Development of the ALOHANET

Invited Paper

NORMAN ABRAMSON, FELLOW, IEEE

Abstract—**The development of the ALOHA system packet broadcasting network at the University of Hawaii illustrates a number of general principles about the relationship between information theory and the design of real information systems. An anecdotal account is given of the development of the ALOHA system, emphasizing the interaction of the theory and the design of the experimental network. In addition, some preliminary results are described on the comparison of ALOHA channels and spread spectrum channels when used for packet communications.**

I. Introduction

IN THE LATE 1960's a number of efforts were in progress to use the existing worldwide telephone network to provide remote access to computer systems for terminals and, in some more ambitious cases, to provide connections among large information processing systems for resource sharing. The term "resource sharing" at that time was often taken to mean a sharing of hardware which would be considered primitive by today's standards. Nevertheless it was becoming apparent that the existing telephone network architecture was not well suited to the rapidly emerging data networking needs of the 1970's. Indeed it would have been surprising had such a network architecture, shaped by the requirements of voice communications at the end of the nineteenth century, been compatible with the emerging requirements of data communication networks at the end of the twentieth century. The original goal of the ALOHA system was to investigate the use of radio communications as an alternative to the telephone system for computer communications and to "determine those situations where radio communications are preferable to conventional wire communications" [1].

At that time the University of Hawaii was composed of a main campus in Manoa Valley near Honolulu, a four year college in Hilo, Hawaii, and five two year community colleges on the islands of Oahu, Kauai, Maui, and Hawaii, all within a radius of about 300 km from Honolulu. In September 1968 we began to plan for an experimental radio linked computer network which would be able to connect all of these locations together in order to permit sharing of the computer resources on the main campus. Even after the decision to build a radio data network, however, a number of basic design decisions dealing with the choice of network architecture arose.

From the beginning it was clear that the nature of the radio channel provided new system design options not available in systems using conventional point-to-point telephone channels. As the network planning proceeded, this key point assumed greater and greater significance. It would be gratifying to be able to report that this key difference between radio channels, with broadcasting or multiple access capabilities, and conventional point-to-point wire (or microwave) channels was appreciated as we began the project. Unfortunately, such an appreciation developed among the members of the project only with time; as is the case in many real-world situations, our foresight was not as clear as our hindsight.

Even at the beginning of the project, however, it was understood that the intermittent operation typical of interactive computer terminals was a convincing argument against the assignment of point-to-point channels in a conventional frequency-division multiple access (FDMA) or time-division multiple access (TDMA) manner. Some form of sharing of a common communication channel resource appeared necessary. The classical spread spectrum techniques of pseudonoise sequences or frequency hopping seemed to be one way of sharing channel resources, and a spread spectrum architecture was considered for the system. But the implementation of a spread spectrum format in each user's system at a useful data rate and at a reasonable cost seemed a formidable task in 1969, given the existing state of technology. Furthermore (with the exception of certain military applications), it was not clear then, and it is still not clear, what advantages could be obtained by the generation of a spread spectrum signal rather than the direct generation of a signal with a channel bandwidth considerably greater than its information bandwidth (see Section VI).

II. System Design

The first year of the project was spent in the study of the above questions, in the specification of a number of basic system parameters (e.g., frequency band, bandwidth, data rate) and in the organization of a laboratory for both the radio and the digital circuit parts of the system. In addition to myself, the faculty involved in the project at that time included. Thomas Gaarder, Franklin Kuo, Shu Lin, Wesley Peterson and Edward Weldon. The key decision to use the direct form of transmitting user information in a single high speed packet burst, now known as an ALOHA channel, was made at a meeting of the project participants in 1969. The operation of the channel was not well understood at that time, and the derivation of the capacity of

Manuscript received April 12, 1984.
The author is with the Department of Electrical Engineering, Holmes Hall 436, University of Hawaii at Manoa, Honolulu, HI 96822, USA.

that channel was to come only several weeks after the decision to use an ALOHA channel. Nonetheless, the criterion of design simplicity which led to that decision is often a reliable guide to designing data networks that work. The cancellation of plans for two large commercial data networks during the past few years, as well as the difficulties experienced by a well-known satellite data network [2, Chapter 20], may indicate that the converse of this "theorem" is equally valid.

The remaining system design decisions were easier to make. In retrospect, the only one which was crucial from the point of view of getting some sort of acceptable system operation was the choice of frequency band. That decision was made on the basis of some helpful advice from Dr. David Braverman of Hughes Aircraft Company. The ALOHA system was assigned two 100 KHz bandwidth channels at 407.350 MHz and 413.475 MHz and, after the usual unexpected software delays, the first ALOHA packet broadcasting unit went into operation in June of 1971. Formatting of the ALOHA packets as well as the retransmission protocols used in the ALOHANET were accomplished by a special purpose piece of equipment (designed by Alan Okinaka and David Wax) called the terminal control unit (TCU). A terminal was attached to the TCU by means of a standard RS232 interface and a user was connected to the central system at a data rate of 9600 bits/s anywhere within range of the radio system (about 100 km for the ALOHANET).

The design of the original TCU was constrained by the need to provide extensive debugging tools in this first packet broadcasting interface and by the technology of the times. Some understanding of the distance we have travelled since then may be obtained by considering how one key decision in the design of the original network was affected by the cost of memory. It was generally understood in the project that a full duplex mode of operation for the TCU would be desirable both from the point of view of system protocols and from the point of view of simpler hardware design. Full duplex operation, however, required the use of two independent buffers to store the packets flowing into and out of the TCU, and at that time the cost of memory for an additional packet buffer (704 bits) was about $300. It was felt that full duplex operation was not important enough to justify that expenditure, since we were interested in a network which might eventually contain hundreds of user stations. As a point of comparison, I have just purchased nine 64K RAM chips (841 times the 704 bits of an ALOHA packet) for $55.

During 1971 and 1972, as additional TCU's were built and a network came into existence, our understanding of the distinction between point-to-point channels and broadcast channels deepened; it became clear that the key innovation of the ALOHANET was not the use of radio communications for computers, but the use of a broadcast communications architecture for the radio channel. Furthermore, we began to understand that the advantages which we were seeing in our network could be obtained in any network with intermittent (bursty) transmitters as long as the network was built around a broadcast channel. The imminent launch of domestic satellites by the United States (Westar I was launched in 1974) suggested the use of satellite channels as broadcast channels, and we turned our attention to the use of ALOHA channels in satellite systems at about that time [3].

III. Satellite Networks

As one step in this process, it was decided to obtain a satellite link to connect the ALOHA system in Hawaii to the rapidly expanding ARPANET packet switching network on the U.S. mainland. The ALOHA system funding at about this time was switched to the Advanced Research Projects Agency (ARPA), under the direction of Dr. Lawrence Roberts, and the connection to the ARPANET seemed sensible for a variety of reasons. When the initial contacts with possible carriers were made in order to lease a 56 kbits/s satellite channel for the connection from the ALOHA system to a node of the ARPANET at NASA Ames Research Center in California, it was suggested that we lease twelve satellite voice channels under existing tariffs for voice grade channels, interface each of these channels with a modem operating at 4800 bits/s and multiplex the outputs to obtain the data rate desired. Such a procedure had obvious economic benefits to those providing the twelve leased satellite channels, but it also required that ARPA lease twelve voice channels to send only 56 kbits/s. The leasing of twelve voice channels seemed wasteful when we pointed out that the standard COMSAT technique for transmitting a single voice channel at that time was to transform the analog voice signal into a 56 kbits/s pulse code modulation (PCM) data signal. In other words, the existing tariff structure would have had us convert our 56 kbits/s data signal into twelve 4800 bits/s signals for transmission to the satellite earth station where these channels would have been transformed into twelve 56 kbits/s signals, transmitted to the other earth station as twelve 56 kbits/s signals, and transformed back into twelve 4800 bits/s signals for delivery to the network. These signals would then be multiplexed to form the 56 kbits/s data stream desired. It is a pleasure to report that common sense finally prevailed here and that, after some discussions with Roberts, a new 56 kbits/s tariff was introduced by COMSAT (which showed up in the COMSAT annual reports for several years as the "ARPA tariff") allowing the lease of digitized 56 kbits/s channels as data channels.

ALOHANET was the first commercial satellite link to employ a single satellite voice channel to transmit 56 kbits/s of data. However, this channel was not operated in a packet broadcast mode, but in a conventional point-to-point mode. The first network to utilize packet broadcasting in a satellite channel was put into operation in the ALOHA system in 1973, using the ATS-1 satellite in an experimental network which included the NASA Ames Research Center in California, the University of Alaska, Tohoku University in Sendai, Japan, the University of Electro-communications in Tokyo, and the University of

Sydney in Australia. This network, called PACNET, operated at 9600 bits/s in an unslotted ALOHA mode and used inexpensive earth stations to show the potential of data networks with large numbers of small earth stations [4].

Financial support for the research of the ALOHA system while the network was being built was provided by the Information Processing Techniques Office of ARPA, under the direction of Roberts. In addition to this financial support of the ALOHANET, Roberts also contributed to the success of the project in a way not ordinarily obtained from funding agencies. In a real sense, Roberts acted as another member of the research staff of the project, contributing a number of major technical results (including the first derivation of the capacity of the slotted ALOHA channel and the first analysis of the capture effect in ALOHA channels [5]).

Sometime in 1972, I was visiting Roberts' office in Washington for discussions dealing with both technical and administrative matters in the ALOHA system when he was called out of his office for a few minutes to handle a minor emergency. 1972 was a year of rapid growth for the ARPANET as the interface message processors (IMP's) which defined the nodes of the network were installed in the first network locations. While waiting for Roberts' return, I noticed on the blackboard in his office a list of the locations where ARPA was planning to install IMP's during the next six month period, together with the installation dates. Since I planned to bring up the question of installation of an IMP at the ALOHA system laboratory in Hawaii to be used with the satellite channel discussed above, I took the chalk and inserted "the ALOHA system" in his list and beside it placed the date of December 17 (chosen more or less at random). After Roberts' return, we continued our discussion but, because of the rather long agenda, we did not discuss the installation of an IMP in Hawaii, and I forgot that I had inserted an installation date of December 17 for us in the ARPA schedule on his blackboard. 1972 was a busy and productive period for those of us working in the ALOHA system, and I never did get the opportunity to discuss the installation of an IMP in our laboratory. Instead, about two weeks before the December 17 date, we received a phone call from the group charged with the responsibility of installing the IMP's asking us to prepare a place for the equipment. On December 17, 1972, an IMP connecting the ALOHA system to the ARPANET by means of the first satellite channel in the network was delivered and installed.

IV. Commercial Applications

By 1974, with the advent of the first microprocessors, it was clear that much of the logic of the TCU could be handled by a microprocessor, and several new terminal controllers, based on the INTEL 8080, were designed by Christopher Harrison and put into operation. Because of the fact that the network protocols could now be implemented in software, and this software could be simultaneously modified for all operating units by having the central station broadcast new protocol parameters or even a completely new protocol, these new controllers were called programmable control units (PCU's).

Once the basic protocols of the ALOHA channel were analyzed and demonstrated in the ALOHANET UHF radio channels, it was not long before other groups began to look into the possibility of designing packet broadcasting networks on other media as well. One of the first and most successful of these efforts was begun in the doctoral dissertation of Robert Metcalfe at Harvard University [6]. Metcalfe was probably the first person to appreciate the distinction between the communications architecture of circuit switched (or packet switched) data networks and the broadcast architecture of the ALOHANET. In his doctoral dissertation, he coined the term "packet broadcasting" to emphasize this distinction. When his dissertation was complete, he spent several months with the ALOHA system, working with Richard Binder who had developed our software and network protocols. Metcalfe then joined the Xerox Palo Alto Research Center, and his development of Ethernet (with David Boggs) at Xerox in Palo Alto demonstrated the effectiveness of packet broadcasting on a cable based medium.

The first use of packet broadcasting in an operational civilian satellite network took place, shortly after the Ethernet effort, in the COMSAT Marisat system. The request channel used to allocate voice and telex channels in Marisat required some technique capable of sharing a single medium-speed (4800 bits/s) channel among hundreds of possible users, and an unslotted ALOHA channel seemed to be the only possibility. Much of the theoretical foundation for both the cable based and satellite based systems was developed at about the same time by Kleinrock, Tobagi and Lam at the University of California, Los Angeles. Once these developments took place, it was astonishing how rapidly the acceptance of packet broadcasting techniques proceeded. By 1976, at least one book had been published which referred to the unslotted ALOHA channel as the "classical ALOHA channel" [7, p. 587].

The commercial use of packet broadcasting in UHF radio based systems, as accomplished in the experimental ALOHANET, took more time because of regulatory constraints on the use of frequency assignments. In the early 1970's only the first stirrings of the deregulation process which was to uproot the communications industry in the United States had taken place. With a more tolerant view of new methods of communication by the U.S. Federal Communications Commission, however, it has become possible to use new digital UHF radio assignments in packet broadcast mode rather than point-to-point mode. In January 1984 Motorola announced the introduction of the PCX personal computer employing an unslotted ALOHA channel in the UHF band at a data rate of 4800 bits/s. With the introduction of new broadcasting channel assignments in U.S. cellular radio telephone systems, we can expect to see some of these systems operate with a small number of these channels serving a large number of users in packet broadcasting mode. In addition, the broadcasting characteristics of the direct broadcasting satellites for television, now being put into operation, suggest that the advent of the personal earth station is not far off.

V. Strategic Theoretical Realities

The place of information theory and statistical communication theory in the design of real-world data networks is not always recognized. Yet the last few years have seen the announcement of a surprising number of major data networks that have subsequently been cancelled. In other cases, major data networks designed by large organizations with unquestioned capabilities in the technology of communications have encountered long and embarassing delays and difficulties. In each of these cases, it is possible to point to strategic management decisions which have been taken in the face of what might be termed strategic theoretical realities. Of these strategic theoretical realities, perhaps the most important is an appreciation of the basic capacity of the channels being used and the matching (or mismatching) of that capacity to the fundamental information rate of the signals which are to be transmitted. In the case of data networks, it is necessary to distinguish between the average data rate of the signal and the burst data rate. And in the design of networks meant to handle different kinds of signals from different sources, it is necessary to retain a respect for the enormous disparities in the information rates required by data, voice, and video. While these disparities, which can amount to as much as six orders of magnitude in information rates, may be well understood by the readers of this journal, the fundamental nature of this difference has not always been given due weight in management decisions dealing with network design.

Another of the strategic theoretical realities which can affect network design deals with the problem of scaling for large systems. For while the economies of scale which exist for large channels are generally appreciated, the diseconomies of scale which exist for large switches are sometimes forgotten. And in a network which uses switching and point-to-point channels for achieving connectivity among network hosts, the consequence of these two facts means that network growth may not simply be an exercise in scaling resources, but may require an examination of the total network architecture. A packet broadcasting network, on the other hand, achieves network connectivity by its broadcast architecture rather than by switching, and thus such a network can reap the benefits of the economies of scale of communication channels up to the point where the channel data rate begins to limit the cost of interfacing to the channel.

The importance of a theoretical understanding of the basic limits of operation for packet broadcasting channels is demonstrated by the articles in this special issue. The experience we have had with the ALOHANET has convinced me of the value of this kind of understanding as a guide to the resolution of many of the issues that arise in the design of real networks. But the converse of this observation may also be true—the operation of a real network can be a valuable guide to the selection of significant theoretical problems. As an example of this kind of synergism between theory and practice, consider the relationship between packet transmitters using conventional pseudonoise or frequency hopping spread spectrum multiple access and transmitters using packet broadcasting multiple access.

VI. Spread ALOHA

Spread spectrum communications is often defined as communications which uses a "bandwidth well beyond what is required to transmit digital data" [8]. Using this definition, we would include a packet broadcasting channel operating at a low duty cycle as just another form of spread spectrum [9]. The term "spread spectrum," however, is often not precisely defined and, when the signals in a spread spectrum system are bursty, some of the advantages ascribed to spread spectrum systems seem to follow from the time spreading as well as the frequency spreading of the transmitted signal. Properties such as the ability of spread spectrum signals to meet international spectrum allocation regulations and to minimize detectability [10] fall into this category. It is possible, however, to provide a trivial transformation of unslotted ALOHA channels to obtain the same properties. Such a transformation, corresponding to what we might call a "spread ALOHA" channel, is shown in Fig. 1.

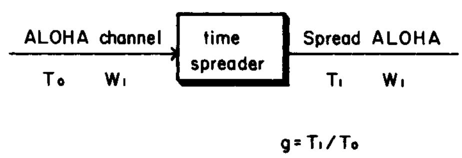

Fig. 1. Spread ALOHA transmitter.

In Fig. 1, the conventional unslotted ALOHA packets flowing out of the packet broadcasting transmitter are processed by a time invariant linear operator in order to spread the packet in time. The ALOHA channel is assumed to have a bandwidth W_1; the duty cycle G of the channel satisfies $G \ll 1$, and the packet duration is T_0. The input packets are given by

$$P_i(t) = \sum_{k=1}^{N} X_k p(t - kT - t_i) \tag{1}$$

where $p(t)$ is the pulse waveform, the X_k are the information symbols carried by the packet, N is the packet length, and $T = T_0/N$ is the pulse duration. One possible form for the output (spread) packets is given by

$$Q_i(t) = \sum_{k=1}^{g} A_k P_i(t - kT_0) \tag{2}$$

where the binary sequence A_k has the usual autocorrelation properties desired for a pseudonoise spreader in a spread spectrum system, and the time spreading factor g in (2) is analagous to the usual processing gain in spread spectrum systems [10].

We can compare the parameters of the spread ALOHA channel in Fig. 1 with the usual spread spectrum model in Fig. 2. The essential difference between the transmitters in the two figures is that in one case (spread spectrum) the signal spreading is caused by a multiplication in time (and thus a convolution in frequency), while in the other case (spread ALOHA) the signal spreading is caused by a convolution in time (and thus a multiplication in frequency). But the theoretical duality of this comparison should not be allowed to mask some important practical differences. For example, in the conventional spread spectrum channel, different transmitters would ordinarily use different spreading sequences with good crosscorrelation properties [11] to achieve multiple access. In the spread ALOHA case, however, only one pseudonoise sequence is needed in the system, since the separation required for multiple access is provided by the channel itself when $G \ll 1$. Moreover, both the spreading and the despreading operations for spread ALOHA are linear operations, rather than nonlinear operations as required for conventional spread spectrum [14].

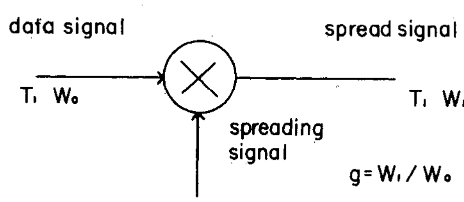

Fig. 2. Spread spectrum transmitter.

Finally, consider the asymptotic behavior of a conventional spread spectrum system and an ALOHA spread spectrum system when there are N packet transmitters sharing a channel and each transmitter has a duty cycle of G. Assume an additive white Gaussian noise channel for both cases with an average signal power per transmitter (while it is actively transmitting a packet) of P, a noise power spectral density of N_0 and a bandwidth of W. Assume for both cases that the packet length is long enough so that the asymptotic results of Shannon theory apply. Then, if we can treat the interference of the $N - 1$ other transmitters as additive white Gaussian noise [12], the capacity of a given transmitter while it is transmitting is

$$W \log\left(1 + \frac{P}{(N-1)GP + N_0 W}\right) \text{ b/s} \qquad (3)$$

(where all logarithms are taken to the base two).

We multiply this quantity by G to get the average data rate per user, and then multiply by N, the number of users, to obtain as the total capacity of the channel

$$NGW \log\left(1 + \frac{P}{(N-1)GP + N_0 W}\right) \text{ b/s.} \qquad (4)$$

Now if we let the number of users N increase, we get an asymptotic capacity for the total channel of

$$W \log e = 1.44W \text{ b/s} \qquad (5)$$

so that the capacity of the spread spectrum channel with a large number of users (the interference limited channel) is just proportional to the bandwidth of the channel. Note that the capacity of the channel is independent of the signal power, as we would expect for the interference limited channel, and that the capacity per user is inversely proportional to the number of users (again as expected).

We can contrast this result with an equivalent result for the conventional unspread ALOHA channel. In this case the data rate is limited by the noise and not by the interference since the original (unspread) ALOHA packets were assumed to have a low duty cycle. In [13, Sec. V], it is shown that the (Shannon) capacity of the low duty cycle average power limited ALOHA channel is the same as if the channel were being used in a point-to-point mode between only two users.

References

[1] N. Abramson, "The ALOHA System—Another alternative for computer communications," in *Proc. 1970 Fall Joint Computer Conf.*, AFIPS Press, vol. 37, 1970, pp. 281–285.

[2] R. D. Rosner, *Packet Switching: Tomorrow's Communications Today.* Wadsworth, 1982.

[3] N. Abramson, "Packet switching with satellites," in *AFIPS Conf. Proc.*, Nat. Computer Conf., New York, June 1973, vol. 42, pp. 695–702.

[4] ——, "Satellite data networks for national development," *Telecommun. Policy*, vol. 8, no. 1, pp. 15–28, Mar. 1984.

[5] L. G. Roberts, "ALOHA packet system with and without slots and capture," *Comput. Commun. Rev.*, vol. 5, no. 2, pp. 28–42, Apr. 1975.

[6] R. M. Metcalfe, "Packet communication," Massachusetts Institute of Technology, Cambridge, MA, Rep. MAC TR-114, July 1973.

[7] J. Martin, *Future Developments in Telecommunications*, 2nd ed. New York: Prentice-Hall, 1977.

[8] A. J. Viterbi, "Spread spectrum communications—myths and realities," *IEEE Commun. Magazine*, vol. 17, no. 3, pp. 11–18, May, 1979.

[9] R. A. Scholtz, "The spread spectrum concept," *IEEE Trans. Commun.*, vol. COM-25, No. 8, pp. 748–755, Aug. 1977.

[10] R. L. Pickholtz, D. L. Schilling, and L. B. Milstein, "Theory of spread spectrum communications—A tutorial," *IEEE Trans. Commun.*, vol. COM-30, no. 5, pp. 855–885, May 1982.

[11] R. Gold, "Optimal binary sequences for spread spectrum multiplexing," *IEEE Trans. Inform. Theory*, vol. IT-13, no. 4, pp. 619–621, Oct. 1967.

[12] M. B. Pursley, "Performance evaluation for phase-coded spread-spectrum multiple-access communication—Part I: System analysis," *IEEE Trans. Commun.*, vol. COM-25, no. 8, pp. 795–799, Aug. 1977.

[13] N. Abramson, "The throughput of packet broadcasting channels," *IEEE Trans. Commun.*, vol. COM-25, no. 10, pp. 117–128, Jan. 1977.

[14] J. Y. N. Hui, "Throughput analysis for code division multiple accessing of the spread spectrum channel," *IEEE J. Select. Areas Commun.*, vol. SAC-2, no. 4, pp. 482–486, July 1984.