

AD-A075 369

NAVAL RESEARCH LAB WASHINGTON DC

F/G 17/2.1

PERFORMANCE STUDIES OF THE DISTRIBUTED CPoDA PROTOCOL IN THE MO--ETC(U)


UNCLASSIFIED

NRL-MR-4084

SBIE-AD-E000 328

NL

| OF |
ADA
075369



END
DATE
FILMED
11 -79
DDC

12 LEVEL III

ade 000328

NRL Memorandum Report 4084

AD A 075369

Performance Studies of the Distributed CPODA Protocol in the Mobile Access Terminal Network

CARL E. LANDWEHR

Information Systems Staff

September 28, 1979

DDC
RECEIVED
OCT 10 1979
B



DDC FILE COPY

NAVAL RESEARCH LABORATORY
Washington, D.C.

Approved for public release; distribution unlimited.

79 09 26 008

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER NRL Memorandum Report 4084	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) <u>PERFORMANCE STUDIES OF THE DISTRIBUTED CPODA PROTOCOL IN THE MOBILE ACCESS TERMINAL NETWORK</u>	5. TYPE OF REPORT & PERIOD COVERED Interim report on a continuing NRL problem	6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) <u>Carl E. Landwehr</u>	8. CONTRACT OR GRANT NUMBER(s) <u>11 28 Sep 79</u>	
9. PERFORMING ORGANIZATION NAME AND ADDRESS Naval Research Laboratory Washington, DC 20375	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS NRL Problem 75R11-09.201 28010N, X-0701-CC	
11. CONTROLLING OFFICE NAME AND ADDRESS Naval Electronic Systems Command ELEX 310 Washington, DC 20360	12. REPORT DATE September 28, 1979	
13. NUMBER OF PAGES 41	14. SECURITY CLASS. (of this report) UNCLASSIFIED	
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) <u>14 NRL-MR-4084</u>	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE	
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report) DDC RECEIVED OCT 10 1979 B		
18. SUPPLEMENTARY NOTES <u>18 SBIE 19 AD-E000 328</u>		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Packet switching Packet protocol Demand access Packet reservation CPODA Contention Broadcast channel Simulation Satellite communication Mobile Access Terminal		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) A Contention Based, Priority Oriented, Demand Access (CPODA) protocol has been proposed for use in a network of shipboard computers and terminals designed to provide access to the Advanced Command and Control Architectural Testbed (ACCAT) via ARPANET. A model of the protocol was implemented on NRL's Satellite Communication Simulator and performance studies were conducted. Although delays in the network are likely to exceed levels tolerable in a shore-based communications network for interactive computer use, they may be acceptable for the purposes of the project.		

DD FORM 1 JAN 73 1473

EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102-014-6601

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

79 09 26 008

Table of Contents

<u>Section</u>	<u>Page</u>
1. Introduction	1
2. MAT Environment	1
3. Distributed CPODA Protocol	4
4. The CPODA Model	7
5. Traffic Model	10
6. Determination of Simulation Parameter Values	12
7. Simulation Studies	15
8. Conclusions and Recommendations	36
9. References	39

ACCESSION for	
NTIS	White Section <input checked="" type="checkbox"/>
DDC	Buff Section <input type="checkbox"/>
UNANNOUNCED	<input type="checkbox"/>
JUSTIFICATION _____	
BY _____	
DISTRIBUTION/AVAILABILITY CODES	
Dist.	AVAIL. and/or SPECIAL
A	

PERFORMANCE STUDIES OF THE DISTRIBUTED CPODA PROTOCOL
IN THE MOBILE ACCESS TERMINAL NETWORK

1. Introduction

This report documents an investigation of the performance of the Contention-based, Priority Oriented Demand Assignment (CPODA) protocol (1). The protocol is evaluated in the environment planned for the Mobile Access Terminal (MAT), which is being developed by the Naval Electronic Systems Command (NAVELEX) to provide secure shipboard access to the Advanced Command and Control Architectural Testbed (ACCAT).

The structure of the communication system between the shipboard terminal and the ACCAT is shown in Figure 1-1, taken from [(2)]. The top of the figure displays the general structure of the network, and the bottom is a detailed block diagram for a ship to shore link. The CPODA protocol is proposed for use on the satellite link between the ships and the shore station.

The purpose of this investigation is to assess the probable delays that will occur to messages transmitted on the satellite link operated with the distributed CPODA protocol. Only queueing and transmission delays for the satellite links are considered; delays in moving characters between the Private Line Interface (PLI), and ARPANET components are not considered.

Section 2 of this report provides a description of the MAT environment. The CPODA protocol and the model of it implemented in the simulation are described in Sections 3 and 4, respectively. Section 5 discusses the assumptions made concerning traffic in the MAT network. The final three sections describe the actual parameters used in the simulation runs, the results observed, and conclusions and recommendations based on the study.

It is recommended that the MAT implementation proceed essentially as proposed in [(3)], but with special attention to (1) inclusion of the backup FTDMA mode to assure reliable ACCAT access and (2) inclusion of sufficient software to enable realistic experiments to be conducted on communication protocols. The software should record and report at least the set of performance parameters enumerated in Section 8.2, and should provide information to users on the current state of network queues and delays. It should also provide for the generation of suitable dummy traffic loads.

2. MAT Environment

This section describes the planned environment for MAT operations. Further details are available in [(2), (3), (4)]. The primary goal of the MAT is to provide a means for experimental systems implemented on the ACCAT to be demonstrated and tested in an actual shipboard environment. In order to allow the MAT to be used from as many different ships as possible, the hardware required is to be constructed as a small number of portable units that may be carried on board a ship and clamped in place for the duration of a test or

Note: Manuscript submitted July 23, 1979.

PERFORMANCE STUDIES OF THE DISTRIBUTED CPDPA PROTOCOL
IN THE MOBILE ACCESS TERMINAL NETWORK

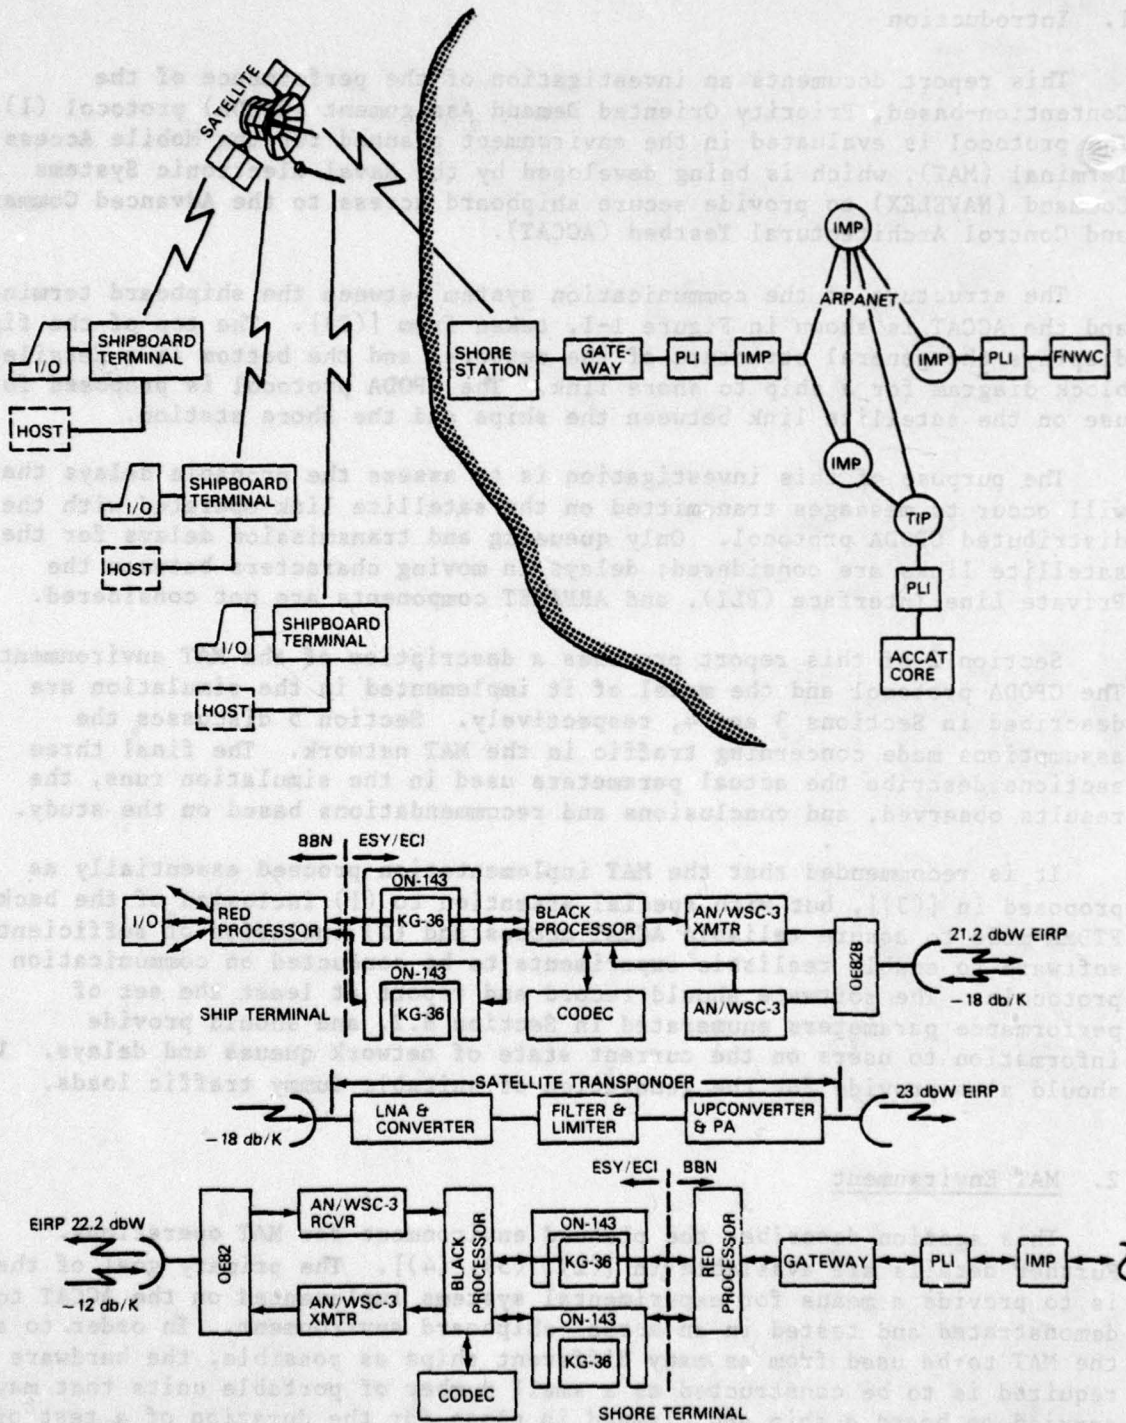


Fig. 1-1 - MAT concept and block diagram

experiment. The present design [(2), (3)] requires the use of two WSC-3 radios on the ship, both sharing the same satellite channel. Two radios are needed to provide full duplex access to the channel. The shore station requires a complementary configuration.

2.1 Equipment

The equipment that is being developed and acquired for the MAT provides a secure access for one or more computer terminals per ship to the satellite link and controls the sharing of the satellite link among several ships. This equipment includes:

1. A processor (or perhaps two microprocessors) to handle the computer terminals, provide support for ARPANET protocols, and control the satellite link using the CPODA (or other) access protocol. This processor is also to furnish a port to which a shipboard ARPANET host computer could be connected. This part of the system is referred to as the red processor, since it handles unencrypted data.
2. An interface unit between the red processor and the crypto equipment.
3. The crypto equipment (2 KG-36's, for full duplex data flow) and an ON-143 interconnection group. (Cryptos already on board the ship may be used.)
4. A processor to control the data flow between the cryptos, the coder/decoder, and the WSC-3's (called the black processor, because it handles the encrypted data).
5. A coder/decoder (codec) with an interleaver to provide error detection and correction on the satellite link.
6. A minor modification package for the WSC-3's, based on TRIDENT developments, to allow a transmission rate of 19.2 Kilo-symbols-per-second (Ksps).

The ship and shore MAT configurations are symmetric, except that at the shore stations, the red processor is connected to an ARPANET gateway processor that provides internetwork protocol handling between the MAT satellite network and the ARPANET. The gateway is connected in turn to an ARPANET Private Line Interface (PLI) that provides secure access to the ARPANET. A complementary PLI is placed between the ACCAT and its local ARPANET Terminal Interface Processor (TIP) to complete the chain from the shipboard terminal to the ACCAT.

2.2 Traffic

In the planned mode of operation, a MAT user on board a ship will access ACCAT computer resources interactively. Thus, traffic over the satellite link is expected to resemble the traffic between users and computers in land-based interactive systems. Previous studies [(5), (6)] indicate that such traffic tends to be bursty, with short messages (5-15 characters) sent from the user

to the computer and longer messages (50-150 characters) from the computer to the user. Shipboard users may be expected to be as impatient as any, so minimizing delays on the satellite link is an important consideration.

Graphics terminals might be expected to receive longer transmissions from the computer than the typical video display terminal receives. File transfers would also cause longer individual transmissions, but should not require as low delay times as graphics or interactive traffic. Both of these types of support have been discussed for MAT, but both would require a shipboard host processor. Although the MAT design includes a host port on the red processor, there are no plans to implement such a host as yet.

2.3 Satellite Channel Access Protocols

Three different satellite channel access protocols are proposed for MAT in [(3)]: fixed assignments, centralized CPODA, and distributed CPODA. Fixed assignments is simply Fixed Time-Division Multiple Access (FTDMA) with a leader station providing timing and user stations transmitting data once per frame. This protocol is intended for use only for hardware checkout and channel testing.

The distributed CPODA protocol is described in detail in the following section. It is a flexible, packet-based reservation protocol that includes contention and has controls for stability. Centralized CPODA is essentially the same as distributed CPODA, except that there is a controlling station that listens to the reservations and broadcasts frame allocations to some or all users. In distributed CPODA, each station listens to all reservations and schedules data transmissions, so no channel time is required for the transmission of allocations. Because it is intended to be the primary protocol used by the MAT network, distributed CPODA was chosen for further study.

3. Distributed CPODA Protocol

Distributed CPODA can be viewed as a particular member of the family of PODA protocols. Descriptions of CPODA have appeared in [(1), (3), (8), and (10)]. This section describes the CPODA protocol as it is presently planned for use in the MAT, based on the descriptions in [(1) and (3)], and on conversations with representatives of BBN. A non-contention version of a PODA protocol (Fixed PODA, or FPODA), as well as CPODA, is discussed in [(10)]. FPODA employs a fixed assignment of users to control subframe slots instead of the contention-based assignment described below. Reservation synchronization algorithms for CPODA are discussed in [(8)].

3.1 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 (see Fig. 3-1). 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.

3.2 Reservations

CPODA is a reservation-based protocol: a subscriber with traffic to send transmits a reservation indicating its identity and its traffic characteristics. The reservation is queued by all network participants, and when the traffic corresponding to pre-existing and higher priority reservations has been transmitted, the subscriber transmits (in an information subframe) the data corresponding to its reservation.

In distributed CPODA 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 transmissions on the channel and predict which subscriber will transmit next.

Although not planned for implementation in the MAT, [(1) and (10)] also define a "stream" reservation. A single stream reservation requests that all users periodically enter a reservation into the queue without any explicit reservation request. The motivation for this type of reservation is to provide service to traffic demands that are inherently periodic (e.g., voice traffic).

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 contention subframe shrinks, fewer subscribers need to use the contention subframe.

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.


3.3 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.

Each frame is divided into two subframes: the information subframe and the contention subframe (see Fig. 3-1). 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.

3.1 Reservations

CPODA is a reservation-based protocol; a subscriber with traffic to send transmits a reservation indicating its frame number and its traffic characteristics. The reservation is queued by all network participants, and when the traffic corresponding to the reservation is ready to be transmitted, the subscriber transmits its information subframe (the data corresponding to its reservation).



- L = LEADER PACKET
- I = INFORMATION PACKET
- C = CONTENTION PACKET
- T_I = INFORMATION SUBFRAME
- T_C = CONTENTION SUBFRAME

Fig. 3-1 - CPODA frame structure

Message acknowledgments, when required, may be sent either by piggybacking them on a data transmission or by sending them in the contention subframe. The format method is preferred if the subscriber already has a reservation in the queue. An acknowledgment generally contains a frame number/frame number and is 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.

3.2 Contention and Retries

Contention problems are handled as follows: The station transmitting a packet in the contention subframe monitors the channel to hear its echo. If the echo is heard correctly, the subscriber assumes that the reservation for acknowledgment has been heard by the other network participants. 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 pseudo-random 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.

3.4 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 downlink 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 re-enters 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 a few frames 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, subscribers wishing to receive traffic must listen anyway.)

4. The CPODA Model

In any simulation, some model of the system to be simulated must be defined. Generally, 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!). This section describes how the model for CPODA used in the simulations differs from CPODA itself.

4.1 Frame Structure

The frame defined for CPODA has the form shown in Figure 3-1. The simulation uses the identical frame structure, except that the leader packet precedes the information subframe (see Figure 4-1). This is more convenient for the simulator, since, during the leader slot, the simulator can schedule the stations that are to transmit during the information subframe based on the reservations just received in the preceding contention subframe. The remainder of the current subframe is then automatically devoted to contention. (Notice that, if the leader is small relative to the frame length, the two schemes 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 by priority (lowest numbered priority first); the frame is filled completely with information slots subject to the constraint that there be at least two contention slots per frame. For most of the tests run, the shore-originated traffic is assigned priority one and all ship-originated messages are assigned priority two.

4.2 Reservations and Acknowledgments

Reservations are handled in the simulator essentially as described in Section 3.2 above. Each contention packet may contain up to two reservations and four acknowledgments. In the proposed MAT system such packets may contain two reservations or several acknowledgments. The actual number of acknowledgments and reservations 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.

The protocol implements an ARQ mechanism; in the simulation, one acknowledgment is required for each data block transmitted. An information packet may contain from one to sixteen 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.

4.3 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 ack). If the node already has reservations pending in the queue, it can piggyback the new reservation (or ack) on its next transmission. In the simulator, this choice is made on the basis of the number of acks and data blocks to be sent, the number of reservations pending for this node in the queue, and the expected waiting time until this first outstanding reservation for this 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 echo and reservation from the packet are requested and the contention threshold is halved, to reduce the probability of contention packet retransmission. If the echo is received successfully, the reservation are deleted into the global reservation queue. The contention threshold is adjusted (the distance between the threshold and 1 is halved) to increase the probability of contention packet reservation, and the echo timeout is cancelled.

4.4 Reservation Start-up

It is in this part of the protocol that the simulated CPODA is most departed from the standard CSMA reservation. To include the synchronization status in detail, would require keeping track of the predictors made by each node for each reservation and transmission. The complexity and overhead that would be incurred would outweigh the present simulator and would not in any way seem to justify the minor increase in precision.




Fig. 4-1 - Simulated CPODA frame structure

The model for interactive traffic is included in the MAI Intern Report [1]. This model, which is already in adopted form, is based on a reservation system. The reservation system is also used in the Intern Report. The investigation of system behavior under such a model is unlikely to handle such traffic in any case, since an algorithm based computer is planned as well.

Table 5-1 displays the differences between the Intern Report model and the one developed here. The arrival rates are identical in both models and are based on the assumption that system response times and human factors considerations will constrain users to an average request rate of three per minute. Since each request from a shipboard user generates a response from the shore-based system, the traffic generation rate at the shore station is equal to the sum of all the ship-based traffic generation rates.

The message lengths assumed in the MAI Intern Report have been revised slightly. The motivation for the Intern Report was message lengths parameters include assumptions about user typing speed (characters/second).

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 acks and 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 (the distance between the threshold and 1 is halved) to increase the probability of contention packet transmission, and the echo timeout is cancelled.

4.4 Subscriber Start-up

It is in this part of the protocol that the simulated CPODA is most abstracted from the planned 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 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 modelled 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 modelled similarly. While a station is out of sync, it does not transmit any messages or acknowledgements. All stations are assumed to start in reservation synchronization, and the initial acquisition state is omitted.

5. Traffic Model

A model for interactive traffic is included in the MAT Interim Report [(4)]. That model, slightly revised, is adopted here. Some assumptions concerning file transfer traffic are also made in the Interim Report; limitations of time and funds have prevented the refinement of that model and the investigation of system behavior under such loads. The initial MAT is unlikely to handle such traffic in any case, since no shipboard host computer is planned as yet.

Table 5-1 displays the differences between the Interim Report model and the one developed here. The arrival rates are identical in both models and are based on the assumption that system response times and human factors considerations will constrain users to an average request rate of three per minute. Since each request from a shipboard user generates a response from the shore-based system, the traffic generation rate at the shore station is equal to the sum of all the ship-based traffic generation rates.

The message lengths assumed in the MAT Interim Report have been revised slightly. The motivation for the Interim Report user message length parameters includes assumptions about user typing speed (5 characters/second),

	CPODA Simulation	MAT Interim Report [(4)]																
Distribution of message inter-arrival times for ship-shore messages at each user terminal:	negative exponential, mean = 20 sec.	Same																
Distribution of ship-shore message lengths:	Tabular distribution, mean = 11 characters, variance = 3.1: <table border="1"> <thead> <tr> <th>x</th> <th>cdf(x)</th> </tr> </thead> <tbody> <tr><td>0</td><td>0.0</td></tr> <tr><td>5</td><td>.3740</td></tr> <tr><td>10</td><td>.4360</td></tr> <tr><td>15</td><td>.9271</td></tr> <tr><td>20</td><td>.9987</td></tr> <tr><td>25</td><td>.9997</td></tr> <tr><td>35</td><td>1.0</td></tr> </tbody> </table>	x	cdf(x)	0	0.0	5	.3740	10	.4360	15	.9271	20	.9987	25	.9997	35	1.0	Uniform distribution from 1 to 73 characters
x	cdf(x)																	
0	0.0																	
5	.3740																	
10	.4360																	
15	.9271																	
20	.9987																	
25	.9997																	
35	1.0																	
Distribution of message inter-arrival times of shore-ship messages, for a network with N user terminals:	Negative exponential, mean = $\frac{N}{20}$ sec.	Same																
Distribution of shore-ship message lengths:	Tabular distribution, mean = 110 characters, variance = 31. : <table border="1"> <thead> <tr> <th>x</th> <th>cdf(x)</th> </tr> </thead> <tbody> <tr><td>0</td><td>0.0</td></tr> <tr><td>50</td><td>.3740</td></tr> <tr><td>100</td><td>.4360</td></tr> <tr><td>150</td><td>.9271</td></tr> <tr><td>200</td><td>.9987</td></tr> <tr><td>250</td><td>.9997</td></tr> <tr><td>350</td><td>1.0</td></tr> </tbody> </table>	x	cdf(x)	0	0.0	50	.3740	100	.4360	150	.9271	200	.9987	250	.9997	350	1.0	Uniform distribution, 1 to 360 characters
x	cdf(x)																	
0	0.0																	
50	.3740																	
100	.4360																	
150	.9271																	
200	.9987																	
250	.9997																	
350	1.0																	
Number of addressees per message (all ship/shore messages addressed to shore station; addresses of shore/ship messages distributed uniformly over all ships)	1	Same																

Table 5-1 Traffic Model Parameters

think times (10 seconds), and system response time (3 seconds). The upper bound of 73 characters was based on the interleaver length, and the uniform distribution was apparently chosen arbitrarily. Shore message lengths were derived by applying an assumed 5:1 multiplication factor, and by limiting packets to 2048 bits (=256 characters).

We prefer to base our assumptions for message lengths on the traffic measured in actual interactive systems and presented in [(5)] and [(6)]. These studies include measurements from four different interactive systems. The mean and variance we assume for the user message length are derived from the "number of bursts per user burst segment" measures in [(6)]. The tabular distribution was generated by a program [(7)] that applies entropy maximization techniques to determine a distribution with a specified mean and variance. The authors of [(5)] observe that the average number of characters sent from computer to user is generally an order of magnitude greater than flow in the opposite direction. Hence, we have merely scaled the user message length distribution by a factor of 10 to get an approximate distribution for the length of computer-generated traffic.

6. Determination of Simulation Parameter Values

The parameters actually used to operate the NRL Satellite Communication Simulator (SCS) with the CPODA protocol are displayed in Table 6-1. Some of the parameters, such as the number of ships, the message arrival rates, and the CPODA frame length, were varied during the studies, but most of them were held constant. This section briefly describes how the parameters were determined.

6.1 Message Generation Parameters

The inter-arrival and length distributions for messages were discussed in Section 5. To minimize queues in the shore processor, all shore to ship messages are defined to be priority one, and all ship to shore messages are priority two. These priorities control the CPODA reservation queues and the information slot allocation algorithm. Node 1 in the simulation represents the shore station; all messages generated by other nodes are sent to node 1. Messages generated by node 1 are addressed randomly (with a uniform distribution) among all other nodes. The number of nodes and the arrival rate per node were varied to explore system behavior under various channel loadings.

6.2 Equipment Related Parameters

These parameters are derived primarily from the data in the ECI proposals [(2), (4)]. Since ranging delays are not expected to be a problem, the simulator assumes that all ships have a round-trip propagation delay of .25 seconds on the satellite channel. The maximum propagation delay in any case would not exceed .28 seconds; additional simulation runs may be performed with the higher value if desired.

FEC coding rate .5
 Modem preamble length (FEC does not apply) 208 bits
 Crypto preamble length (before FEC) 64 bits
 Guard bits between user transmissions 20 bits
 Channel transmission rate 19.2 K sps
 (=9.6 K information bits/sec.)

CPODA Frame Format

Frame length .5 sec.
 Minimum number of contention slots per frame 2

Leader Packet

Length of data 100 bits
 Total packet length (with sync and FEC) 536 bits

Contention Packet

Length of data 100 bits
 Maximum number of reservations/packet 2
 Maximum number of acknowledgments/packet 4
 Total packet length (with sync and FEC) 536 bits

Information Packet

Length of overhead data (incl. ack/resv.) 280 bits
 Length of data per information block 225 bits
 Maximum number of information blocks/packet 16
 Minimum total packet length (1 data block, with sync and FEC) 1366 bits
 Maximum total packet length (16 data blocks, with sync and FEC) 8096 bits

Timeouts

Maximum tolerable time to wait for a piggybacking opportunity 30 sec.
 Block timeout (time to wait before retransmitting an unacknowledged block) 31.5 sec.

Table 6-1 CPODA Simulation Parameters

Modem turnaround time is assumed zero, since MAT operates as a full duplex system. Modem synchronization bits include the modem preamble and the unique word; the crypto preamble assumes the KG-36 key is compressed and half-rate encoded. Both of these assumptions, and the assumption of 20 bits for guard time between user transmissions as well, are taken directly from [(2)]. The channel transmission rate of 19.2Ksps and the codec operation at half rate are based on the same source.

The projected rate of retransmissions required due to detected transmission errors (other than collisions) is difficult to estimate. Consequently, a baseline of zero transmission errors to assess the relative performance of the system under different traffic loads has been used. Some additional runs with non-zero retransmission probabilities have been made (see Section 7) to assess the basic stability of the performance of the protocol.

6.3 Protocol Related Parameters

The fundamental parameters for CPODA are the frame length, the length of a data block (the minimum number of data bits in an information packet), the maximum number of data blocks in an information packet, and the length of a contention packet. The frame length was varied in a number of test runs, (see Section 7) and the value of 0.5 seconds per frame provided a reasonable balance between reduced delay in a lightly loaded system and increased overhead under heavy load. The present frame length used in the CPODA experiments on SATNET (successor to DARPA's Atlantic Packet Satellite Experiments) is approximately .25 seconds. SATNET employs a 64Ksps channel, so a longer frame length chosen for the lower rate MAT channel seems appropriate. Reduction of the frame length to below one round-trip propagation delay time might require some minor revisions to the CPODA scheduling algorithms.

The packet lengths are based on the MAT definitions in [(2)]. Recent information on the codec evaluation indicates that interleaving may not be necessary, particularly when there is no RFI. In this case, fill bits would not be required to pad packets to the interleaver length, allowing slightly shorter control packets and a more uniform number of additional information bits per additional unit of information packet length. The bit counts given in the table do not include the half rate coding or crypto and modem sync bits; these overheads are added automatically by the simulator.

The contention packet data length of 100 bits is based on an assumption of roughly 50 bits for each of two reservations. An information packet is assumed to include a contention packet plus additional overhead of 180 bits (corresponding to Transmission Control Protocol (TCP) and inter-network headers) and from one to 16 data blocks of 225 bits each. The block size was determined by dividing the number of data bits in a MAT "16 data packet" [(2)] by sixteen and rounding to an even 5 bits. The individual packet check sums have been neglected. A leader packet is assumed to be the same length as a contention packet.

7. Simulation Studies

More than thirty separate simulation runs have been performed, each representing an hour of simulated activity over the MAT channel. Before we present the results of these experiments, we describe the structure of the studies and some limitations on the results.

The studies began with several runs to establish an appropriate value for the CPODA frame length. Once this parameter was determined, studies of the message delay as a function of the number of ships and the number of terminals per ship commenced. Finally, a few studies of the effect of channel errors and the performance of a single priority system were performed.

Limitations on the results arise from two sources: the inherent nature of simulation and the practical restrictions of time and money on the number and length of the experiments performed. The nature of simulation is such that the results are generally in the nature of a feasibility demonstration, not a guarantee of performance. A result observed in a given simulation run may be achieved in the real world if the assumptions about the environment and the protocol embodied in the computer program are correct and if the particular stream of events generated during the simulation (or in the real world) is not, by chance, highly unusual. This type of limitation leads us to be explicit about our assumptions (see Sections 2-6) and to interpret measures from the simulation only as statistical measures, not as precise predictions. Relative measures from simulations ("situation A has less delay than situation B") are generally more reliable than absolute measures ("the mean response time in situation A is 8.263 seconds").

The limitations of time and money have restricted the number of protocols investigated in this project to one, CPODA, and also the number of experiments with that protocol. For this reason, multiple replications of the same experiment with different random number streams have been performed only in a few cases of primary interest. To preserve comparability among other cases, a standard random number stream has been used for message generation. (A single random number stream corresponds to a single pattern of message arrivals over a one hour period.)

In the first set of studies, for example, different values for the frame length were tested against the identical sets of message arrivals. In another series of runs, for a baseline configuration of six ships, two terminals per ship, and no errors, multiple seed sets were used. This approach allows reasonable confidence in relative results for all runs with the "standard" message arrival stream, and also a reasonable assessment of absolute performance in the experiments where multiple seed sets were used in the baseline environment.

7.1 Experimental Baseline

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.

	<u>Case 1</u>	<u>Case 2</u>
Probability of modem synchronization	1.0	.999
Probability of crypto synchronization	1.0	.999
Probability of error in header	.0001	.0001
Probability of error in data block	.0025	.01
Mean time to loss of reservation synchronization (sec.)	100,000	180.
Mean time to regain reservation synchronization (sec.)	.0001	2.0

Table 7-1 Parameters for Error Environment

- A. All runs are made for one hour of simulated time. Preliminary studies 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 minutes of simulated time, depending on the traffic load.
- B. An error-free environment is assumed. In most of the studies, it is assumed that all traffic is received error-free, that modems and cryptos always sync correctly and that reservation synchronization is never lost. Collisions, of course, can still occur on slots transmitted in the contention subframe. The assumption of an error-free environment is probably not too severe for the MAT system when operating over FLTSAT in good weather and in the absence of scintillation, and it provides a common base for comparison. To explore the sensitivity of the system performance with respect to this assumption, a model of an error environment was developed and used in a set of simulation studies. The results are reported in Section 7.4.1, below.
- C. Propagation delays are assumed to be .25 seconds (round trip) for all stations. This assumption corresponds to operation with the satellite approximately overhead. Worst case conditions (satellite on the horizon) could lead to round trip delays of .28 seconds.
- D. Parameters for protocol overheads, frame formats, and message arrival rates per terminal are fixed. The determination of these parameters was described in Section 6. Each additional terminal in a configuration adds traffic to the ship on which it is placed and adds corresponding traffic (computer responses) to the shore station. Except in the studies of single priority configurations, the shore-originated traffic has priority over ship traffic in the CPODA queues. (Delays for the ship-originated traffic may still be less than for traffic from the shore, however, since messages are queued in the particular station before they enter the CPODA queues).

7.2 Frame Length Studies

The initial series of experiments measured the delay for decreasing values of the CPODA frame length. A network of one shore station and six ships, with two terminals per ship was chosen as the environment for these tests. The results are displayed in Figure 7-1. From these tests it is clear that, for this loading, the shortest frame length yields the shortest delays.

The linear relation between frame length and delay for a system with low utilization can be interpreted as follows: each message must wait, on the average, one half frame after its arrival for the start of the next frame. A reservation message transmitted in that frame is not served until the information subframe which occurs at the beginning of the second frame following the message's arrival. Thus a message incurs an average delay of roughly 1.5 frames (plus a round trip transmission delay and any queuing delay).




Fig. 7-1 - Mean message delay vs. CPODA frame length

As the frame length approaches the round trip delay time, no further decrease in message delay is to be expected, since the minimum delay possible is two round trip propagation delays (one for the reservation plus one for the message). In addition, shorter frames have more overhead bits per frame, since the leader packet is fixed length.

From the test results for this baseline, the .30 second frame yielded the shortest delays, but a frame of .5 seconds was chosen for the remainder of the tests. Two reasons lead to this choice: (1) at higher utilizations, the .5 second frame would lead to less overhead, and (2) the simulation overhead (the time spent by the simulator generating dummy leader packets) is greater per simulated hour for the .30 second frame. In actual MAT operation, the CPODA frame length should be a parameter that can be tuned over a range, so the network can be adapted for different loads.

7.3 Traffic Loading Studies

The central group of studies investigates the performance of the protocol for a given seed set under varying traffic loads. These experiments fall into three subcategories:

- A. Fix the number of terminals per ship and increase the number of ships.
- B. Fix the number of ships and increase the number of terminals per ship.
- C. Fix the total number of terminals and distribute them equally among an increasing number of ships.

In this section we examine first the delay measurements from these experiments and then some other measures of interest.

7.3.1 Delay Measurements

Figures 7-2 through 7-5 display the average delays observed for various numbers of ships and terminals per ship in four different formats. Delay includes the elapsed time from the generation of a message at the source (ship or shore station) until it is successfully received at its destination. An input/response sequence from a shipboard terminal to a shore-based computer thus requires both a ship/shore delay and a return shore/ship delay (plus any additional delays to transmit the request through ARPANET, process it at the host computer, and transmit the response via ARPANET to the shore station).

As the graphs show, for the experiments conducted the mean delay increases only gradually with the load. The range of mean delays observed is from 1.3 to 2.0 seconds, but additional replications with other seed sets would be required to establish confidence in these absolute values (see below).

Plotting delay as a function of traffic load (number of terminals), as shown in Figure 7-5, yields a multiple valued function because of the effects of the distribution of terminals among ships. That is, if twenty terminals are on a single ship, the traffic for all twenty terminals is queued and




Fig. 7-2 - Mean message delay vs. number of ships

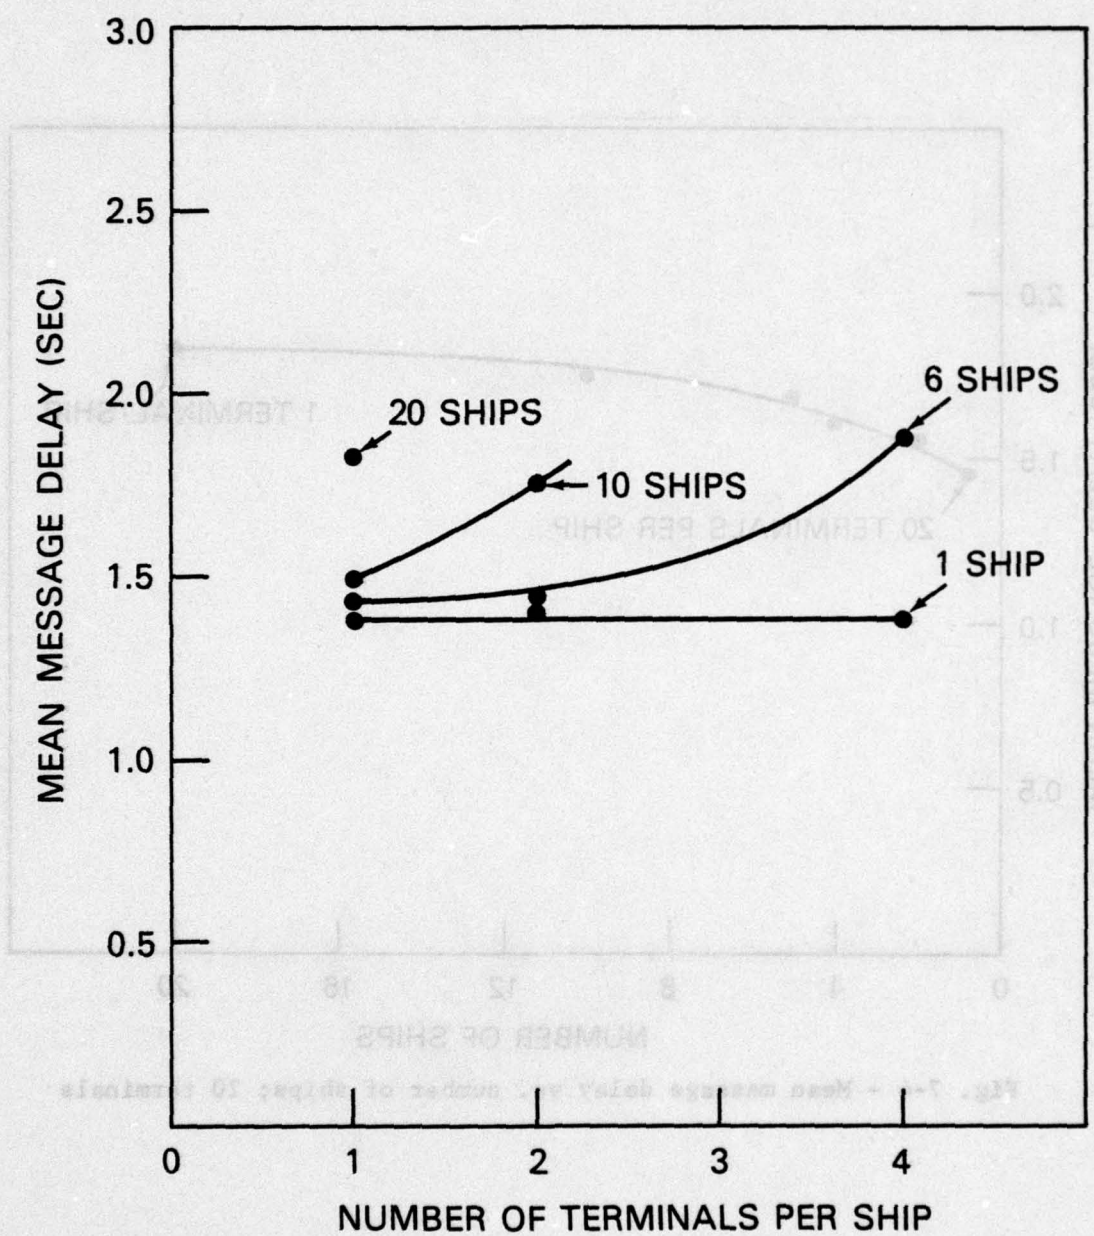


Fig. 7-3 - Mean message delay vs. number of terminals/ship




Fig. 7-4 - Mean message delay vs. number of ships; 20 terminals

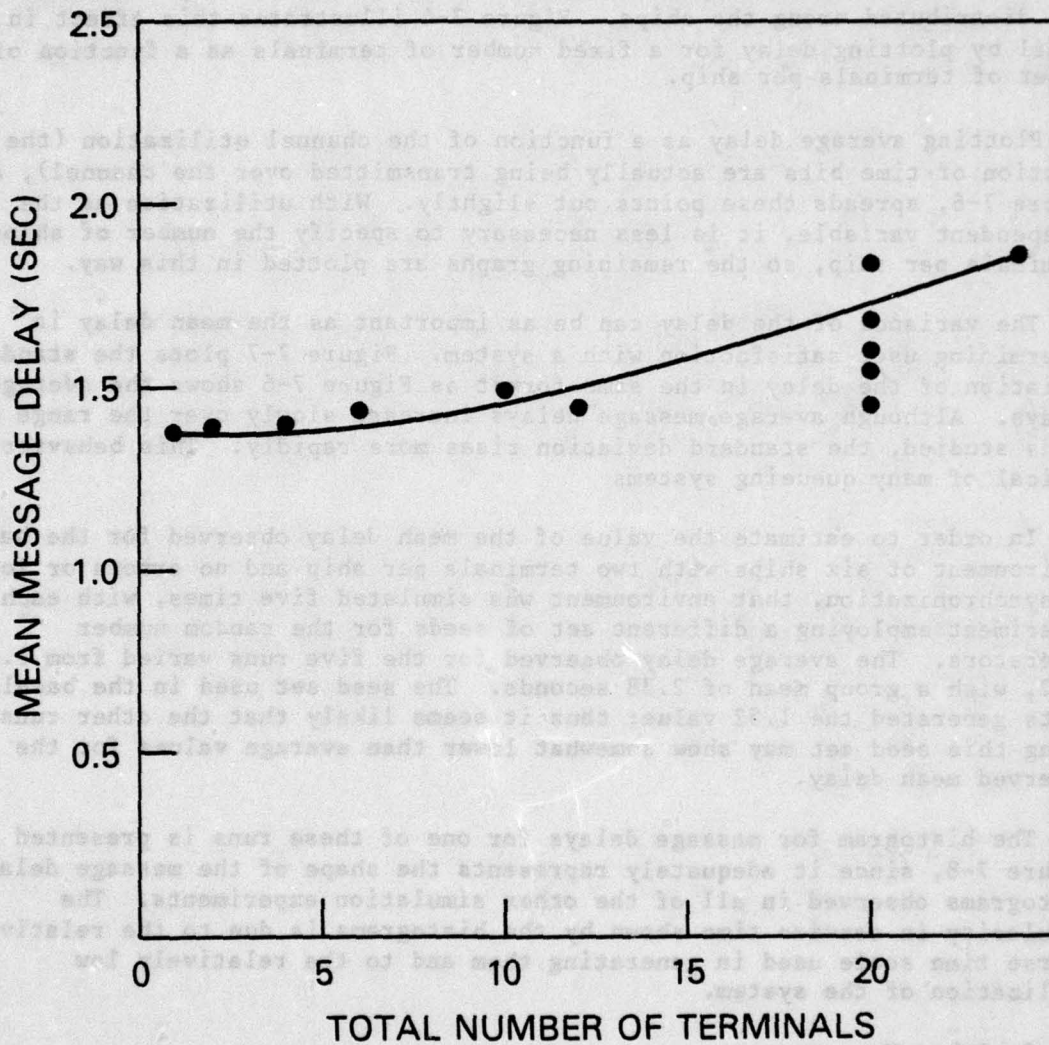


Fig. 7-5 - Mean message delay vs. total terminal load

transmitted by that ship without interference (collisions) occurring among those terminals. If the same twenty terminals are distributed among twenty ships (one per ship), then each ship must broadcast reservations, collisions may occur, and the terminals may experience longer average delays.

Thus, the group of points plotted in Figure 7-5 representing the delays at load = 20 terminals is dispersed because of the different ways the terminals were distributed among the ships. Figure 7-4 illustrates this effect in detail by plotting delay for a fixed number of terminals as a function of the number of terminals per ship.

Plotting average delay as a function of the channel utilization (the fraction of time bits are actually being transmitted over the channel), as in Figure 7-6, spreads these points out slightly. With utilization as the independent variable, it is less necessary to specify the number of ships and terminals per ship, so the remaining graphs are plotted in this way.

The variance of the delay can be as important as the mean delay in determining user satisfaction with a system. Figure 7-7 plots the standard deviation of the delay in the same format as Figure 7-6 shows the average delays. Although average message delays increase slowly over the range of loads studied, the standard deviation rises more rapidly. This behavior is typical of many queueing systems

In order to estimate the value of the mean delay observed for the baseline environment of six ships with two terminals per ship and no errors or losses of synchronization, that environment was simulated five times, with each experiment employing a different set of seeds for the random number generators. The average delay observed for the five runs varied from 1.52 to 4.32, with a group mean of 2.38 seconds. The seed set used in the baseline tests generated the 1.52 value; thus it seems likely that the other runs made using this seed set may show somewhat lower than average values for the observed mean delay.

The histogram for message delays for one of these runs is presented in Figure 7-8, since it adequately represents the shape of the message delay histograms observed in all of the other simulation experiments. The regularity in service time shown by the histograms is due to the relatively coarse time scale used in generating them and to the relatively low utilization of the system.

7.3.2 Other Measurements

In addition to recording message delays, the simulator records measures of many other aspects of the system operation. This section highlights a few of these that illuminate the behavior of the MAT network with the distributed CPODA protocol.



Fig. 7-6 - Mean message delay vs. utilization




Fig. 7-7 - Standard deviation of message delay vs. utilization

EACH X = 41.0000 OBSERVATIONS

	UPPER BOUND	NUMBER OBS	
	0.0000	0.0000	
	2.0000	2029.0000	XX
	4.0000	90.0000	XX
	6.0000	10.0000	
	8.0000	1.0000	
DELAY	10.0000	1.0000	
(SEC)	12.0000	1.0000	
	14.0000	0.0000	
	16.0000	0.0000	
	18.0000	0.0000	
	20.0000	0.0000	
	22.0000	0.0000	
	24.0000	0.0000	
	26.0000	0.0000	
	28.0000	0.0000	
	30.0000	0.0000	
	32.0000	0.0000	
	34.0000	0.0000	
	36.0000	0.0000	

PRIORITY 1

EACH X = 42.0000 OBSERVATIONS

	UPPER BOUND	NUMBER OBS	
	0.0000	0.0000	
	2.0000	2052.0000	XX
	4.0000	166.0000	XXX
	6.0000	9.0000	
DELAY	8.0000	1.0000	
(SEC)	10.0000	1.0000	
	12.0000	0.0000	
	14.0000	0.0000	
	16.0000	0.0000	
	18.0000	0.0000	
	20.0000	0.0000	
	22.0000	0.0000	
	24.0000	0.0000	
	26.0000	0.0000	
	28.0000	0.0000	
	30.0000	0.0000	
	32.0000	0.0000	
	34.0000	0.0000	
	36.0000	0.0000	

PRIORITY 2

Fig. 7-8 - Sample histograms for message delay

30

A. Throughput

The throughput is the number of information bits transmitted during a period. The total number of bits transmitted (as used in calculating the utilization) includes overhead bits as well as information bits. The relation between throughput and utilization, as displayed in Figure 7-9, is approximately linear. This relation is not surprising; it merely indicates that the number of overhead bits transmitted is a linear function of the number of information bits. If the two scales are resolved into common units, the slope of the line is approximately .35 information bits per bit actually transmitted (neglecting the half rate FEC coding).

B. Message Backlog

The average number of messages queued for transmission affects the buffer sizes required in the ship and shore stations. A message is considered backlogged until the sending node receives an acknowledgment for that message. The simulator records the total number of messages backlogged throughout the network, so the average backlog figures correspond to a global buffer requirement. Figure 7-10 displays the average backlog observed as a function of the utilization.

C. Use of Piggybacking

The distributed CPODA protocol includes a mechanism for piggybacking reservations and acknowledgments on information slots in order to reduce use of the contention subframe. As the utilization of the system increases, the proportion of reservations and acknowledgments piggybacked should increase. Figure 7-11 plots the ratio of piggybacked to contended reservations and acknowledgments and the ratio of the total number of information to contention packets transmitted as a function of traffic load. The results do indicate an increase in the piggybacking of reservations as the utilization increases, but these measures generally fluctuate widely. Further investigation seems necessary to explain these observations.

7.4 Additional Studies

To allow the establishment of a baseline performance for CPODA in an error-free environment, the studies reported in Section 7.3 made the assumptions that crypto and modem sync always occur correctly and that the codec can correct all transmission errors. In addition, the shore station was always given priority in the CPODA transmission queues. The studies reported in this section test the effects of allowing some transmission and synchronization errors and explore the behavior of a single priority system in a limited number of test cases.

7.4.1 Studies with Transmission and Synchronization Errors

A simple model for the error environment was developed and the simulator was used to test the operation of the MAT system in such an environment. Precise modelling of the sources and effects of errors on a satellite channel




Fig. 7-9 - Throughput (message bits) vs. utilization




Fig. 7-10 - Mean message backlog vs. utilization

with a shipboard receiver is beyond the scope of this project. Consequently, the goal of these studies was to examine the behavior of the MAT network in a "moderate" error environment, so that the general effect of errors on delay could be observed. No attempt was made to assess at what noise level the system would begin to fail.

7.4.1.1 Parameter Development for the Error Environment

To successfully receive an information or contention packet, a receiver must: (1) establish modem synchronization; (2) establish crypto synchronization; (3) receive the header and data portion of the packet error-free (after FEC 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 (following a timeout) by the sender. In addition, the distributed CPODA protocol 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.

Two graded environments were defined to examine the effects on errors: in the first environment, detected errors were allowed to occur in data and header blocks, but loss of crypto, modem, or reservation synchronization did not occur. The second environment allowed crypto, modem, and reservation synchronization errors as well as errors in data and header blocks. A final test examined the second environment under a higher load (6 ships with 4 terminals per ship) and a correspondingly higher rate of loss of reservation synchronization. The parameter sets used are shown in Table 7-1; their derivation is given below.

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 225 bits long; so, at a rate of 10^{-5} , roughly one block in 400 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 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 simulation of a six ship MAT network with two terminals per ship, the number of reservations per hour was approximately 5,000. Assuming that each reservation represents a separate trial, this figure (together with the modem, crypto and header error probabilities) indicates that each ship would lose reservation sync approximately 10 times per hour, or once every 6 minutes. For the six ship network with four terminals per ship, roughly 10,000 reservations per hour are transmitted, leading to a mean time between loss of reservation sync of about 3 minutes.

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 an 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 seconds (4 half-second frames) was chosen as a reasonable value for the mean time to regain scheduling sync in both cases.

7.4.1.2 Observations of Error Studies

The baseline environment (six ships, two terminals per ship) was tested in the environments described (with increasing error problems), and the six ship, four terminals per ship configuration was tested in the full error environment. The observed mean delays for these tests are shown in Figure 7-12, along with corresponding delays for the environments without errors.

The degradation in mean delay between the no error environment and the full error environment (in the absence of scintillation) for the six ship, two terminals per ship configuration is slightly less than 25%. In the more heavily loaded configuration, the degradation added by errors is more severe--approximately 50%. The total delays in both cases would be easily tolerable in a message transmission system; they are at the borderline of acceptability for typical interactive computer use.

The increase in delays observed with the error rate may be partly controllable with the block time-out parameter. This parameter determines the length of time a node waits for an acknowledgment before retransmitting the message. By decreasing the parameter, the delay introduced by an incorrect transmission may be reduced, but the probability of sending a duplicate packet (in case of a correct reception but a slow acknowledgment) increases. Another approach to reducing the delay added by errors would be to introduce a negative acknowledgment (NAK) to be transmitted when a packet was received in error from a known source node.

The error model used in these studies is a relatively unsophisticated one. If more accurate statements about the impact of errors on the system are desired, a more precise model will be required.

CASE 1: BASELINE (6 SHIPS, 2 TERMINALS PER SHIP) NO ERRORS
CASE 2: 6 SHIPS, 4 TERMINALS PER SHIP, FULL ERROR ENVIRONMENT
CASE 3: 10 SHIPS, 1 TERMINAL PER SHIP, NO ERRORS




Fig. 7-13 - Comparison of one and two priority systems

7.4.2 Single Priority Studies

The basic MAT environment modelled called for shore-originated messages to be assigned priority 1 and ship-originated messages to be priority 2, in order to minimize queues at the shore station. A few studies were run to assess the behavior of a single priority system as well; these are documented in Figure 7-13. The mean message delays (averaged over messages of all priorities) are approximately the same in both cases, as one would expect. The average message backlog (summed over all nodes) is only slightly increased in the single priority case. Unfortunately, the simulator did not collect statistics on queue length on a node by node basis; thus we cannot infer that the queue length at the shore node increased in the single priority case.

8. Conclusions and Recommendations

This section provides a summary and assessment of the observed system behavior of the MAT network with the distributed CPODA protocol and recommends a course of development for the MAT network.

8.1 CPODA Performance in the MAT Network

- A. Capacity: The simulation studies demonstrate that the MAT network, operating with the distributed CPODA protocol, has sufficient capacity for up to twenty ships with one terminal per ship or six ships with up to four terminals per ship. No larger configurations than these were tested. The traffic loads for these tests called for an average of one message per terminal per 20 seconds from ship-to-shore and an equal number in the reverse direction, with message lengths of several characters ship-to-shore and several tens of characters shore-to-ship.
- B. Delay: For the baseline environment of 6 ships and 2 terminals per ship (and assuming no transmission or synchronization errors), mean delays varied from 1.52 seconds to 4.32 seconds per message in the course of five different experiments, each of one simulated hour. The group mean was 2.38 seconds. The addition of normal transmission and synchronization errors to the environment will increase these delays. Variances in delays were also significant. For a normal land-based computer-communications link, such values would be uncomfortably high. An average round trip transmission delay of more than 4 seconds would irritate most users of interactive computer systems. Nevertheless, these delays are considerably lower than might be experienced if presently operational Navy communication systems were employed. When lightly loaded, the MAT system with CPODA can achieve average one-way delays of less than 1.5 seconds. It should be noted that the figures cited here are for channel delays only; the total system delay can be significantly longer than the channel delay if, for example, host processing is slow.
- C. Stability: For the limited range of cases tested, no instabilities appeared in the protocol operation. Increases in traffic naturally increased utilization, delays, and delay variances, but these increases were gradual. The highest utilization observed was

slightly under .4. (Utilization is the fraction of time that bits are actually being transmitted over the channel, and includes all synchronization and overhead bits.)

- D. Sensitivity to Parameter Settings: Parameter settings that can affect the performance of the protocol include the CPODA frame length, the maximum and minimum information packet lengths, the minimum number of contention slots per frame, the length of the timeout period before a packet is retransmitted, the number of acknowledgments and/or reservations allowed per contention packet, and the priority queue structure for reservations. This study found that the frame length had a decisive effect on delays if it was much longer than a round-trip propagation delay. Values for the other parameters were determined based on the traffic model and the existing MAT design and appeared to be satisfactory.

3.2 Recommendations

The principal goal of the MAT project, as stated in [(2)] is to provide a shipboard access to ACCAT facilities. The performance studies reported here indicate that the distributed CPODA protocol operated over a FLTSAT channel at a 19.2KSps transmission rate (9600 information bits per second coded at half rate) can support this type of access. Delays will be somewhat longer than is desirable for a comparable land-based transmission path, but should be tolerable.

Comparison of delays between CPODA and alternative satellite channel protocols is beyond the scope of this project; however, the implementation of a simple Fixed Time Division Multiple Access (FTDMA) protocol for initial link testing and for backup (as proposed in (3)) is advised. Distributed CPODA operating in a shipboard environment still must be viewed as experimental.

A goal for the MAT that is secondary at present, but still important, is to provide a test bed for evaluating new access protocols for satellite broadcast channels. Navy communication systems do not generally face the stringent requirements for low delay that interactive computer access imposes. A study performed recently for OP-943 [(9)] has investigated the CPODA protocol as a vehicle for the consolidation of systems presently using dedicated FLTSAT channels. The MAT network offers an attractive environment for the testing and evaluation of this concept.

To be of maximum benefit as an experimental test bed for protocols, the MAT implementation should include a flexible facility for the measurement of protocol performance. Measures of interest in the CPODA protocol include the following:

1. Packet and message delivery delays.
2. Rate of collisions.
3. Error rate for data blocks.

4. Error rate for header information.
5. Rate of loss of scheduling synchronization.
6. Rate of use of piggybacking (for reservations and acknowledgements).
7. Number of information bits per packet.
8. Number of information bits per message.

Instrumentation for these measures may already exist in the CPODA software developed for the ARPA packet satellite experiments, but no explicit mention of measurement capabilities has appeared in MAT documentation to date.

In addition, facilities for the generation of dummy traffic loads would enable studies of system behavior under identical loads but with different error conditions and parameter settings. It should be straightforward to build such facilities (or to include software arrangements for their later implementation) in the shipboard red processors.

Finally, because delays in the MAT network may be expected to exceed those in comparable shore-based systems and users may on occasion grow impatient, we recommend that some means for providing users with information on the present state of the network queues and delays be implemented. The red processor must maintain up-to-date information on queue lengths for CPODA in any case; relaying this information to the waiting user (either periodically or at his request) should add a negligible burden to the system.

In summary, it is recommended that the MAT implementation proceed essentially as proposed in [(3)], but with special attention to (1) inclusion of the backup FTDMA mode to assure reliable ACCAT access and (2) inclusion of sufficient software to enable realistic experiments to be conducted on communication protocols. The latter software should record and report at least the performance parameters enumerated above and should provide information to users on the current state of network queues and delays. It should also provide for the generation of dummy traffic loads.

9. References

1. Jacobs, I., et al. CPODA - A Demand Assignment Protocol for SATNET. Proc. Fifth Data Comm. Symp., ACM-IEEE, Sept., 1977.
2. Proposal for Phase 2 Mobile Access Terminal for ACCAT. E-Systems, ECI Division, April, 1978.
3. BBN Technical Proposal, Mobile Access Terminal (MAT), proposal No. P78-CSY-83, April 1978, submitted to NAVELEX Code 310.
4. Interim Report, Mobile Access Terminal for ACCAT, E-Systems, ECI Division, Nov., 1977.
5. Fuchs, E. and P. E. Jackson. Estimates of the Distribution of Random Variables for Certain Computer Communications Models. CACM 13, 12 (Dec 1970) 752-757.
6. Jackson, P. E. and C. D. Stubbs. A Study of Multiaccess Computer Communications. AFIPS Conf. Proc. Vol 34 (SJCC 1969) 491-504.
7. Johnson, R. W. Computer Programs for Determining Probability Distributions by the Principles of Maximum Entropy and Minimum Cross-Entropy. In preparation.
8. Hsu, Nai-Ting and Lin-Nan Lee. Channel Scheduling Synchronization for the PODA Protocol. 1978 International Conf. on Comm., Conf. Rec. Vol 3, pp. 42.3.1-42.3.5, IEEE Cat #78CH1350-8CSCB.
9. LEASAT Consolidation Study Final Report (draft), NRL, submitted to OP-943 Dec. 1978.
10. Jacobs, I. M., R Binder, and F. V. Hoversten. General Purpose Packet Satellite Networks. Proc. IEEE 66, 11 (Nov. 1978) 1448.