

AD-A115 133

RICE UNIV HOUSTON TEX DEPT OF ELECTRICAL ENGINEERING F/G 17/2
DATA LATENCY IN TIME MULTIPLEXED BUS SYSTEMS FOR MISSILE APPLIC--ETC(U)
MAR 82 J B SINCLAIR AFOSR-81-0178

UNCLASSIFIED

AFOSR-TR-82-0415

NL

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1

1-1


1-1

END

DATE

7 82

DTIC



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS 1963-A

14

AFOSR-TR- 82 - 0415

AD A115133

DATA LATENCY IN TIME MULTIPLEXED BUS
SYSTEMS FOR MISSILE APPLICATIONS

Dr. J.B. Sinclair
Department of Electrical Engineering
Rice University
Houston, Texas 77251

February 22, 1982

FINAL REPORT
Grant No. AFOSR 81-01/8

DTIC FILE COPY

DTIC
ELECTE
JUN 04 1982
E

Approved for public release;
distribution unlimited.

82 00 03

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER AFOSR-TR- 82-0415	2. GOVT ACCESSION NO. AD-A115133	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) DATA LATENCY IN TIME MULTIPLEXED BUS SYSTEMS FOR MISSILE APPLICATIONS		5. TYPE OF REPORT & PERIOD COVERED FINAL REPORT 01 JAN 81 Thur 22 FEB 82
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) J.B. SINCLAIR		8. CONTRACT OR GRANT NUMBER(s) AFOSR - 81-0178
9. PERFORMING ORGANIZATION NAME AND ADDRESS RICE UNIVERSITY DEPARTMENT OF ELECTRICAL ENGINEERING HOUSTON, TX 77251		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 61102F 2305/D9
11. CONTROLLING OFFICE NAME AND ADDRESS AIR FORCE OFFICE OF SCIENTIFIC RESEARCH /NE BOLLING AFB, WASHINGTON, DC 20332		12. REPORT DATE 11 Mar 82
		13. NUMBER OF PAGES 22
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Data Bus, Time Multiplexed Bus Systems, Computer communication bus protocols, Round robin passing protocol, Carrier-sense, Multiple-access with collision detection, Guided weapon, Mid-course guidance.		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This research examines the performances of three computer communication bus protocols: round robin passing protocol (RRPP), modified round robin passing protocol (MRRPP), and carrier-sense multiple-access with collision detection (CSMA/CD). These are compared through the simulation of a medium-range air-to-surface guided weapon federated computer network, specifically the mid-course guidance phase. The performance of each protocol is measured by the average and maximum message waiting time and latency. - OVER -		

DD FORM 1473

EDITION OF 1 NOV 65 IS OBSOLETE

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

Under these performance measures, the CSMA/CD protocol is clearly superior, with average message waiting time as low as 22% of that for RRPP and 8% of that for MRRPP. Maximum waiting times are also lower for CSMA/CD, although by a much narrower margin. This indicates the potential for the use of CSMA/CD communication channels in low channel utilization, delay-sensitive applications.

Accession For	
NTIS GRA&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A	



ABSTRACT

This research examines the performances of three computer communication bus protocols: round robin passing protocol (RRPP), modified round robin passing protocol (MRKPP), and carrier-sense multiple-access with collision detection (CSMA/CD). These are compared through the simulation of a medium-range air-to-surface guided weapon federated computer network, specifically the mid-course guidance phase. Four different configurations, each involving a different channel throughput, are investigated.

The performance of each protocol is measured by the average and maximum message waiting time and latency. The message latency is the interval between the time that a message becomes ready for transmission and the time that the (successful) transmission begins. The message waiting time is the sum of the message latency and the message transmission time.

Under these performance measures, the CSMA/CD protocol is clearly superior, with average message waiting time as low as 22% of that for RRPP and 8% of that for MRKPP. Maximum waiting times are also lower for CSMA/CD, although by a much narrower margin. This indicates the potential for the use of CSMA/CD communication channels in low channel utilization, delay-sensitive applications.

AIR FORCE OFFICE OF SCIENTIFIC RESEARCH (AFSC)
NOTICE OF TRANSMITTAL TO DTIC
This technical report has been reviewed and is
approved for public release JAW AFR 190-12.
Distribution is unlimited.
MATTHEW J. KERPER
Chief, Technical Information Division

1. Introduction

The purpose of this report is to compare the performances of three time-division multiplexing schemes for access control of a serial bus in a guided weapon application. The three access control methods are:

- (1) Round-robin passing protocol (RRPP)
- (2) Modified round-robin passing protocol (MRRPP)
- (3) Carrier-sense multiple-access control with collision detection (CSMA/CD)

The principle criterion to be used in comparing performance will be the average message waiting time, where message waiting time is defined to be the length of the interval between the time at which a message is ready to be transmitted and the time that the message transmission is successfully completed. We are also interested in the maximum message waiting time as well as average and maximum message latencies, where message latency is the message waiting time minus the message transmission time. Each of these will in general be a function of the average channel or bus utilization. The channel utilization is the ratio of the number of message and routing control bits transmitted over the channel during some interval to the total number bits that could have been transmitted during that interval.

2. Description of access control methods

This section briefly describes some of the important features of each of the three access control methods that were considered. For more complete description, see the references in the bibliography.

2.1. Round-robin passing protocol

This method is also known as the DISMUX bus or MIL-STD 1765 [1][2][4]. There is no centralized bus control. Every processor is connected to the bus via a Bus Interface Unit (BIU) which is responsible for acquiring bus access when its associated processor wishes to transmit information, and for reading information from the bus which its processor should receive. Each BIU occupies a fixed position in a cycle which includes all BIU's in the system. This cycle determines the sequence in which BIU's gain control of the bus for message transmission.

A BIU which acquires control of the bus but has no information to transmit broadcasts a minimum-length message relinquishing control to the next BIU in the cycle. A BIU which wishes to transmit information prefaces each message with appropriate control information and transmits these messages consecutively, terminating the string of messages by relinquishing control of the bus.

There are several important points to be emphasized here. First, control of the bus is passed from one BIU to another cooperatively. That is, a BIU in control of the bus voluntarily releases the bus, and then another BIU (the next BIU in the cycle) acquires control. This means that control of the bus is distributed across all BIU's in the cycle. Second, the cycle in which BIU's acquire control is predetermined and cannot be modified during normal operation of the bus. This is basically an implementation decision, since one can easily imagine a system allowing dynamic cycle modification. Third, the possibility exists that a BIU may malfunction and fail to respond to its

turn in the cycle. Since this would halt all bus communication in the system, provision is made for a supervisor processor in the system to monitor the round-robin sequencing. If a BIU fails to initiate a transmission within a specified interval after its predecessor releases control of the bus, the supervisor will transmit the bus-release message for the mute BIU, allowing its successor to resume normal sequencing of control.

RRPP is a "fair" scheme (as are the other two protocols), in that it allows all processors equal access to the bus. BIU's are restricted in the total amount of information that they may transmit during one cycle, thereby limiting the maximum length of a cycle to the number of BIU's times the maximum transmission interval per BIU. The distributed control of the bus, together with the monitor role of the supervisor, makes the system relatively insensitive to the failure of a single BIU in terms of disrupting the flow of control from one BIU to the next in the round-robin cycle. Even the failure of multiple BIU's can be overcome in this scheme.

The failure of the supervisor's BIU is fatal, however. Since the supervisor's BIU is in the round-robin cycle, its failure will in fact completely disrupt bus communications.

Under fairly heavy loading conditions, the utilization of the DISMUX bus is quite good, since the only overhead is the message by which each BIU releases control of the bus to the next BIU in the cycle. The only difficulty is the delay between the time at which a message becomes available for transmission and the time at which the BIU able to begin the transmission. For heavily loaded systems, the round-robin

passing protocol keeps average message delays reasonably small. In lightly loaded systems, the average delay is unnecessarily large because of the protocol overhead. Furthermore, in heavily loaded systems, some messages may be more time-critical than others and the transmission of these messages should receive higher priority. This is difficult to do in the context of the DISMUX bus.

2.2. Modified round-robin passing protocol

MRRPP is a variation of MIL-STD 1553B [3][4]. This scheme is similar to the RRPP scheme described above, with the exception that passing control from one BIU to another on the bus requires an acknowledgement on the part of the BIU receiving control. This introduces extra overhead and results in a lengthened minimum bus cycle time. MRRPP also requires that MIL-STD 1553B be used in a broadcast mode, a concept rather foreign to the original intent of the protocol.

Because of its similarity to RRPP, MRRPP has many of the same advantages and disadvantages. The principal difference between the two is the increased overhead in MRRPP which causes longer average delays and maximum delays.

2.3. Carrier-sense multiple-access with collision detection

CSMA/CD differs from the two methods listed above in a number of ways. In a CSMA/CD system, there is no centralized control of the communication channel. Control is again distributed among all BIU's

connected to the channel. Message transmission is accomplished by a BIU first gaining exclusive access to the channel, so CSMA/CD is a TDM access method, also.

The most important difference between CSMA/CD and both RRPP and MRRPP is the mechanism for obtaining access to the channel. In CSMA/CD, BIU's may request channel access whenever they can detect no activity (no transmissions) on the channel. A request for access is made simply by starting to transmit on the channel. A conflict arises when two or more BIU's begin to transmit at the same, or very nearly the same, time. This is called a collision.

In the absence of a collision, a BIU transmits data until a message transmission is complete. If a collision occurs, each of the conflicting BIU's detects the collision by "listening" to the channel as it transmits. If the data that it "hears" is not the data it is transmitting, the BIU knows that a collision has occurred and terminates the transmission. The access control algorithm in each BIU will then determine a retransmission time for the message, usually through a randomizing procedure to minimize the probability of participating in another collision. The term "carrier sense" is something of a carry-over from radio packet communication systems, where activity on the channel was indicated by the presence of a carrier wave. In a CSMA/CD system, information is transmitted digitally using a Manchester or phase encoding of the data, which requires a signal transition in the center of each bit window. The occurrence of these transitions is the digital equivalent of the presence of a carrier waveform in a radio system.

In its simplest form, a CSMA/CD system is "fair"; that is, each BIU uses the same access algorithm and hence competes on an equal basis for access to the channel. This project had as one objective the consideration of systems in which BIU's had different priorities, enforced by the access algorithm. This might be extremely beneficial in a system in which it was necessary to minimize the delay for some messages at the expense of increasing the average delay for other, less time-critical messages. An example might be a guided weapon vehicle in which navigational information is delivered to the autopilot via the channel rather than via a special purpose interface. As the results of the simulation show, for the estimated amount of message traffic in the system under consideration, priority transmissions would be of relatively little use because delays in message transmissions are relatively small in a CSMA/CD system.

3. The organization of a CSMA/CD system

3.1. Introduction

In this section we consider the details of the CSMA/CD system used as a model for the simulation. To be able to compare results directly, we have assumed identical channel characteristics for all three access methods. The channel is a 1 megabit per second bus. Information on the channel is phase-encoded, with each 16 bits of information requiring 20 bit times (20 microseconds) for transmission. In each of the three access methods, messages are not destination-specific. Rather, each

message includes a message identifier in the message header word which precedes the transmission of a message. A BIU continually monitors the channel for the transmission of any message in which it is interested. The purpose of this was to allow broadcast of messages to several destinations without the overhead of specifying all destinations in the message header. A similar result can be achieved by having the message header carry a destination field which identifies a set of processors (which may consist of a single processor or even all processors). In this approach, messages are divided into classes according to their destination sets. If each class consists of a single message, this approach is equivalent to the one actually simulated.

3.2. Implementation of a simple CSMA/CD system

The delay algorithm used for this simulation is a binary exponential backoff algorithm similar to that used in Xerox's Ethernet [5][10]. When a message is involved in a collision, the BIU which attempted to transmit the message calculates a maximum backoff time which is twice the previous backoff time (the backoff time for the initial transmission attempt is assumed to be 1 time unit). The BIU then schedules the retransmission attempt by randomly and uniformly selecting a retransmission time between 1 and the maximum backoff time. Of course, if a transmission is in progress when the retransmission time arrives, the retransmission is queued until the channel becomes idle again.

The time unit for determining retransmission times was chosen to be 10 microseconds. In an Ethernet network, the time unit (also called the

slot time) is a function of several parameters, including the maximum span of the channel, the turn-around time in the transceivers, signal rise times, encoding time, etc. Since we are dealing with a bus of length a few meters at most, the slot time does not depend on round trip propagation time and is almost entirely a function of the signal characteristics and the time required to detect that a collision has occurred. An 10 microsecond slot time is felt to be conservative.

The simulation assumes that each BIU has three queues: a primary message queue, an auxiliary queue, and a collision queue. New messages are placed in the primary message queue. All messages in the primary message queue are transmitted as a single packet. When a collision occurs, all messages in the packet involved in the collision are placed in the collision queue until the scheduled retransmission time. Messages are transferred from the collision queue to the primary message queue when their retransmission time has elapsed, unless the BIU is currently transmitting, in which case the messages are transferred to the auxiliary queue. When the BIU completes transmission from the message queue, it examines the auxiliary queue. If the auxiliary queue is nonempty, its contents are transferred to the primary message queue and the BIU begins to contend for the channel again.

If a packet is involved in a collision but the BIU already has messages waiting in the collision queue for retransmission, the simulation is aborted with a message indicating that the collision queue is busy. This was done only to make the simulation a little simpler. In this way, a single value associated with the collision queue specifies the retransmission count for all messages in the queue. If we want to allow messages with different retransmission times to be in the

collision queue simultaneously, each message could have an associated retransmission count. We felt that this would probably be unnecessary for the purposes of simulation because the occurrence of a collision when a collision buffer was nonempty was a low probability event. Simulation results support this assumption.

The details of collision detection and enforcement are not specified. However, it is assumed that a collision is detectable within a window of 10 bit times. A packet transmission is preceded by the transmission of an illegal bit string of 3 microseconds duration as a synchronization mechanism. This involves a transition on the channel in the middle of the second bit window, and hence there is a delay of approximately 2 bit times before other BIU's are able to sense the carrier. Additional delays are incurred due to the encoding, bit transmit, bit receive, and decoding times in each BIU.

If a collision is detected in the first 10 bit times from the start of transmission (including the 3 bit times for synchronization), each BIU detecting a collision while transmitting will attempt to enforce recognition of the collision by all BIU's by jamming the channel with some predetermined bit pattern to complete the 10 bit time collision enforcement window. It is assumed that any collision will be detected during this window. Collisions which are not detected will often cause parity errors, and hence the packet would be rejected anyway. If this error control is not sufficient, a longitudinal check code or CRC could be appended to each packet to eliminate the possibility of accepting an incorrect message as valid data.

It is desirable to make the collision enforcement window as short as possible, since once a collision has occurred, the data involved in the collision will have to be retransmitted at a later time and continuing to transmit it at the time of the collision would be an unnecessary use of channel capacity.

Additional details of the CSMA/CD simulation are found in [7] and [11]. Emphasis is placed on structural differences between the CSMA/CD simulator and the simulator used for the round-robin protocols.

4. Simulation results

The simulation for each communication protocol was run using the model of message traffic developed during the summer of 1980 at the Air Force Armaments Laboratory at Eglin AFB, Florida [6]. This model dealt with the midcourse guidance phase of a air-to-surface guided weapon vehicle with a low-cost inertial guidance system (LCIGS) in which navigational data is filtered to improve its accuracy. The filter is derived by comparing the vehicle's navigational system output with superior quality information from the transporting aircraft's navigational system and also possibly from a radar-based terrain matching system or from a global positioning satellite system after the weapon vehicle has separated from the aircraft.

Four cases were considered for each protocol:

- (1) no midcourse navigational updates, only filtered LCIGS data on channel
- (2) no midcourse navigational updates, all LCIGS data on channel

- (including all data to autopilot)
- (3) midcourse navigational updates from an onboard radar-based terrain matching system (TERCOM), only filtered LCIGS data on channel
- (4) midcourse navigational updates from an onboard radar-based terrain matching system, all LCIGS data on channel (including all data to autopilot)

These four cases cover a wide range of channel loads.

The results of the simulations are shown in Figures 1-4. These figures compare the average and worst case message waiting times for each of the four cases. Table 1 summarizes and more precisely quantifies these results. Both average and maximum wait times are given. The wait time is the queueing time for a message; that is, it is the length of the interval between the time at which the message first became available to be transmitted and the time at which its (successful) transmission began. For the CSMA/CD scheme, this includes any time spent in a collision and waiting for retransmission.

Table 1: Comparison of Message Wait Times

		UTG		TERCOM	
		No LCIGS	LCIGS	No LCIGS	LCIGS
RRPP	Average Wait	80.2	96.3	94.9	109.2
	Maximum Wait	524.0	568.0	597.8	772.0
MRRPP	Average Wait	224.6	256.4	253.3	286.5
	Maximum Wait	711.5	852.0	972.0	1148.0
CSMA/CD	Average Wait	24.8	23.2	25.3	23.8
	Maximum Wait	216.0	452.0	420.0	628.0

Appendix A in [11] gives a much more detailed picture of the performance of each scheme. For each of the three protocols and for each of the four system configurations, the simulation gives information on the maximum and average delay, maximum and average wait, frequency of occurrence, and message length. It also gives a histogram of wait times in intervals of 100 microseconds, up to a maximum of 600 microseconds. The wait time differs from the delay time in that the delay time includes the time to actually transmit the message. The delay time can be calculated from the wait time, knowing the message length and the transmission baud rate.

The average wait time per message is relatively unaffected by channel traffic in the CSMA/CD scheme. The average wait times vary from 23.2 to 25.3 microseconds. This agrees with our expectations, since even for the TERCOM configuration with "raw" LCIGS data transmitted on the channel, the channel utilization is quite low (approximately 12%). Only when channel utilizations approach 50% should we see serious degradation of wait times due to channel contention. It is the relative insensitivity of average wait times plus the fact that CSMA/CD had lower maximum wait times for each test case than either RRPP or MRKPP that caused us to reconsider attempting to simulate priority transmissions.

Another set of important performance measures is the average and maximum waits for specific messages, in particular for those messages which are most time-critical, i.e., those messages for which delays have the most affect on system performance. One such message in the system being simulated is the LCIGS data message which is transmitted directly to the digital autopilot in two of the simulated cases. Since this information is part of a feedback control loop, delay in the

transmission of the information from the navigation subsystem to the autopilot is manifested as reduced phase margin in the control system, and hence it is desirable to keep the delay for this message small. Table 2 gives a comparison of the wait times for this message in the TERCOM configuration case, for each of the protocols. The average wait for CSMA/CD is only 20% of the average wait for RRPP and only about 9% of the value for MRRPP. Even maximum wait time favors CSMA/CD, which has a maximum wait time 17% less than RRPP and 51% less than MRRPP.

5. Conclusions

In comparing CSMA/CD to RRPP and MRRPP, only the average and maximum message delays are used. CSMA/CD is a 70% to 78% improvement over RRPP with respect to average message waits and as much as a 91% improvement over MRRPP. CSMA/CD also performs better with respect to maximum waits, although the improvement is much less dramatic. Because the cases simulated involved relatively low bus utilization (about 12% in the worst case), very little would be gained by going to a priority transmission scheme for time-critical messages.

Table 2: Comparison of Wait and Delay Times
for LCIGS Data Message (TERCOM)

	Average Wait	Maximum Wait	Average Delay	Maximum Delay
RRPP	107	592	287	772
MRRPP	278	1012	458	1192
CSMA/CD	21	492	201	672

CSMA/CD also represents an improvement in system robustness, or performance degradation due to an individual processor or BIU failure. In both RRPP and MRRPP, control must be passed from one processor to another in round-robin fashion, and any BIU not prepared to transmit a message must still take control of the bus briefly to pass control to the next BIU in the cycle. A supervisor processor monitors the passage of control, and a momentary or permanent failure on the part of any other BIU to participate in the passage of control causes the supervisor to intervene in the cycle to insure that the flow of control is uninterrupted. However, the supervisor itself is in the round-robin cycle, and a failure on its BIU's part will cause the flow of control to be halted.

In contrast, the CSMA/CD scheme has no such orderly flow of control from processor to another. Processors attempt to access the bus when a message is available to be transmitted. No supervisor is needed to perform a bus monitoring function. In the CSMA/CD scheme as well as both RRPP and MRRPP, it is necessary to be somewhat paranoid about the design of the BIU, to insure that a BIU failure will be non-disruptive to other BIU's attempting to transmit data on the bus.

The CSMA/CD scheme also allows for simple system expansion, since it is not necessary for each processor to have some information about the round-robin cycle as in RRPP and MRRPP. In particular, a diagnostic station could be added to the network with no disruption of the system. For critical processors in the system, it might even be feasible to have a second BIU, which is normally passive, but which can be switched on in the event of the failure of the primary BIU.

The problems with the CSMA/CD scheme seem to be primarily concerned with the feasibility of implementation in the restricted circumstances of a guided weapon. The absence of any round-robin scheduling scheme is countered by the complexity of the carrier-sense, collision-detection, and collision retransmission schemes. These problems are well understood and there is considerable effort underway in industry to develop low cost CSMA/CD interfaces for systems like Ethernet. Intel is rumored to be undertaking the development of an Ethernet interface on a single chip, a much more complicated undertaking than that necessary for the system proposed here.

Another question about CSMA/CD concerns its viability in those cases in which channel utilization is high. The simple CSMA/CD protocol described above may lead to unacceptably long delays or even an unstable system (an unstable system is one in which an attempt to increase channel utilization actually decreases the net throughput because of the increased number of collisions). Limited contention protocols have been proposed which will force the CSMA/CD protocol to assume the characteristics of a round-robin scheme in the presence of large amounts of channel traffic, while still retaining the desirable characteristics of the simple CSMA/CD protocol at low channel utilizations [8][9]. The cost of this is increased complexity in the control algorithms and typically some predetermined knowledge about other BIU's in the system.

In summary, the CSMA/CD scheme may be a viable alternative to either RRPP or MRRPP for many applications, including the one simulated. If the delay characteristics of the message traffic are of critical importance, CSMA/CD represents a distinct improvement over round-robin communication protocols. The CSMA/CD network is also inherently more

robust than any system which either uses centralized control of the communication medium or uses active control transfers like round-robin cycling. These advantages may make it worthwhile to investigate the feasibility of a CSMA/CD protocol implementation.

There are at least three important questions that deserve additional consideration. The first concerns the complexity of the entire CSMA/CD physical layer protocol and whether it is necessary in the context of low throughput systems such as the one used as a model in this research. Much of the complexity of a CSMA/CD system derives from the collision detection and enforcement algorithms. Given that collisions are rare events, perhaps a simpler CSMA system with no collision detection would be adequate. Other means of error control may be sufficient to handle garbled transmissions.

Second, given that CSMA/CD outperforms RRPP and MRKPP at these low levels of channel throughput and is likely to be worse at relatively high levels of throughput, at what point does the crossover occur? Although this would seem to be within the capabilities of the present simulator to determine, it unfortunately is not. The problem is the implementation of the simulator as described in Section 3.2. It was noted that the collision queue was assumed to hold only one collided message at a time, and if this was ever violated during the course of a simulation, the simulation would be aborted. This works well for low values of channel utilization, but as one attempts to artificially increase channel utilization, for instance by decreasing intermessage intervals, the probability of having more than one collided message at a given station waiting to be retransmitted goes up. The simulation is therefore almost always aborted before very much simulated time has

elapsed. Although the present simulator can be modified to avoid this problem, the modification is nontrivial (the reason that it was not done initially).

The third question deals with the use of adaptive protocols to minimize channel contention for heavily loaded systems, and the relative merits of these protocols compared to some sort of priority scheme that recognizes that all messages are not equal but that it is more critical to avoid long delays with some messages than with others. An additional complication of the adaptive limited contention protocols is that they require estimates of the number of ready stations in the network.

Bibliography

1. N. Nathan et. al., "Digital Processing Analyses/Partitioning," Report No. R-1121, Charles Stark Draper Laboratory, Cambridge, MA., November 1977.
2. - , "Proposed Military Standard: Weapon Internal Time Division Multiplex Data Bus," Draft, AFATL/DLMM, Eglin AFB, FL., March 1980.
3. - , "Military Standard: Aircraft Internal Time Division Command/Response Multiplex Data Bus," MIL-STD-1553B, Dept. of Defense, Washington, D.C., September 1978.
4. J.W. Jones, "DISMUX RRPP and 1553 Modified RRPP," Memorandum, AFATL/DLMM, Eglin AFB, FL., January 1980.
5. R.M. Metcalfe and D.R. Boggs, "Ethernet: Distributed Packet Switching for Local Computer Networks," Comm. ACM, Vol. 19, No. 7, July 1976.
6. J.B. Sinclair, "Characterization of Quasi-Correlated Mux Bus Traffic in a Guided Weapon," Final Report, AFOSR Contract No. F49620-79-0038, AFATL/DLMM, Eglin AFB, FL, August, 1980.
7. J.B. Sinclair, "Mux Bus Traffic Simulator", AFATL/DLMM, Eglin AFB, FL, August, 1980.
8. J.I. Capetanakis, "Tree Algorithms for Packet Broadcast Channels," IEEE Trans. Inf. Theory, Vol. IT-25, pp. 505-515, September 1979b.
9. L. Kleinrock and Y. Yemini, "An Optimal Adaptive Scheme for Multiple Access Broadcast Communication," Proc. ICC, pp. 7.2.1-7.2.5, 1978.
10. - , "The Ethernet: A Local Area Network - Data Link Layer and Physical Layer Specifications," Ver. 1.0, Intel Corporation, September 1980.
11. J.B. Sinclair, "Data Latency in Time Multiplexed Bus Systems for Missile Applications," Tech. Rept. No. 8104, Department of Electrical Engineering, Rice University, Houston, TX, December 1981.




Figure 1

Comparison of message wait times for
UTG-configured system with no LCIGS
data on the channel. All times are
in microseconds.




Figure 2

Comparison of message wait times for
UTG-configured system with LCIGS data
on the channel. All times are in
microseconds.




Figure 3

Comparison of message waiting times for TERCOM-configured system with no LCIGS data on the channel. All times are in microseconds.




Figure 4

Comparison of message wait times for TERCOM-configured system with LCIGS data on the channel. All times are in microseconds.

DATE
FILMED
7-8