

TR-87-14

A Distributed Reservation Scheme for
Spread Spectrum Multiple Access Channels

by

Jeffrey E. Wieselthier
Anthony Ephremides



**A DISTRIBUTED RESERVATION SCHEME FOR
SPREAD SPECTRUM MULTIPLE ACCESS CHANNELS**

Jeffrey E. Wieselthier and Anthony Ephremides

A DISTRIBUTED RESERVATION SCHEME FOR
SPREAD SPECTRUM MULTIPLE ACCESS CHANNELS

JEFFREY E. WIESELTHIER

Naval Research Laboratory
Washington, D.C.

ANTHONY EPHREIMIDES

University of Maryland
College Park, Maryland
and
Naval Research Laboratory

ABSTRACT

A distributed reservation-based multiple access protocol is proposed, analyzed and evaluated. This scheme is based on the quasi-orthogonality property than can be achieved through the use of spread spectrum signaling, which permits the correct reception of a signal despite the simultaneous presence of other channel users. The effect on performance of the ability to tolerate varying levels of other-user interference is demonstrated. No acknowledgment or feedback information from the destination is required, nor is any form of coordination among the users necessary.

INTRODUCTION

We consider the problem of channel access by a population of bursty users. In such systems only a small fraction of the total number of users will typically have a packet ready for transmission at any time. It is usually assumed that if more than one user transmits a packet in any slot, then all packets involved in this "collision" are destroyed. A large variety of schemes have been proposed to handle such multiple access problems, including contention-based schemes (e.g., ALOHA), fixed allocation schemes (e.g., TDMA), channel sensing schemes (e.g., CSMA), reservation schemes, as well as hybrids of two or more of these classes. Reservation schemes are often attractive because stable operation can be maintained at relatively high throughput levels, provided that a reservation mechanism can in fact be implemented. In this paper we present, analyze, and evaluate a distributed reservation scheme that takes advantage of the quasi-orthogonality property that can be achieved through the use of spread spectrum code division multiple access (CDMA) signaling. This property permits the successful reception of a packet despite the simultaneous transmission by other users. The reservation scheme is in fact totally distributed in that it does not require the transmission of acknowledgment or feedback information from the destination to the users. Furthermore, no coordination among the users is required, nor do they even have to monitor each others' transmissions.

RESERVATION SCHEMES

We first discuss reservation schemes in general. It is clear that if each terminal knew the number of packets waiting in queue at every other terminal at all times, it would be possible to schedule transmissions so that the channel would never be idle when some terminal had a packet ready to transmit; it also follows that at most one user would transmit in any slot. If reservations for these packets could be made without expenditure of bandwidth and with zero time delay, we would be able to obtain this ultimate performance limit for multiple access schemes. Such an idealized scheme is called perfect scheduling, and represents an ideal First-In-First-Out (FIFO) queueing system. However, in practice each reservation scheme suffers the consequences of allocating bandwidth for making the reservations, as well as the time delay associated with the reservation process.

In typical reservation schemes (such as Roberts' Reservation scheme [1]), each terminal first transmits reservation minipackets (which are much smaller than the actual information packets) to request the assignment of a time slot for each packet in its queue. Upon the successful transmission of a reservation minipacket, the corresponding information packet(s) will in effect join a common queue from which they may be transmitted, without danger of collision and according to whatever service discipline has been chosen (e.g., first come first served).

Under ideal conditions (i.e., a noiseless channel where all users are within communication range of each other and can thus monitor each other's reservation requests) no central controller would be needed. The Interleaved Frame Flush-Out (IFFO) protocols [2], for example, are a class of schemes in which transmission schedules are generated by the users in a distributed fashion, based upon the reservation requests of the population of users. All scheduling decisions could be made unambiguously by the users themselves in a distributed fashion because they would have the same data base after receiving the reservation transmissions. In many realistic situations, however, it cannot be assumed that all users have the same information. Inconsistent transmission schedules can then be generated resulting in the simultaneous transmission by two

or more users. Such collisions are usually assumed to destroy all packets that are involved. Therefore a central controller will usually be needed to allocate slots to the users requesting them to ensure that at most one user transmits in any slot.

In this paper we present, analyze, and evaluate a distributed reservation scheme that can in fact operate in an environment of noisy channels characterized by lack of complete connectivity. It is especially important that the control of a channel access scheme be distributed when multiple users are attempting to communicate with a central station that cannot communicate back to the users and thus cannot provide schedules. Such a situation can occur when the destination is required to maintain radio silence (either by operational doctrine or as a result of transmitter power limitations or malfunction) or if the link from the destination back to the multiple access users is jammed.

SYSTEM MODEL

We assume fixed length packets and a fixed length slotted frame structure, as shown in Figure 1. As is typical of slotted systems the slot duration is equal to the length of a packet, and all packet transmissions start at the beginning of a time slot. Each user is assumed to have at most one packet to transmit in any frame. Spread spectrum signaling is used, thereby permitting the use of CDMA techniques, which permit the correct reception of signals despite the simultaneous transmission by other users. A receiver is assumed to be able to monitor only one of these signals at a time, however, as in standard multiple access systems.




Fig. 1 Slotted frame structure.

We assume a contention-free reservation process. Such a scheme can be implemented if the size of the total user population is not too great by designating the first slot of each frame as a reservation slot and dividing it into TDMA minislots. Alternatively, a separate reservation channel could be implemented. Contention-based reservation procedures can also be considered. In such cases only the reservation minipackets (which are much smaller in length than message packets) are competing for channel access, and significantly higher data throughput can be maintained than in a purely contention-based channel access scheme. The impact on performance of contention in the reservation process is not investigated here.

We propose a reservation procedure that is quite different from conventional schemes. In the first version considered each user with a packet to transmit chooses one of the L slots in the frame at random. He sends a reservation minipacket that consists not only of a declaration

of intent to transmit in the coming frame, but in addition the actual slot number in which he will transmit. Since the users are uncoordinated it is possible for two or more of them to choose the same slot. The receiver, however, has full knowledge of the transmitters' intentions. In conventional "time-domain" schemes such interference would result in collisions that destroy each packet that is involved. We assume, however, the use of spread spectrum code division techniques that permit the simultaneous sharing of a wideband channel by a number of signals that use different quasi-orthogonal codes. Whenever two or more users declare their intent to transmit in the same slot it is up to the receiver to decide which of these signals it will in fact monitor. There is thus a successful transmission in any slot chosen by one or more users. By analogy to standard reservation schemes we sometimes use the term "assigned slot" to refer to the slot in which a user is successful, i.e., is being monitored, despite the fact that explicit assignments are not made. The ability to correctly receive one signal, despite the presence of others, results in considerable performance improvement as compared with conventional time-domain ALOHA-type schemes, as we shall demonstrate.

A sample realization of a frame of protocol operation is illustrated in Figure 2 for the simple case of frame length $L = 5$ slots and $M = 5$ users transmitting in the frame. Only the actual data slots are shown, and not the reservation slot (or subchannel). The slots chosen by the users have been shaded. Users #1, #3, and #5 have chosen slots #2, #1, and #5 respectively; all of these users are successful because they are the only ones to transmit in their respective slots. Users #2 and #4, however, have both chosen slot #4. Only one of these is successful; the decision of whom to monitor is left up to the destination.

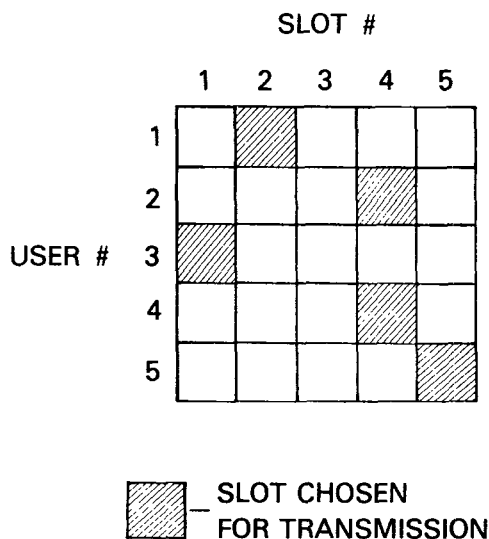


Fig. 2 Sample realization of a frame of protocol operation for $L = 5$, $M = 5$.

There are two basic ways to handle unsuccessful packets. They either may be retransmitted at a later time (e.g., in the next frame) or they may be simply dropped from the users' buffers. The former approach is the one most often taken in contention-based schemes. However, we have assumed that there is no feedback information of any type transmitted by the destination or among the users. The success or failure of individual transmissions cannot be determined, and so there is no information available on which to base a decision to retransmit. We thus assume that unsuccessful packets are dropped from the user's buffer, and therefore lost. The probability of packet loss is easily evaluated in the course of evaluating channel performance. (We note that the case in which acknowledgment information, but not scheduling, is available can also be modeled by the approach of this paper.)

The basic reservation scheme has been extended by permitting each transmitting user to designate several (say Q) slots in which he will transmit the same packet, thereby providing packet diversity. The destination node, after receiving all of these reservations, will attempt to generate a monitoring schedule that maximizes the number of distinct packets it receives correctly. Alternatively, a priority structure can be implemented in which certain users are more likely to be monitored, while others are monitored only if the high priority users are taken care of. In this paper we use the expected throughput, i.e., the expected number of successful packets per slot, as the performance criterion.

We can consider an extreme case in which each transmitting user transmits in all L slots of the frame. If the spread spectrum codes employed by each user were truly orthogonal, then such a scheme would provide optimum performance. However, only a quasi-orthogonality normally exists among such codes, and channel errors will result if too many signals attempt to share the channel simultaneously. The optimum value of the packet diversity parameter Q therefore depends on the degree of other-user interference that can be tolerated.

We first consider the case in which other-user interference is not troublesome (i.e., any number of other simultaneous users can be tolerated). One packet is thus successful in every assigned slot. We then introduce the effects of other-user interference by means of a simplified model in which a packet is never received correctly if the number of other users in its assigned slot is equal to or greater than some threshold, but always received correctly if the number is lower than the threshold. A more realistic model would be based on a probability of success that is a function of the number of users transmitting in the slot. Such a probabilistic model can easily be used with our model. The simplified model used here permits the presentation of the basic concept of operation, however, without undue complication.

ANALYSIS OF DISTRIBUTED RESERVATION SCHEME

We evaluate the conditional probability distribution of the number of successful transmissions in a frame consisting of L slots,

given that M users transmit packets, and each of these transmits its packet in Q slots chosen at random. We do not consider the mechanism used to make reservations; it is assumed that the destination receives error free reservation information from all users without incurring any overhead costs. We also do not consider the statistics of the arrival process. In this section we consider only the case of perfectly orthogonal codes; i.e., a packet that is numbered is received correctly regardless of the number of other users transmitting in the same slot. We later consider the case in which a packet is received correctly if and only if the number of other users transmitting in the slot does not exceed some threshold.

We first examine the case of $Q = 1$, i.e., each packet is transmitted once in the frame. This case can be treated exactly. One way to do so is to use combinatorial techniques to determine the probability distribution for the number of non-empty time slots (see e.g., [3]). We consider, however, the following Markov chain approach, which can be extended to obtain a lower (pessimistic) bound on system performance for $Q > 1$.

The problem is formulated as follows. We are given L slots and M users, each of which transmits in exactly one slot. The number of successful packets is equal to the number of slots in which one or more packets are transmitted. We approach this problem from the viewpoint of the M users, each of which in turn independently places a packet into one of the L time slots (rather than from a slot-by-slot viewpoint in which we would consider how many users transmit in each slot). As each user picks a slot we determine whether this slot has already been chosen by another user. Note that a limited priority mechanism can be implemented by considering the requests of higher priority users first; however, such priorities are implementable only among the users that choose the same time slot. Alternatively, fairness can be maintained by randomizing the assignments, or by granting access first to those who have not been successful for the greatest amount of time. We define the transition probability for the number of successes in the frame as the number of users is increased from k to $k+1$, for $1 < k \leq M-1$:

$$P(n|i) = \Pr(n \text{ successes by first } k+1 \text{ users} | \text{ given } i \text{ successes by first } k \text{ users}). \quad (1)$$

Clearly, the only possible transitions from i successes are to $n = i$ and $n = i+1$. An unsuccessful transition occurs if user $k+1$ chooses one of the i slots chosen by the first k users:

$$P(i|i) = i/L. \quad (2)$$

A successful transition occurs if user $k+1$ chooses one of the $(L-i)$ slots not chosen by the first k users:

$$P(i+1|i) = 1 - i/L. \quad (3)$$

We define,

$$p_j(i) = \Pr(i \text{ successes in first } j \text{ user attempts}). \quad (4)$$

This probability can be expressed in terms of the transition probabilities as,

$$p_j(i) = p_{j-1}(i)P(i|i) + p_{j-1}(i-1)P(i|i-1), \quad (5)$$

with initial condition $p_1(1) = 1$. The distribution for $p_j(i)$ is evaluated recursively until we obtain,

$$p_M(i) = \Pr(i \text{ successes in first } M \text{ user attempts}) \\ = \Pr(i \text{ successes in frame}). \quad (6)$$

For $Q > 1$ the distribution for the number of successful transmissions depends on the strategy used by the destination to determine whom it will monitor in each slot. It is difficult to evaluate this distribution for an optimal monitoring strategy. The case of $Q = 1$ was quite simple because the criterion for successful packet transmission was simply whether or not the user's slot had already been chosen. In Figure 3 we illustrate the difficulty for $Q > 1$ in extremely simplified form for the case of $Q = 2$. In this example, if the destination decides to monitor user #1 in slot #1 and user #2 in slot #2 then no assignment is possible for user #3. If, however, user #1 is monitored in slot #4, then slot #1 would be available for user #3. For large values of L , M , and Q it is considerably more difficult to create an optimum set of slot assignments (i.e., one that maximizes the number of successful packets). We have therefore considered a non-optimal scheme, that is amenable to analysis, which we describe as follows.




Fig. 3 Simplified illustration of the difficulty of slot assignment.

In the non-optimal scheme, which we consider for $Q > 1$, the destination assigns slots before he has complete knowledge of the reservations for all users. As in the case of $Q = 1$ we consider the transition probabilities as we add users, until a total of M users have been considered. The first user ($k=1$) is always successful. The destination chooses one of his Q slots at random; his remaining $Q-1$ slots are treated as empty slots. No effort is made to coordinate his assignment with those of the other users that follow him in sequence. Therefore some inefficiencies can result as discussed above. The analysis therefore provides a pessimistic estimate of the system performance as compared with that of a more intelligent decision maker. Note that a limited priority mechanism can again be implemented by making assignments in decreasing order of priority.

Note that the first Q users are always successful, even if they all choose the same set of Q slots. In general, user k will be successful if one or more of his Q slots has not already been assigned to another user. The destination randomly assigns one of these (not previously assigned) slots to him. Therefore,

$$P(i|i) = \begin{cases} \frac{i(i-1)\dots(i-Q+1)}{L(L-1)\dots(L-Q+1)}, & i \geq Q \\ 0 & i < Q \end{cases} \quad (7)$$

and,

$$P(i+1|i) = 1 - P(i|i). \quad (8)$$

We again use the recursion defined by eq. (5) to evaluate the probability distribution for the number of successful transmissions in the frame.

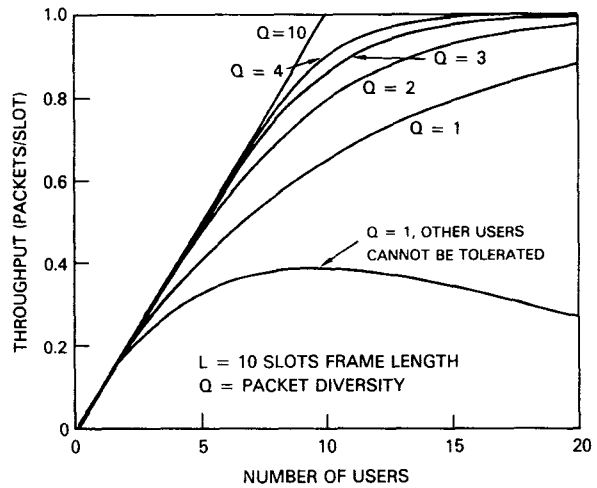


Fig. 4 Throughput performance of distributed reservation scheme; any number of other users can be tolerated in same slot.

PERFORMANCE EVALUATION

Case 1: Orthogonal Codes

We have used as a performance criterion the throughput, which we define as the expected number of successful packets received per slot. Any overhead caused by the reservation process is neglected, but could easily be added to our model for any of a number of specific reservation schemes. In Figure 4 we illustrate throughput as a function of the number of users M for packet diversity values of $Q = 1, 2, 3,$ and 4 . The frame length is $L = 10$ slots. These performance curves were generated under the assumption that the spread spectrum codes are in fact orthogonal, thus permitting other-user interference to be ignored. An upper bound on throughput (corresponding to the case of $Q = L$, i.e., all users transmitting in every slot) is also provided. Throughput of course increases as Q increases, with the most significant increase occurring as Q is increased from 1 to 2. The throughput of an ALOHA-type system (i.e., $Q = 1$, but transmission is

unsuccessful if two or more users transmit in the same slot) is also shown to illustrate the considerable improvement that is obtained as a result of the ability to tolerate other-user interference by means of spread spectrum CDMA signaling.

An alternative performance criterion, which we do not consider here, is the probability of successful packet delivery, or equivalently probability of packet loss.

Case 2: Quasi-orthogonal Codes

We now consider the effect of other-user interference on channel throughput. A model for such interference is presented in [4,5] for the case of frequency hopped signaling with Reed-Solomon coding used to correct errors caused by frequency hits (i.e., the simultaneous transmission by two or more users in the same frequency bin). The probability of correct packet reception in such systems depends on code rate, number of users, number of frequency bins, and availability (or lack of it) of side information (i.e., knowledge of which hops have been corrupted by hits).

In this paper we have used a simplified model for other-user interference in which a packet is never received correctly if the number of other users transmitting in the same slot is equal to or greater than some threshold, but always received correctly if it is lower than the threshold. While a more detailed probabilistic model can also be considered, it is felt that the model discussed here is adequate to illustrate the dependence of system performance on the degree of ability to tolerate other-user interference. Use of the model of [4,5] would require a detailed specification of system parameters, while we prefer to keep the scope of this paper quite general.

We also note that while it may be possible for the destination to change slot assignments to avoid slots with many users, our model assumes use of the same slot assignments obtained in the orthogonal code (no interference) case. In this sense our model is pessimistic, because an intelligent receiver might be able to make such a decision.

An exact system description is difficult to obtain. We would need the conditional joint probability distribution for the number of assigned slots in which *i* users transmit, for *i* = 1,2,...,M. We have simplified the model by assuming a Bernoulli transmission sequence of rate *Q/L* at each user that transmits in the frame in every slot. This Bernoulli model results in the same average number of users transmitting per slot as in the original system model. Under this assumption the probability distribution for the number transmitted in each slot is then

$$q(t) = \Pr(t \text{ users transmit in a slot} | M \text{ users})$$

$$= \binom{M}{t} \left(\frac{Q}{L}\right)^t \left(1 - \frac{Q}{L}\right)^{M-t} \quad (9)$$

We actually need

$$q(t|A) = \Pr(t \text{ users transmit} | \text{given } A, M \text{ users}) \quad (10)$$

where,

$$A = \text{event that the slot is assigned to some user.} \quad (11)$$

We make the simplifying assumption that the probability of a slot being assigned to some user (i.e, for a user to be monitored by the destination in that slot) is independent of the number of users transmitting in that slot, provided that at least one actually transmits. (The validity of this assumption will be examined in the future by comparison with a more detailed analytical model that is currently under development.) Therefore,

$$q(t|A) = q(t)/(1-q(0)). \quad (12)$$

The probability of successful packet reception, given that a slot is assigned, is equal to the probability that fewer than *T* users transmit in the slot, i.e.,

$$\Pr(t < T | M) = \sum_{t=1}^{T-1} q(t|A). \quad (13)$$

An approximation for expected throughput is obtained by multiplying the values presented in Figure 3 by this probability.

In Figures 5 and 6 we illustrate throughput performance for threshold values of *T* = 2 and 4, respectively, again for the case of *L* = 10. Figure 4, which represents the case in which other-user interference can be ignored, corresponds to *T* > *M*.

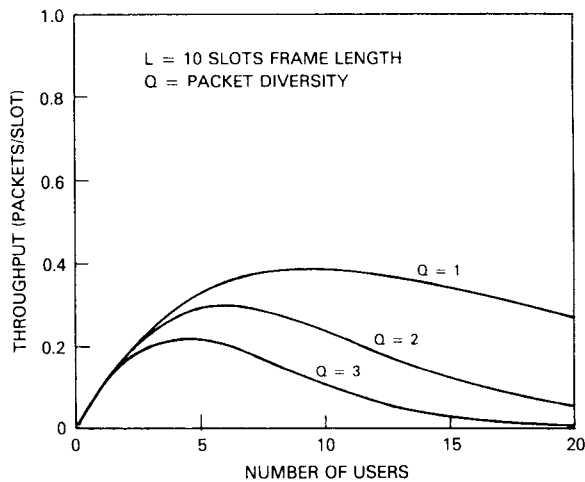


Fig. 5 Throughput performance of distributed reservation scheme; no other users can be tolerated in same slot.

We see that the optimum value of *Q* for a given threshold *T*, depends on the number of users. For the case of *T* = 2, however, in which the presence of one or more other users causes a packet error (as in ALOHA) a packet diversity value of *Q* = 1 is best for any number of users.

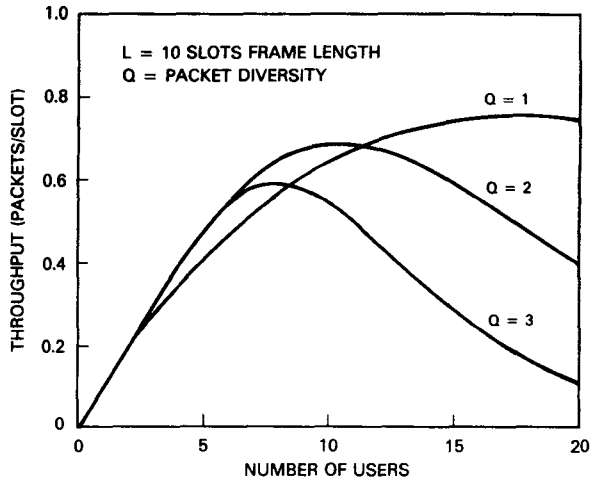


Fig. 6 Throughput performance of distributed reservation scheme; two other users can be tolerated in same slot.

In Figure 7 we illustrate throughput performance for $Q = 3$ as T is varied from 2 to 6. As T is increased the throughput increases, until it reaches the limiting case for $T \gg M$, which is in fact the orthogonal code case in which other user interference can be neglected.

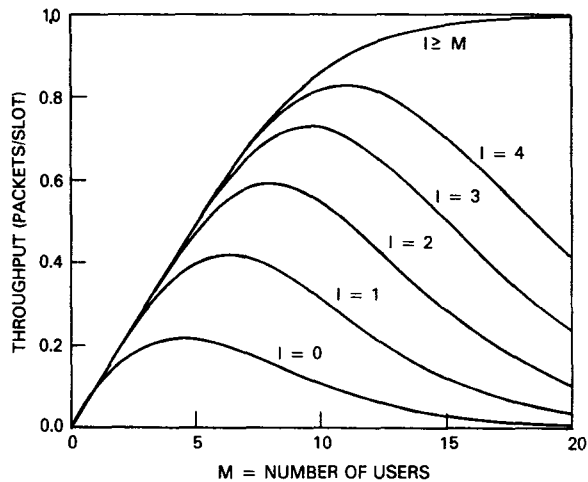


Fig. 7 Throughput performance of distributed reservation scheme for several levels of tolerable interference.

A NOTE ON ROBUSTNESS AND SURVIVABILITY

In our analysis we have assumed that the destination receives all reservations correctly, and is therefore able to construct a (not necessarily optimal) monitoring schedule. We now consider qualitatively the effect of incomplete or

erroneous reservation information. First, consider the case of missed reservations. Each user will transmit in the Q slots he has selected independent of whether or not his reservation is actually received. The destination will then simply create a schedule consisting of as many of the users it has received reservations from as is possible. The loss of one or more reservations will tend to make it easier to schedule those for whom reservations have been received, because there are fewer users to schedule. Thus, failure to receive reservations will usually adversely affect only those whose reservations are not successful, and not the remainder of the population. It is straightforward to model a system in which reservations were correctly received with some probability, rather than the perfectly reliable reservation mechanism assumed in this paper.

A crucial feature of the distributed reservation protocol is that the destination does not have to broadcast schedules, and can thus maintain radio silence. Therefore, in order to disrupt protocol operation one must disrupt the actual link from user to destination, since there is no (potentially weak) feedback or acknowledgment channel from destination to users. Another feature aiding survivability is the fixed frame length. One does not have to monitor either data traffic or control traffic to know frame boundaries.

AN EXTENSION TO MULTIPLE DESTINATIONS

We have thus far considered a single destination to whom all packets are directed. We now consider multiple destinations, as shown in Figure 8. We assume that some of the users communicate with more than one destination, and that not all users are within communication range of all destinations. As in the single destination case the use of multiple transmissions permits greater flexibility in slot assignment. It is certainly possible for one destination to monitor one transmission of a packet while another destination monitors one of the other redundant transmissions.

In contrast, we can consider a conventional centrally controlled reservation scheme (i.e., one in which the destination prepares and distributes schedules to the users) operating in a multiple destination environment. In such a system multiple transmissions from users would often be required (i.e., one to each destination), unless the destinations were able to coordinate their schedules, a task which requires the exchange of information among the destinations. Such coordination would have to be done each frame because of the assumed bursty nature of the traffic process, and in many cases would not be feasible. The distributed reservation scheme, on the other hand, is very well suited for communication from a population of bursty users to a group of geographically separated uncoordinated destinations.

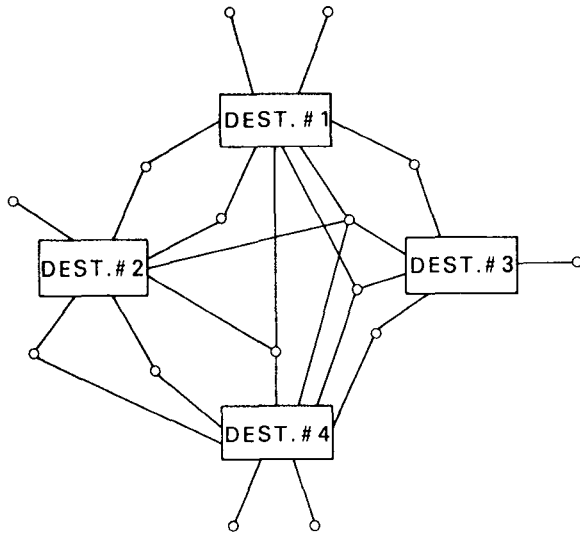


Fig. 8 Sample multiple destination geometry.

CONCLUSION

We have proposed, analyzed, and evaluated a new distributed reservation-based multiple access scheme for bursty users. This scheme is based on the quasi-orthogonality property that can be achieved through the use of spread spectrum signaling, which permits the correct reception of a signal despite the simultaneous presence of other signals. A fixed frame length of L slots is assumed. Each user with a packet ready for transmission picks Q (>1) slots at random out of the L . He sends a reservation consisting of these slots numbers, and then transmits his packet in these Q slots (the same packet each time). The destination, upon receiving the transmission schedules from all users, determines a monitoring schedule that attempts to maximize the number of users it actually monitors in the frame.

Performance depends heavily on the degree of orthogonality that can be achieved among the simultaneous channel users. We have used a simple threshold model to illustrate the dependence of achievable throughput on the number of other users that can simultaneously share the channel without resulting in packet errors. A high degree of robustness and survivability are provided by this scheme as a result of its distributed nature and its lack of need for acknowledgments or other forms of user coordination. Furthermore, it is easily applicable to the case of geographically separated multiple destinations.

REFERENCES

1. L. G. Roberts, "Dynamic Allocation of Satellite Capacity Through Packet Reservation," AFIPS Conference Proc., 1973 National Computer Conference, Vol. 42, pp.711-716, 1973.
2. J. E. Wieselthier and A. Ephremides, "A New Class of Protocols for Multiple Access in Satellite Networks," IEEE Trans. Automatic Control, vol. AC-25, pp. 865-879, October

- 1980.
3. W. Feller, An Introduction to Probability Theory and its Applications, Third Edition, revised 1970, John Wiley & Sons, Inc. New York.
4. B. Hajek, "Recursive Retransmission Control -- Application to a Frequency-Hopped Spread-Spectrum System," Proc. 1982 Conference on Information Sciences and Systems, Princeton, Univ., pp.116-120, March 1982.
5. M. B. Pursley, "Coding and Diversity for Channels with Fading and Pulsed Interference," Proc. 1982 Conference on Information Sciences and systems, Princeton, Univ., pp.413-418, March 1982.

Jeffrey E. Wieselthier was born in Brooklyn, NY, on March 14, 1949. He received the S.B. degree from the Massachusetts Institute of Technology, Cambridge, in 1969, the M.S. degree from the Johns Hopkins University, Baltimore, MD, in 1971, and the Ph.D. degree from the University of Maryland, College Park, in 1979, all in electrical engineering.

He was employed at the Naval Surface Weapons Center, White Oak, Silver Spring, MD, between 1969 and 1979. Since 1979 he has been with the Naval Research Laboratory, Washington, DC. His current research interests include multiple access, routing, and flow control in communication networks.

Code 7521
 Naval Research Laboratory
 Washington, DC 20375

Anthony Ephremides was born in Athens, Greece, in 1943. He received the B.S. degree from the National Technical University of Athens in electrical engineering in 1967 and the M.A. and Ph.D. degrees from Princeton University, in Princeton, NJ, in 1969 and 1971, respectively.

He is a Professor in the Department of Electrical Engineering, University of Maryland, College Park. During 1978-1979 he was on sabbatical leave at INRIA, France, and at the University of California, Berkeley. He has edited two volumes on random processes in the DH & R Series of Benchmark Papers. He has lectured at the CISM Summer School in Udine, Italy, and was a member of the U.S. delegation in the 1975 Information Theory Workshop, Moscow, U.S.S.R. He has been a member of the IDC and Associate Editor for Estimation of the IEEE Transactions on Automatic Control. He has served as the Treasurer, and since 1981 has been a member of the Board of Governors of the Information Theory Group of the IEEE. He is also a consultant at the Information Technology Division of the Naval Research Laboratory, and has taught several short courses in the Continuing Engineering Education Program of the George Washington University. His current research interests include routing, multiple access, and flow control in computer communication networks and statistical communication theory and applications.

Department of Electrical Engineering
 University of Maryland
 College Park, Maryland 20742

