by a fixed amount. By removing several modules, a less expensive AM LOS version of the WSC-3 can be built. This LOS WSC-3 can be converted to a satellite WSC-3 by installation of the appropriate modules.

The WSC-3 is packaged for conventional 19-in. rack mounting. The number of WSC-3s stacked vertically in a rack can vary from one to four. Most installations have a single channel capacity, but others combine two, three, or four WSC-3s, as on large ships.

The maximum RF power transmitted by a WSC-3 is 100 W for PSK, FSK, and FM; and 30 W for AM. The transceiver is presently configured with a single synthesizer, and operates in a half-duplex mode. The turnaround time for transmit/receive is 7 s (maximum). The front end of the receiver is protected by a T/R switch which keeps the transmitter from being inadvertently connected to the receiver. In addition there is a blanker in the receiver IF which blanks the main IF in the event of high-level pulsed interference (as from shipboard radars).

Internal to the WSC-3 is a modem which is the equivalent of the OM-43 modem. The WSC-3 interfaces at baseband to incoming data rates of 75, 300, 1200, 2400, 4800, and 9600 symbols per second. For satellite IXS application BPSK modulation (differentially encoded) is used, and error rates are within 2 dB of theoretical limits. An additional modem using FSK at 75 symbols per second is also provided internal to the WSC-3.

At shore stations, the AN/WSC-5 transceiver is used to provide (up to) an eight-channel capability. This transceiver is equivalent to a WSC-3, with the exception that the OM-43 modem is external to the transceiver. All Naval Communications Area Master Stations are or will be equipped with WSC-5 transceiver installations. Typically, an eight-channel capability will require four full racks of equipment, including WSC-5s, OM-43s, and control equipment.

In addition to the WSC-3, there has been developed an auxiliary WSC-3 receiver. This receive-only (R/O) unit is identical in performance (as a receiver) to the WSC-3. The WSC-3 (R/O) is smaller and less expensive than the WSC-3, but uses common WSC-3 modules for the synthesizer and demodulator/transmitter chain and by using a power supply and self-test package with lesser requirements than the full WSC-3. The simultaneous use of a WSC-3 as a transmitter and a WSC-3 (R/O) as a receiver permits a full-duplex T/R capability, and this combination has particular importance to this study.

The WSC-3 radio can also be modified to increase the maximum transmission rate to 19.2 ksps. This can be accomplished easily and economically, and can be made switch-selectable with the 19.2 ksps transmission rate being put in the place of the infrequently used 300 bps switch position.

Both the R/O and the 19.2 ksps modifications have been performed in connection with the Trident communications program. As such, it is feasible to carry forward this implementation in the CPODA consolidated approach discussed in this report.

Finally, it should be mentioned that modifications to the WSC-3 will also be required in the UHF DAMA program. Using the WSC-3 (with a single synthesizer) in a half-duplex mode requires that the transceiver be rapidly turned around in order to transmit and receive adjacent DAMA slots. As presently configured, the settling time of the synthesizer is not adequate to meet these requirements and synthesizer modifications will have to be performed in order to overcome this difficulty.

Appendix E
THE UHF SATCOM TRANSMISSION CHANNEL: POWER AND BANDWIDTH UTILIZATION

Network communications problems include several aspects—demand generation, queue management, and information transmission. The purpose of this appendix is to summarize certain aspects of the transmission problem—understanding the physical environment so as to utilize the available power and bandwidth in an efficient manner. From a transmission viewpoint, the characteristics of interest are power, bandwidth, and the channel parameters. These ultimately determine the data rate, bit-error-rate, and error statistics. In turn, these lower level performance measures contribute toward the higher level performance measures such as message delay, throughput, and character-error rate.

The Link Power Budget

For existing FLTSAT channels under normal conditions, the uplink power is sufficient to capture the hard-limiting amplifier in the satellite, and the overall channel performance is limited by the noise and interference at the downlink terminal (*i.e.*, a downlink limited channel). The downlink power budget for a 2.4 kbps transmission to a shipboard terminal with 26 dBW satellite effective radiated power (ERP) is given in Table E-1.

Table E-1—Downlink Power Budget for a Shipboard Terminal

ERP	26	dBW
Path Loss	−173	dB
Receive Signal Power	−147	dB
Receive Antenna Gain	10	dB
Noise Temperature (1200°K)	30.8	dB°K
kT	−198	dBW/Hz
C/kT (available)	61	dB-Hz
E_b/N_o (at BER = 10^{-5})	13	dB
Date rate (2.4 kbps)	33.8	dB-Hz
C/kT (required)	48.8	dB-Hz
Link Margin	14.2	dB

In Table E-1 the receiver noise temperature is derived from noise contributions from the receive antenna (OE-82), a low-noise pre-amp, and low level shipboard radio frequency interference (RFI). High level pulsed noise contributions, such as from shipboard radar, are excluded from this link budget. (To deal with radar RFI, additional coding protection is required, as discussed in Appendix C.) With a reasonably white Gaussian noise background, the required energy per bit to noise power density ratio, E_b/N_o, for 10^{-5} bit-error-rate is chosen as 13 dB. This assumes differentially encoded (coherent) BPSK with an appropriate implementation loss.

The link margin for this 2.4 kbps transmission is the difference between the available and required C/kT's (carrier power to noise density ratios), in this case 14.2 dB. On the other hand, the link margin for a disadvantaged user is nearly 10 dB less than for the shipboard terminal. The reason for this is that submarine and airborne terminals have nondirectional receive antennas. To compensate somewhat, disadvantaged users generally exist in a less severe RF environment so that a small portion of the link margin loss is regained in the form of a lower receiver noise temperature. Conversely, for a COMMSTA receiver site, the receive antenna gain is higher than shipboard (by about dB) but the RFI environment is far more severe. On balance, the COMMSTA enjoys a slightly greater link margin than the shipboard terminals (by a few dB).

Spectral Utilization

The presently configured 25 kHz FLTSAT channels utilize coherent BPSK (differentially encoded) modulation. This modulation possesses a power spectrum which is proportional to

$$\frac{\sin^2\left[\dfrac{\pi(f-f_o)}{r_D}\right]}{\left[\dfrac{\pi(f-f_o)}{r_D}\right]^2},$$

where f_o is the (suppressed) carrier frequency (Hz), and r_D is the data rate (bps).

For data rate equal to 2.4 kbps (uncoded), as exists in the present CUDIXS system, the power spectrum occupies a relatively small portion of the allocated 25 kHz channel. This is shown in Fig. E-1. The main lobe, residing within $f_o \pm 2.4$ kHz, contains slightly more than 90% of the transmitted power. Nearly 80% of the total transmitted power is confined to within $f_o \pm 1.2$ kHz, that is, a bandwidth equal to the data rate and centered at f_o.

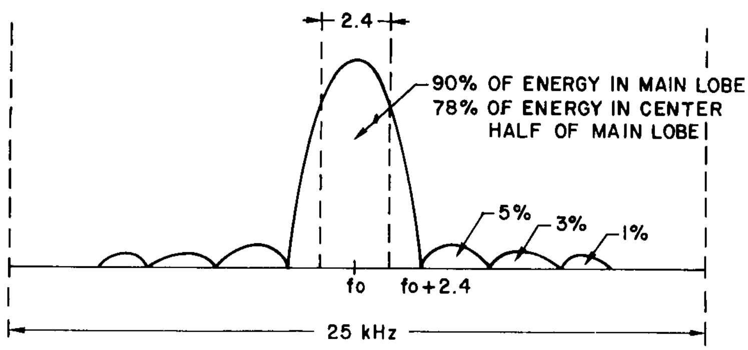

Fig. E-1—Spectrum of a 2.4 kbps (uncoded) BPSK signal

It is apparent that the 25 kHz FLTSAT channels, as present configured, are underutilized from a bandwidth point of view. The WSC-3 radio has switch-selectable transmission rates which go up to 9.6 ksps and can be modified to accommodate 19.2 ksps. At 19.2 ksps, more than 80% of the transmitted energy still resides within the 25 kHz channel.

Conclusions

Existing 25 kHz FLTSAT channels are not efficient in their use of bandwidth and power. With the use of half-rate error correction codes it is possible to incorporate several of the present FLTSATCOM users into a single channel with an overall data rate of 9.6 ksps and coded symbol transmission rate of 19.2 ksps. Such a transmission could be handled by the WSC-3 radio (modified to accommodate 19.2 ksps).

At a transmission rate of 19.2 ksps (corresponding to a data rate of 9.6 ksps, one-half rate coded), the link power budget of Table E-1 would be revised to that given in Table E-2. Note that the increased data rate reduces the received C/kT by 6 dB, but that a 4 dB coding gain partially offsets this loss. Also it is noted that the link power budget is applicable even in the presence of radar pulse RFI, because the coding protection extends over pulsed interference.

Table E-2—Downlink Power Budget for a Shipboard Terminal
(at 9.6 kbps using rate-1/2 coding)

ERP	26	dBW
Path Loss	−173	dB
Receive Signal Power	−147	dBW
Receive Antenna Gain	10	dB
Noise Temp (1200°K)	30.8	dB°K
kT	−198	dBW/Hz
C/kT (available)	61	dB-Hz
E_b/N_o (at BER = 10^{-5}; includes 4 dB coding gain)	9	dB
data rate (9600 bps)	39.8	dB-Hz
C/kT (required)	48.8	dB-Hz
Link Margin	12.2	dB

In summary, 9.6 kbps (coded to a transmission rate of 19.2 ksps) can be accommodated within a 25 kHz channel and with realistic downlink ERPs can provide a link power budget with adequate link margin.* Accordingly it is viewed as a viable transmission rate for the present study.

*UHF DAMA has been operated at this higher data rate during its recent TECHEVAL. Results are not available for inclusion in this report.