

ALOHA PACKET BROADCASTING - A RETROSPECT

BY

R. BINDER, N. ABRAMSON, F. KUO, A. OKINAKA, D. WAX

AD A122775

DTIC FILE COPY

JANUARY 1975

THE ALOHA SYSTEM

TECHNICAL REPORT B75-41

UNIVERSITY OF HAWAII

HONOLULU, HAWAII 96822 28 090

APPROVED FOR PUBLIC RELEASE
DISTRIBUTION UNLIMITED

DTIC
SELECTED
DEC 29 1983

(1)


ALOHA PACKET BROADCASTING - A RETROSPECT

by

R. Binder, N. Abramson, F. Kuo, A. Okinaka, D. Wax
University of Hawaii, THE ALOHA SYSTEM, Honolulu, Hawaii

January 1975

Sponsored by
Advanced Research Projects Agency
ARPA Order No. 1956



The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Advanced Research Projects Agency of the United States Government.

To be presented at NCC 75, Anaheim, California

ALOHA PACKET BROADCASTING - A RETROSPECT

by

R. Binder, N. Abramson, F. Kuo, A. Okinaka, D. Wax
University of Hawaii, THE ALOHA SYSTEM, Honolulu, Hawaii

ABSTRACT

✓
In this paper we discuss the lessons learned in the design and implementation of the ALOHANET, a packet broadcasting radio network in operation at the University of Hawaii since 1970. The major part of the paper consists of a detailed discussion of the communications protocol choices that have evolved since the initial stages of the design of ALOHANET. Choices concerning the design of the radio communication subsystem are then examined, followed by an evolutionary view of the important impact that microcomputer technology has had on the user interface design and resulting system capabilities. The concluding section summarizes our present views with respect to the basic system configuration and properties of packet broadcasting networks.

11

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



This report was supported by THE ALOHA SYSTEM, a research project at the University of Hawaii, which is supported by the Advanced Research Projects Agency of the Department of Defense and monitored by NASA Ames Research Center under Contract No. NAS2-8590.

ALOHA PACKET BROADCASTING - A RETROSPECT

Contents

	Page
INTRODUCTION	1
THE ALOHANET	2
<i>Operation</i>	3
<i>ALOHANET Remote Units</i>	6
PROTOCOL CHOICES	7
<i>Random Access Channel Control</i>	8
<i>Broadcast Channel Queueing</i>	10
Priority Queues	11
Fairness	11
Turnaround Delay	13
Repeater Scheduling	13
<i>Packet Length</i>	14
Line Transmissions	14
Character-by-Character	15
<i>Addressing</i>	16
User Nodes	16
Repeaters	17
Resource Addressing	18
<i>Error Control</i>	19
Random-Access Channel	19
Broadcast Channel	20
Repeaters	22
Single-Channel Configurations	23
<i>Flow Control</i>	23
The Initial System	23
Some Terminal Complications	24
Satellite Complications	26
Multiple-Line Packets	27
Full Duplex Interaction	29
<i>File Traffic</i>	

RADIO SUBSYSTEM CHOICES	.31
<i>RF Channels and Modulation</i>	.31
<i>Radio Range</i>	.33
<i>Data Synchronization</i>	.35
USER INTERFACE CHOICES	.36
<i>The Original TCU</i>	.37
<i>Minicomputer Nodes</i>	.38
<i>Microprocessor Technology</i>	.39
<i>Size and Power</i>	.41
CONCLUSIONS	.41

LIST OF FIGURES

- 1 THE ALOHANET
- 2 BROADCAST CHANNEL MULTIPLEXING
- 3 BROADCAST CHANNEL FLOW CONTROL (Original System)
- 4 BROADCAST CHANNEL/ARPANET FLOW CONTROL
- 5 ALOHA SYSTEM UHF RADIO COMMUNICATION SYSTEM

ALOHA PACKET BROADCASTING - A RETROSPECT

by

R. Binder[†], N. Abramson, F. Kuo, A. Okinaka, D. Wax
University of Hawaii, THE ALOHA SYSTEM*, Honolulu, Hawaii

INTRODUCTION

Packet broadcasting is a technique whereby data is sent from one node in a net to another by attaching address information to the data to form a packet - typically from 30 to 1000 bits in length. The packet is then *broadcast* over a communication channel which is shared by a large number of nodes in the net; as the packet is received by these nodes the address is scanned and the packet is accepted by the proper addressee (or addressees) and ignored by the others. The physical communication channel employed by a packet broadcasting net can be a ground based radio channel, a satellite transponder or a cable.

Packet broadcasting networks can achieve the same efficiencies as packet switched networks,¹ but in addition they have special advantages for local distribution data networks² and for data networks using satellite channels.³ In this paper we concentrate on those characteristics which are of interest for a local distribution data network. In particular, we discuss the lessons learned in the design and implementation of the ALOHANET, a packet

[†] Now with BBN, Inc., Cambridge, Massachusetts.

* Supported by the Advanced Research Projects Agency of the Department of Defense and monitored by NASA Ames Research Center under Contract No. NAS2-8590.

broadcasting radio network in operation at the University of Hawaii since 1970. A number of design issues which arose in the construction of the system are defined, our solutions are explained, and in some cases they are justified. The lessons learned from the ALOHANET are used to indicate how such a radio packet broadcasting system might best be built using the technology available in 1975.

In the next section a brief description of the ALOHANET and its rationale is given. This is followed by a detailed discussion of the major system protocol choices that have evolved, pointing out some related theoretical work where appropriate. Choices concerning the design of the radio communication subsystem are then examined, followed by an evolutionary view of the important impact microcomputer technology has had on the user interface design and resulting system capabilities. The concluding section summarizes our present views with respect to the basic system configuration and properties of packet broadcasting nets.

THE ALOHANET

The ALOHANET is the first system which successfully utilized the packet broadcasting concept for on-line access of a central computer via radio. Its primary purpose is to provide inexpensive access to one or more time-sharing systems by a large number of terminal users, typically in the hundreds. However, it also allows user-to-user communication within the net and is evolving towards use in a more generally-oriented computer communications environment.

Operation

The present network configuration makes use of a broadcast channel for only one direction of traffic flow. (As we shall see in later sections, the lack of a broadcast capability in the other direction has seriously handicapped the development of effective protocols in certain areas.) Two 100 KHz channels are used in the UHF band -- a *random access* channel for user-to-computer communication at 407.350 MHz and a *broadcast channel* at 413.475 MHz for computer-to-user messages. The original system was configured as a star network, allowing only a central node to receive transmissions in the random access channel; all users received each transmission made by the central node in the broadcast channel. Recently the addition of ALOHA repeaters has generalized the network structure.

A block diagram of the present operational ALOHANET is shown in Figure 1. The central communications processor of the net is an HP 2100 minicomputer (32K of core, 16 bit words) called the MENEHUNE⁴ (Hawaiian for IMP) which functions as a message multiplexor/concentrator in much the same way as an ARPANET IMP.⁵ The MENEHUNE accepts messages from the UH central computer, an IBM System 360/65 running TSO (as of December 1974, a 370/158) or from ALOHA's own time-sharing computer, the BCC 500, or from any ARPANET computer linked to the MENEHUNE via the ALOHA TIP.⁶ Outgoing messages in the MENEHUNE are converted into packets, the packets are queued on a first-in, first-out basis, and are then broadcast to the remote users at a data rate of 9600 baud.

The packet consists of a header (32 bits) and a header parity check word (16 bits), followed by up to 80 bytes of data and a 16-bit data parity check word. The header contains information identifying the particular user so that when the MENEHUNE broadcasts a packet, only the intended user's node will accept it. More will be said about packet formats later.

ALOHANET




FIGURE 1 THE ALOHANET

The random access channel (at 407.35 MHz) for communication between users and the MENEHUNE is designed specifically for the traffic characteristics of interactive computing. In a conventional communication system a user might be assigned a portion of the channel on either an FDMA or TDMA basis. Since it is well known that in time sharing systems, computer and user data streams are bursty,⁷ such fixed assignments are generally wasteful of bandwidth because of the high peak-to-average data rates that characterize the traffic. The multiplexing technique that is utilized by the ALOHANET is a purely random access packet switching method that has come to be known as the *pure ALOHA* technique.⁸ Under a pure ALOHA mode of operation, packets are sent by the user nodes to the MENEHUNE in a completely unsynchronized manner -- when a node is idle it uses none of the channel. Each full packet of 704 bits requires only 73 msec at a rate of 9600 baud to transmit (neglecting propagation time).

The random or multi-access channel can be regarded as a resource which is shared among a large number of users in much the same way as a multiprocessor's memory is "shared". Each active user node is in contention with all other active users for the use of the MENEHUNE receiver. If two nodes transmit packets at the same time, a collision occurs and both packets are rejected. In the ALOHANET, a positive acknowledgement protocol is used for packets sent on the random-access channel. Whenever a node sends a packet it must receive an acknowledgement message (ACK) from the MENEHUNE within a certain time-out period. If the ACK is not received within this interval the node automatically retransmits the packet after a randomized delay to avoid further collisions. These collisions will limit the number of users and the amount of data which can be transmitted over the channel as loading is increased.

An analysis⁸ of the random access method of transmitting packets in a pure ALOHA channel shows that the normalized theoretical capacity of such a channel is $1/2e = 0.184$. Thus the average data rate which can be supported is about one sixth the data rate which could be supported if we were able to synchronize the packets from each user in order to fill up the channel completely. Put another way, this result shows the present 9600 bit/second channel could support between 100 and 500 active teletype users -- depending upon the rate at which they generate packets and upon the packet lengths.

ALOHANET Remote Units

The original user interface developed for the system is an all-hardware unit called an ALOHANET Terminal Control Unit (TCU), and is the sole piece of equipment necessary to connect any terminal or minicomputer into the ALOHA channel. As such it takes the place of two dedicated modems for each user, a dial-up connection and a multiplexor port usually used for computer networks. The TCU is composed of a UHF antenna, transceiver, modem, buffer and control unit.

The buffer and control unit functions of the TCU can also be handled by a minicomputer or a microcomputer. In the present system several minicomputers have been connected in this manner in order to act as multiplexors for terminal clusters or as computing stations with network access for resource sharing. A new version of the TCU using an Intel 8080 microcomputer for buffer and control has been built. Since these programmable units allow a high degree of flexibility for packet formats and system protocols, they are referred to as PCU's (Programmable Control Unit). A more detailed discussion of terminal considerations is given in a companion paper in these proceedings.⁹

Since the transmission scheme of the ALOHANET is by line-of-sight, the radio range of the transceivers is severely limited by the diversity of terrain (mountains, high rise buildings, heavy foliage) that exists in Hawaii. A recent development has allowed the system to expand its geographical coverage beyond the range of its central transmitting station. Because of the burst nature of the transmissions in the ALOHA channel it is possible to build a simple store-and-forward repeater which accepts a packet within a certain range of ID's and then repeats the packet on the same frequency. Each repeater performs identically and independently for packets directed either to or from the MENEHUNE. Two of the repeaters have been built which extend coverage of the ALOHANET from the island of Oahu to other islands in the Hawaiian chain. These repeaters are discussed in more detail in the following section.

PROTOCOL CHOICES

Two fundamental choices which have dictated much of the system protocol are the two-channel star configuration of the original network and the use of random accessing for user transmissions. Investigation of the random accessing principle using radio was in fact the original motivation for constructing the ALOHANET, while the two-channel configuration was primarily chosen to allow this investigation without complication from the relatively dense total traffic stream being returned to all users. An additional reason for the star configuration was the desire to centralize as many communication functions as possible at the MENEHUNE, minimizing the cost of the TCU at each user node.

Within this context, a number of protocol issues must be resolved. The more important of these are:

- random access channel control
- broadcast channel queueing
- packet length
- addressing
- error control
- flow control

Many of the original choices in these areas have undergone significant changes as a result of new user resources and user interfaces, or in some instances due to advancements in theoretical knowledge. The addition of repeaters has (potentially) a particularly significant impact on protocol.

We now discuss some of the considerations and resulting choices made in each of the above areas, with the impacts of new factors introduced within the context of each area. The section concludes with a brief discussion of the problem of integrating file traffic into the random access channel, a subject of current concern in the ALOHANET.

Random Access Channel Control

The retransmission strategy used in the random access scheme plays a central role in the scheme's effectiveness. Its determination directly affects the average delay experienced by users for a successful transmission, given a certain number of users accessing the channel, their traffic statistics, and the channel capacity. It can also be used to prevent the occurrence of channel saturation, a situation in which the channel becomes filled with

retransmissions and the number of successful packets falls to zero. These topics have only recently been quantified^{10,11} and remain subjects of current investigation.

One approach is to use different constant retransmission intervals at each node, with the intervals equal to integer multiples of the maximum packet transmission time to avoid subsequent conflicts. This results in a priority structure, since nodes assigned the longer intervals will experience a correspondingly longer average delay. As the number of nodes becomes large, however, unacceptably large delays result for the majority of users.

A strategy more appropriate for large user populations is to randomize the retransmission intervals used at each node (note that a priority structure can still be introduced if desired by using larger mean values for lower priority users -- in the remaining discussion, equal priorities will be assumed). According to recent results by Lam,¹¹ the resulting channel behavior appears to be relatively insensitive to the exact nature of the randomization, at least when comparing the use of uniform and geometric distributions. In any event, the cost of implementing a particular distribution at each node is an important design consideration. Based on initial estimates of the expected ALOHANET characteristics, a choice was made to use a uniform distribution. This allowed a relatively simple implementation in both hardware and software user nodes.

A simple technique was used in the original system nodes to achieve short delays when the channel is lightly loaded, while preventing channel saturation from occurring due to peak-hour loading or statistical traffic fluctuations: small retransmission intervals are used (relative to the intervals between new packets), but only for a maximum of three successive retransmission attempts.

If the third attempt is unsuccessful, the user is notified of a failure and must manually re-initiate the retransmissions. This in effect introduces a long interval between every three retransmissions, allowing time for retransmissions from other users to succeed. Based on a maximum packet transmission time of 70 milliseconds, the intervals are selected from a range of 0.2 to 1.5 seconds, giving a mean of about 0.7 seconds (ten maximum packet times) per retransmission. The lower bound is chosen to allow sufficient time to receive an ACK from the MENEHUNE if the packet was sent successfully, avoiding unnecessary retransmissions. (This time is based on a direct user-MENEHUNE path; if repeaters form a part of the radio path, the lower limit must be increased accordingly.)

The newer programmable PCU's in the system offer the capability of a more flexible strategy, for example allowing the interval used after each third retransmission to be automatically inserted. The use of different strategies, such as continuously increasing the time range used for selection of successive retransmissions, is also easily implemented by program; these and other strategies are currently under investigation.

Broadcast Channel Queueing

The MENEHUNE acts as a concentrator for the broadcast (F_2) channel, queueing waiting traffic when necessary for sequential transmission to user nodes. Four complicating factors exist, however: a need for priority queueing, fair allocation of the channel, the turnaround delay required by half duplex nodes, and the presence of repeaters.

Priority Queues

It is important that the F_2 channel data traffic not prevent the prompt return of an ACK to a user node, since this could lead to unnecessary user retransmissions and possible degradation of the random access (F_1) channel. Thus, an integral part of the F_2 channel multiplexing is the priority queueing mechanism maintained by the MENEHUNE, as shown in Figure 2. When even a transmission is completed on the F_2 channel the ACK queue is checked, and if not empty the ACK at the head of the queue is sent. Only when the ACK queue is empty is the data packet queue checked for waiting packets. This guarantees that at most one complete data packet plus any previously queued ACK's will be sent ahead of an ACK just placed on the queue. (Because the average rate of successful arrivals on the F_1 channel is limited to one-sixth the rate of F_2 transmissions by the random access technique, the number of previously queued ACK's will be zero most of the time.)

Fairness

A second problem is the possible hogging of the F_2 channel by one or a few users. This problem is eliminated by the queueing discipline used for the data packet queue. Only one packet per user is allowed on the queue at any time, and the queue is serviced on a first-come-first-




FIGURE 2 BROADCAST CHANNEL MULTIPLEXING

served (FIFO) basis. The prevention of more than one packet per user on the queue is handled in conjunction with user flow control, discussed below.

Turnaround Delay

A delay function is used by the MENEHUNE to count off the time required by half-duplex user nodes to switch from a transmit to a receive state. The actual time is determined by the equipment type -- the original off-the-shelf equipment required 100 milliseconds due to its use of mechanical relays; approximately 10 milliseconds is counted off for newer equipment now in use.

Repeater Scheduling

The addition of repeaters to the system introduces a number of new problems into the F_2 channel, both because of radio range overlap and the nature of the repeaters themselves. The latter are store-and-forward devices; a packet which is to be repeated is first received and stored in its entirety, then transmitted on the same frequency on which it was received (preventing reception of a new packet during this time). In order to prevent the loss of a second packet destined to the same repeater, the MENEHUNE must therefore appropriately schedule the packets in its F_2 channel queues.

For efficient scheduling (i.e., to maximize channel utilization), the MENEHUNE must know the repeater routing paths for each user node. This function could thus become quite complicated or even not achievable, depending on the degree of dynamic routing used. Because of the small percentage of traffic currently handled by repeaters in the present ALOHANET, a very simple brute

force method is used: whenever a packet is sent which is forwarded by one or more repeaters, the MENEHUNE counts off sufficient time for it to be repeated once before beginning a new transmission to any node (knowledge of which packets are to be repeated is available from the user address, discussed below). This results in wasted channel capacity, but is not significant due to the capacity available in the system at present.

Packet Length

Three factors having an important impact on the system are the use of variable or fixed-length packets, the way packet length or the number of data bytes is indicated, and the maximum packet length allowed. The choices made must take into account the different traffic characteristics generated by line-oriented and character-oriented user-computer interactions.

Line Transmissions

Fixed-length packets were used in the initial system to simplify the design and construction of system hardware. The data packet length for both channels was chosen to allow up to 80 data bytes (640 bits), based on the user delays introduced by the 9600 bps channel data rates, the line length of the terminals in the system, and the line-oriented characteristics of the IBM 360/65 used as the central time-sharing system. An end-of-line (EOL) indicator consisting of eight zero bits was used within the packet to identify the end of actual data, where the latter was restricted to 7-bit ASCII with the eighth (parity) bit set to one. Since it was anticipated that many of the lines typed by users would be less than 40 characters, a second packet type was also defined which contained a 40-byte data field (a "Half-Packet"). This last step proved to

be a mistake -- the half-packet logic at each end of the link was a significant source of both hardware and software bugs.

The packet formats have since been changed to allow the use of variable-length packets with newer user nodes. An 8-bit count field is used in the packet header to indicate the number of 8-bit data bytes in the packet, with the data parity word immediately following the last data byte. In addition to eliminating the wasted channel capacity of the fixed-length packets, this also removes constraints on the data itself necessitated by unambiguous detection of the EOL indicator within the data stream. The 80 data-byte maximum has been retained for both channels, since it still appears to be a reasonable upper bound with respect to both the multiplexing delays introduced to either channel and node buffering requirements. This should not be construed as an indication that this length is optimal, however; as file-oriented messages are introduced to the total traffic and/or user node storage continues to become cheaper, a larger maximum may be desirable for one or both channels (for a given channel data rate and user response time constraints).

Character-by-Character

The increased flexibility provided by PCU's has allowed the introduction of a 'short' data packet in which a single data byte is sent in the header in place of the byte count, followed only by the header parity word. Although a use for this packet occasionally arises for interactions with line-at-a-time systems, its main use is with the character-oriented ARPANET computers now available to ALOHANET users.

The use of these character-oriented systems can have a considerable impact on the size and frequency of packets sent in the random access channel. This has an important consequence for the buffering strategy and choice of packet length used at each node: since a new transmission cannot begin until an ACK has been received for the last one, all characters typed by the user during the ACK waiting time should be sent in a single packet. Thus if communication delays tend to overlap inter-character generation times, the affected characters are accumulated at the originating node and sent (more efficiently) in a variable-length packet, without adversely affecting user-computer interaction.

A logical extension of this last strategy is to buffer all characters typed by the user at his node until one is typed which causes some action to be taken by the computer. If the appropriate set of action characters is known at the user node, this allows an optimum use of both channel capacity and system buffering without degrading the user-computer interaction. A scheme which allows this to be done in conjunction with echoing control is given by Davidson,¹² and is currently being introduced into selected ARPANET hosts. Its implementation cost in ALOHANET PCU user nodes appears reasonable, and is anticipated for use as its support by host computers becomes widespread.

Addressing

User Nodes

User addressing is determined by the radio channel configuration and associated multiplexing technique. Ignoring repeaters for the moment, the

two-frequency configuration used in the ALOHANET allows only a single destination in the random access channel (the MENEHUNE), and a single source in the broadcast channel (the MENEHUNE). Thus only the sender's address is required in the random access channel and only the destination address in the broadcast channel, which in both cases is the user address. Concentration of more than one user at a radio node is handled by permanently allocating a block of user addresses to the node, allowing user node multiplexing without introducing another level of addressing complexity to the system. The required address space is determined by the total number of users expected to be supported by the random access channel, and is 2^8 (eight header bits) for the present 9600 bps ALOHANET channel.

Repeaters

The use of repeaters in the system introduces some significant new factors to be considered in choosing an address scheme. Because of radio range overlap and the store-and-forward nature of the repeaters, problems can arise involving conflicts generated by two or more repeaters repeating simultaneously to the same destination, infinite repeating of the same packet (looping), and weak-signal operation due to multiple (but time-sequential) paths. In addition, the addressing scheme directly affects the MENEHUNE's ability to schedule transmissions in order to maximize broadcast channel utilization, as discussed in a preceding section. The ability to eliminate or minimize these problems depends on the degree of mobility desired for user nodes and/or the repeaters themselves.

Because of the small percentage of user nodes which currently require repeaters in the ALOHANET, a simple scheme is in use based on the hardwired

properties of the original repeaters built for the system. A block of user addresses is defined for each repeater, the latter repeating only those addresses in its block. The block assigned to a repeater two hops from the MENEHUNE is a subset of the block assigned to its first hop repeater. User nodes are constrained to operate within the geographic range of their 'assigned' repeater by this scheme, but the node's user address is easily changeable if a relocation becomes necessary. Since only one path choice exists between each user node and the MENEHUNE at present, the optimum path is selected by default. As the number of repeaters in use increases and existing units are replaced by programmable devices, a more flexible repeater addressing scheme is expected to be implemented.

Resource Addressing

This refers to the user's choices regarding which system resource he may communicate with. The system allows users to request a connection to the campus IBM 370/158, the ARPANET, or another ALOHANET user node. This is accomplished by sending special sequences of ASCII characters in the data portion of packets to the MENEHUNE, which may either be typed by a terminal user or automatically generated. If the requested destination is available, its identification is stored in a *Connection Table* entry for the requesting user in the MENEHUNE, and the user's address stored in a similar entry for the destination. All subsequent packets from the user are passed to the stored destination and conversely, until either end requests that the connection be broken.

Two exceptions exist to this connection table routing of packets. The first are commands intended for the MENEHUNE, such as the 'connect' and

'disconnect' above. The second is a capability which allows a user to send a single packet to another ALOHANET user independently of current connection table entries. The originating user simply types a special two-character ASCII sequence followed by the destination user's address (up to three ASCII digits), followed by the desired text.

Note that in the case of a connection to another ALOHANET node, the latter's address is also the resource address. If the node's resource can service more than one user at a time (such as might be the case for a specialized minicomputer or storage device), the present addressing scheme requires either that a block of addresses be allocated to the receiving node (as in the case of a concentrator for sending), or a sub-address be sent in the text portion of every packet. The block allocation suffers from rigidity in that resource addresses cannot be reused dynamically by different users, and does not appear desirable if many such addresses must be allocated in the system.

Error Control

Random-Access Channel

Two distinct error sources exist at the MENEHUNE receiver, the usual random noise and errors due to packet conflicts. Because of the high probability of errors due to conflicts at full loading of the random access channel, a very reliable error detection mechanism is required. To achieve this it was decided to use two 16-bit cyclic polynomial parity check words in each data packet, one following the header and a second following the data. The separate header parity check forms the basis for a highly reliable packet synchronization method discussed in another part of this paper; it also allows reliable

establishment of packet length and other information prior to processing the data portion of a packet. A single header bit is also used in conjunction with the parity check for sequence numbering, allowing the detection of duplicate packets by the MENEHUNE.

Broadcast Channel

Error control for broadcast channel data packets (MENEHUNE to user nodes) involves some special considerations. For efficient operation, the usual positive acknowledgement scheme in which the ACK's themselves are not acknowledged depends on a high probability of the ACK's being successfully received. However, an ACK sent from user nodes must compete with data traffic in the random access channel. At full channel loading each random access packet must be retransmitted an average of 1.7 times, which means each data packet or ACK must be sent a total of 2.7 times on the average before it is successfully received.* But in order to force retransmission of the ACK's, the data packet being acknowledged must also be sent an average of 2.7 times by the MENEHUNE -- even though it was received correctly the first time! The problem is compounded by the typically high ratios of computer/user traffic which exist for most interactive systems, resulting in many more ACK's than data packets in the random access channel. This problem was "resolved" for the initial implementation by simply not sending ACK's from user nodes. Because of the high received signal strengths at the nodes, a very low error rate was anticipated; considering also that user nodes

* This assumes ACK's and data packets are the same length; although the ACK's are in fact shorter, the resulting error rate is still very high compared to a typical conflict-free channel.

consisted only of human terminal users, it was decided that a simple error detection/user notification scheme would be sufficient.

However, this is in general not adequate when more sophisticated data transfer functions take place or significant error rates exist at user nodes. An example of the first case is the loading of programs into core storage of a minicomputer node, where manually initiated error recovery usually requires restarting the loading from the beginning of the file. In the second case, error rates can become appreciable when user nodes are located in weak signal areas caused by distance, multipath interference, or line-of-sight blocking, or in strong signal areas in which strong local noise sources also exist. To allow for these situations, an option which allows user nodes to send positive acknowledgements has been implemented. The scheme works identically to that for the random access channel, but is only used selectively with newer programmable nodes when required (it can be turned on or off by a command from the user node to the MENEHUNE). Its effectiveness is based on the relatively light existing channel loading of the system and its use by only a few of the nodes.

One solution to this problem when all traffic to user nodes must be acknowledged in a loaded random access channel is to use sequence numbering with a large modulus, sending an ACK only when the maximum sequence number is received. This approach suffers from the unpredictable nature of interactive user-computer traffic, however; if the last computer output prior to new user input is missed by the node, a potential deadlock situation is created until the user decides something is wrong and takes manual action. An additional mechanism can be used to circumvent this, such as using automatic timeouts at the user node or sending dummy traffic to the node to 'flush out' missed packets. However, the sequence numbers succeed only in reducing the

number of ACK's sent in the random access channel -- to eliminate the unnecessary repetitions of data packets from the MENEHUNE, it is also necessary to acknowledge the ACK. That is, the ACK sent by a user node is timed out and retransmitted until an acknowledgement for it is received, just as for data packets. If another packet is waiting for transmission to the node at this time, its transmission with the next sequence number constitutes the ACK to the ACK; otherwise, a short ACK-ACK packet is sent by the MENEHUNE. This can be easily shown to result in significantly less total channel overhead, at the expense of more complication in the node implementation.

Repeaters

We have so far ignored the effects of repeaters in this discussion on both random access and broadcast channel error control. The repeaters currently in use in the ALOHANET do not generate acknowledgements in either direction, resulting in only end-to-end acknowledgements between the MENEHUNE and user nodes as above (but with longer minimum retransmission timeouts). This choice was made for initial repeater simplicity; it has been shown analytically, however, that a hop-by-hop acknowledgement scheme is in general superior to an end-to-end scheme, at least in contexts such as ARPANET¹⁰ and the ARPA Packet Radio effort.¹³ Thus we expect to convert to a hop-by-hop scheme when the existing repeaters are replaced by programmable units and/or repeater traffic error rates require it; this area remains a relatively unexplored problem domain within the present ALOHANET implementation.

Single-Channel Configurations

Finally, we note that the problems discussed above concerning ACK's sent by user nodes in the random access channel are effectively non-existent if a single-frequency channel configuration is used (and propagation times are less than the shortest packet transmission times). If all nodes can hear the transmission of all other nodes, it is only necessary that nodes refrain from sending for an ACK packet time following the transmission of a data packet by any node, except for the intended receiver who sends an ACK (if appropriate) during this time. Thus ACK's are sent conflict-free, allowing a simple positive acknowledgement scheme to be used for all traffic. Note that packets sent by the MENEHUNE are treated exactly the same as packets sent by user nodes with respect to ACK's, thus also eliminating any effects due to asymmetric computer-user traffic ratios.

Flow Control

The Initial System

In the initial system environment of a single half-duplex time sharing system, model 33 teletypes, and hardwired user nodes which buffered only the line being displayed, flow control was a relatively simple matter. A user always received at least one output line from the time sharing system (IBM's TSO running on a 360/65) for each input line, and a prompt character when it was ready for more input. The bandwidth between the MENEHUNE and 360 and the latter's I/O response times are such that one or two MENEHUNE buffers are normally sufficient to support transfers of packets received

from the random access channel; in the unlikely event that no buffers are available when a packet arrives, the channel protocol guarantees its retransmission. Thus no explicit flow control was provided to prevent new packets from being sent by a user node. If the user sends one before the 360 is ready, the packet is discarded and a "WAIT" message returned to the user by the MENEHUNE (the status of each 360 connection is known in the MENEHUNE by information routinely passed from the 360).

Broadcast channel flow control was necessary, however, since each line (packet) sent to a (hardwired) user node must be completely displayed before a new line can be received. This was accomplished by the scheme shown in Figure 3, in which the control for each user node is centralized at the MENEHUNE. The latter counts off the required display time following transmission of each packet to a user, inhibiting further transmissions to that user until the time is up. To prevent 360 output from tying up MENEHUNE buffers while packets are being displayed, a handshaking flow control is used; the 360 sends only one line of output for each user, then waits for a *go-ahead* (GA) message with that user's address. The GA is sent by the MENEHUNE whenever a user's display time is up, resulting in at most one buffer required for each user (the MENEHUNE can also hold up acceptance of any packet from the 360 indefinitely until it has buffer space available). Note that this strategy also prevents any user from hogging the broadcast channel, since it allows only one packet per user in the channel queue,

Some Terminal Complications

The introduction of high speed CRT and hardcopy terminals to the system required an expansion of the MENEHUNE's flow control mechanism for




Figure 3 BROADCAST CHANNEL FLOW CONTROL (Original System)

the broadcast channel. A set of display rates was added, with the rate used at each user node stored in a permanent table in the MENEHUNE; a user can change the stored value for his node by typing a special command to the MENEHUNE at any time. The CRT terminals require an additional flow control mechanism to suspend output when the CRT screen has filled, allowing the user to signal when he is ready to proceed. Thus a screensize command was created which allows users to specify a screensize of between one and 99 lines (or an infinite screensize); this value is also stored in MENEHUNE tables for each user node. A counter is maintained for each user with a finite screensize specification and is updated for each line sent to the terminal; when the maximum is reached, the MENEHUNE suspends generation of the GA message until the user inputs a carriage return.

Satellite Complications

The next complication to MENEHUNE flow control processing was caused by the connection of the ALOHANET to the ARPANET. The latter involves a 50 Kbps INTELSAT IV satellite path connecting Hawaii to California; because of its long propagation time (approximately 0.25 seconds) and ARPANET flow control protocol, a large amount of buffering is required at the receive end of the link to support continuous display at higher speed terminals -- in particular, a 9600 bps terminal requires approximately a 1000-byte buffer. (Since in general CRT terminal users do not require continuous output at this rate, a smaller amount of buffering is in fact used.) This required a substantial increase in the size of the MENEHUNE buffer pool and a more complicated queueing structure to support the broadcast channel, since now more than one packet per user must in general be stored in the MENEHUNE during display at the user node.

To maintain the single-packet-per-user policy for the channel queue, a separate queue was created for each user to hold additional packets. The resulting flow control scheme is shown in Figure 4, where the GA's sent to the 360 in Figure 3 are now sent to the internal ARPANET protocol module. The maximum allowed size of each user queue is determined by the user's terminal rate and the available MENEHUNE buffer pool, and in turn defines the parameters used in the ARPANET flow control protocol.

Multiple-Line Packets

A second complication resulting from the ARPANET connection concerns the extra time required by some higher speed displays for certain characters such as carriage return (CR) and/or line feed (LF). Output from the 360 in the initial system contained such characters only at the end of a line (packet), allowing the transmission time and other inter-packet delays to provide any extra time required. However, many ARPANET computers are character-oriented, at times generating many CR and LF characters within a single packet. Thus it was necessary to provide a padding function in the MENEHUNE which inserts dummy characters or otherwise adds a display time delay after each CR or LF occurrence within packets destined for a higher speed (greater than 110 bps) terminal. This necessitates the splitting of packets whenever the maximum 80-byte packet length is exceeded, and in general involves a significant amount of additional processing per packet.




Figure 4 BROADCAST CHANNEL/ARPANET FLOW CONTROL

Full Duplex Interaction

A third complication arising from many ARPANET computers is their full duplex user interaction. Unlike the 360, users do not necessarily receive output in response to each input or an indication of when the computer is waiting for more input. Since no explicit flow control is provided for input from user nodes to the MENEHUNE, users are forced to either interact in a half duplex fashion (guessing as to when the computer has finished its output) or suffer occasional losses of input data and subsequent retyping. The latter can occur frequently with the hardwired TCU's, since they contain a single buffer which is used for both keyboard input and display; if computer output arrives while the user is typing, the typed characters are overwritten in the buffer by the received packet. The newer programmable user nodes now in the system provide full duplex buffering for the terminal, allowing a packet to be received and displayed without disturbing the keyboard buffer.

However, even if user nodes are completely full duplex a flow control problem exists for packets sent to the MENEHUNE. Unlike the case for the 360, users of full duplex hosts may generate successive input packets without receiving responses from the host computer. If the ARPANET or host computer or both slow down, an excessive number of buffers can become queued in the MENEHUNE on behalf of the user. Thus, to prevent user hogging of the buffer pool a count of the number of input buffers queued for each user is now maintained; when equal to the maximum allowed, arriving packets are discarded and a discard notification returned to the user.

File Traffic

The original ALOHANET design was based on a homogeneous population of terminal users generating bursty traffic into the random access channel. However, the connection of minicomputers and other terminals with memory has introduced at least two sources of non-bursty, or 'file', traffic. The first case occurs when users desire to transfer data from a paper tape or other storage media to a host computer. The second occurs when it is desired to transfer program-generated output from a minicomputer at a user node to a display device at a second user node (users can connect to other user nodes through the MENEHUNE in the same way as to the 360 or ARPANET). In either case the resulting traffic must be prevented from hogging or degrading the random access channel, and must also be constrained to the destination's acceptance rate.

The random access technique itself implicitly provides an anti-hogging mechanism, since retransmission timeouts can be used to decrease the user's average rate if conflicts occur. This does not provide for destination flow control, however, and is not necessarily an optimal solution for the random access channel. A second approach is the use of explicit flow control in the form of GA's sent by the MENEHUNE to the sending user node. This provides a solution to both problems at the expense of a small percentage of broadcast channel capacity. Since the MENEHUNE receives GA's from the user's destination, either explicitly from the 360 or ARPANET module or from its display time counting for another ALOHANET node, it can simply relay them to the sending node in a short control packet. This approach also allows centralized optimization of traffic in the random access channel by the MENEHUNE, and is the subject of current investigation.

RADIO SUBSYSTEM CHOICES

The design of the ALOHANET radio communication system required the balancing of a number of performance goals against various system constraints which are peculiar to the use of radio frequencies for data communication channels. These trade-off studies resulted in the selection of our RF channels and modulation method. The determination of operating ranges and the choice of a data synchronization method resulted from the basic channel and modulation selection decisions. In this section we will describe the primary issues related to RF channel selection, modulation design, radio range determination, and data synchronization design.

RF Channels and Modulation

The choice of radio channels for any communication system is a complex task, requiring the trade-off of many factors such as desired bandwidth, area coverage, spectrum availability, potential interference and noise sources, regulatory requirements, and equipment costs. In the case of the ALOHANET, a wide channel bandwidth was considered desirable for the random access channel since user nodes are required to send messages to the MENEHUNE at high peak data rates compared to their average data rate. Wide bandwidth was also deemed advisable for the broadcast channel due to the expected high traffic density from the MENEHUNE. The use of wide channel bandwidth tends to force the use of higher frequencies where spectrum crowding is less severe and the availability of bandwidth is greater. Crowded radio bands are undesirable not only from the standpoint of interference to other users but

also because of potential interference from them. Another disadvantage of lower frequencies is the higher probability of interference from man-made noise sources, particularly in an urban area where the ALOHANET has most of its terminals.

From the above considerations it can be seen that the system's communication requirements tend to emphasize the use of higher radio frequencies. The primary constraint on moving to even higher frequencies is equipment cost and radio range. Above 500 MHz equipment costs tend to escalate rapidly. Area coverage also becomes more difficult due to more pronounced shadowing effects of the radio waves by buildings and hilly terrain. (Above 30 MHz radio propagation tends to be limited to line-of-sight paths.)

Therefore, the 400 to 500 MHz UHF band was selected as the optimum for the ALOHANET radio frequencies. Reasonably priced commercial radio equipment was found to be available in this frequency region and radio band crowding was not severe in Hawaii. Initially, assignments in the 450 to 470 MHz mobile radio band were requested but were rejected by the FCC because of our wide channel bandwidth requirements. (The mobile radio channels are specified at about 15 KHz bandwidth, whereas we were requesting 100 KHz.) We were fortunate enough to receive assignments as an experimental service in the government UHF band of 406 to 420 MHz, where spectrum space was available.

Since most radio equipments available in the UHF bands use frequency modulation (FM), this type of modulation was selected for the RF channels. A slight variation was incorporated in the hardware design to minimize the interface problems between the radios and the data modems. This variation was the use of a subcarrier tone to carry the actual data modulation. This tone is phase-shift-keyed by the data and the resultant signal is used to modulate

the FM transmitter. This modulated tone is recovered from the FM receiver and fed to the demodulator of the modem. This modulation system is referred to as FM/DPSK to indicate frequency modulation by a differentially phase-shift-keyed subcarrier. (Differential phase-shift-keying is used to resolve the problem of received phase ambiguity.) The resultant configuration is shown in Figure 5.

Radio Range

The maximum operating distance between any terminal of the ALOHANET and the MENEHUNE (or a repeater) is specified as the system's radio range. This distance is primarily a function of a transmitter's radiated power, the receiver's sensitivity, and the attenuation of radio signal power for the given distance. Local noise conditions at the receiver location can also affect this distance, but for system planning purposes, range is usually calculated on the basis of some given propagation model. For line-of-sight paths, which exist at VHF, UHF, and higher frequencies, two different models are used depending upon local topographical conditions. In an urban area these paths are partially obstructed and suffer from multipath effects. A power loss proportional to $1/R^4$ is usually assumed for these conditions.¹⁴ Where paths are unobstructed and well clear of the local terrain, a spreading loss proportional to $1/R^2$ can be assumed. Receiver threshold sensitivity in the ALOHANET is defined as that receiver input power level which causes an average bit error rate of 10^{-5} . This bit error rate should provide a packet throughput reliability better than 99 per cent for full-length ALOHA packets.

Assuming a transmitter equivalent radiated power of 10 watts, a simple whip antenna at a user terminal, an elevated antenna at the MENEHUNE or repeater




Figure 5 ALOHA SYSTEM UHF RADIO COMMUNICATION SYSTEM

and a 3 microvolt receiver sensitivity, the radio range works out to about 17 miles in the urban area for the ALOHANET frequencies. Between repeaters and the MENEHUNE terminal, which have well-elevated antennae and good path clearances, the assumed $1/R^2$ model gives a maximum range of 290 miles. The use of high-gain omnidirectional antenna arrays at repeater sites extends these ranges. Tests conducted on a 100 mile path between two ALOHANET repeaters confirmed the $1/R^2$ spreading-loss assumption and indicated a fade margin of 30 db existed (due to the 10 db gain antennae used for the test).

Data Synchronization

Because of the burst nature of radio transmission of ALOHANET packets, special synchronization techniques must be employed in the modem and data terminal equipment. Since the phase-shift-keying used in the ALOHANET modem design is a bit-synchronous technique, bit synchronization must first be performed in the demodulator before packet synchronization can be attempted. Bit sync is performed by a phase-locking circuit, and a lock-indication signal is passed to the data equipment when bit sync has been attained. The bit-sync detection circuit is so designed to provide a very low false detection probability (less than 10^{-6}) and a high probability of packet detection. The narrow bandwidth of the phase-lock circuit presently designed into the ALOHANET modem requires a bit sync preamble of 90 bits to ensure reliable bit sync. Studies have indicated that this preamble can be reduced to about 10 bits by use of a redesigned wide-band phase-lock circuit. In fact, we are presently contemplating doing away with the bit-sync preamble entirely, further reducing packet overhead. The unique characteristics of the ALOHA modem design make such an approach feasible.

Packet synchronization is accomplished in the ALOHANET data terminal buffer by means of the 16-bit parity word contained in the packet header. When the parity check routine accepts the header, the packet is assumed to be synchronized. Since the parity check routine is initiated by the first bit of the header, packets can be missed due to detection of an early error bit before the header. This miss probability is presently controlled by the modem at about 10^{-3} or less, providing a packet detection probability of 99.9 per cent or better. The false detection probability of this circuit is $\sim 1.5 \times 10^{-5}$, which is independent of that of the modem. Thus, the overall probability of false detection is less than 1.5×10^{-11} . Therefore, less than one out of a thousand packets will be lost due to packet sync errors and packet sync false alarms occur with extreme rarity.

USER INTERFACE CHOICES

The development of the ALOHANET user interface has been an evolutionary process, as is typical of most research developments. Since there were expected to be many user nodes (as compared to the single MENEHUNE node), the primary design goals were initially set as simplicity of design and low cost. This led to the design of a hard-wired control unit with limited data storage capability coupled to a modem and radio transceiver. This initial design was termed a Terminal Control Unit (TCU). As experience developed with operation of the net, other functions became evident as being desirable in a TCU. At about this time the first microprocessor chips and low-cost semiconductor memory chips were becoming available in the marketplace. It was decided that a new TCU design should be initiated using these new devices since much greater flexibility and additional functions could be readily incorporated in

a unit having a capability of being programmed. It was also noted that the cost of these new devices was such that a unit could be built for the same cost or less than that of the original design. Thus, the Programmable Control Unit (PCU) was developed, and there are now several operating units in the system. We will now discuss some of the issues involved in designing a terminal control unit for use on the ALOHANET. These issues lie in the general areas of interface considerations, and the technology of microprocessors.

The Original TCU

The ALOHANET was originally envisioned as a terminal network, with the TCU's interfacing human users to a half duplex, line-oriented time-sharing system. At the time of the first TCU design effort memory was relatively expensive; so in order to minimize cost a single buffer was chosen for use with both the terminal keyboard and display. (As noted earlier in this paper, when full duplex computer interactions were available in the system the single buffer was found to be quite a disadvantage.) The buffer was designed for a full line length of 80 characters, which allowed handling of both the 40 and 80 character fixed-length packets defined for the system.

Additional basic functions performed by the TCU's were generation of a cyclic-parity-check code vector and decoding of received parity code words for error-detection purposes, and generation of packet retransmissions using a simple random interval generator. If an acknowledgement was not received from the computer after the prescribed number of retransmissions, a flashing light was used as an indicator to the human user. Since the TCU's did not send acknowledgements to the MENEHUNE, a steady warning light was displayed to the human user

when an error was detected in a received packet. Thus it can be seen that considerable simplification was incorporated into the initial design of the TCU, making use of the fact that it was interfacing a human user into the network.

Other functions hardwired into the TCU were the obvious requirements of checking for and generating its address, packet sequence numbering, checking to see if a received packet is an ACK packet or a data packet, and generating and checking for half- or full-packet conditions. (The control bits for these functions all reside in the header portion of the packet.)

The final consideration was the choice of standard interface signals between the TCU and the user's equipment. This was a relatively simple choice, since most equipment is designed to meet the EIA standard RS 232C interface specification. Therefore, the TCU was designed to meet this standard, which allows direct connection of most terminals in use today.

Minicomputer Nodes

As the ALOHANET developed, some minicomputers were interfaced into the network as concentrators for a number of terminals. Many of the logical functions performed in a TCU were now incorporated into the mini's software, with error detection and parity word generation performed in a special hardware interface unit imposed between the minicomputer and an ALOHA modem. (This unit was very much like the encoder/decoder unit used at the MENEHUNE to interface that minicomputer to the channel.) Parallel-to-serial and serial-to-parallel conversion was also performed in this interface unit.

However, a minicomputer is an expensive device to use for these simple functions, and it requires considerable amounts of power and space. If it already

exists for the purpose of performing various user-oriented tasks, then it is cost-effective to incorporate the software interface and a minimal amount of hardware for use on the ALOHANET.

The advent of the microprocessor chip changed all this. The relatively low-cost processing power demonstrated by these units made it apparent that many system options we had previously considered and discarded because of hardware complexity and cost limitations in the TCU, were now viable in a PCU. Some of these options -- file transfer, remote user ACKs, single frequency operation, character-by-character transmission -- were discussed in previous sections. This trend toward programmable and more powerful TCU's has thus led to the development of the ALOHA PCU, using a microprocessor to handle the TCU buffering and control functions, in addition to more complex and sophisticated functions.

Microprocessor Technology

The development from the hardwired TCU concept to the fully-programmable PCU has closely followed the rapidly changing technology of microprocessors. The availability of lower-cost semiconductor memory has allowed the evolution from half-duplex to full-duplex operation in the PCU, with the beneficial side-effect of decreased logical complexity due to separation of the input and output functions. However, the first PCU developed had a hardware complexity level comparable to the TCU due to the relatively primitive structure of early microprocessor designs. This first PCU, designed with the Intel 8008 CPU, required a considerable amount of circuitry for buffering and multiplexing functions needed with this early microprocessor chip. Because of the slow speed of the chip, bit-by-bit processing was not possible and additional buffering was

also necessary. But, much greater flexibility was introduced into the scope of functions which could be performed, due to its programmability.

Later microprocessor designs, such as the Intel 8080 and National IMP-16, have introduced much greater sophistication into the processor chips accompanied by significant processing speed improvements. A newer PCU design, incorporating an Intel 8080 chip, has demonstrated a considerable reduction in hardware complexity accompanied by an even greater degree of processing flexibility. For example, parity generation and checking are done in software with this prototype design.

Buffering has progressed from the simple shift-register storage devices of the TCU to the use of semiconductor RAM devices used in the microprocessor's random-access memory. All of the micro-instructions for the Intel 8080 microprocessor PCU design reside on four PROM chips, providing 1024 bytes of microcode. The random-access memory consists of 2048 bytes of RAM.

Recent product introductions such as Intel's 3000 series bi-polar chips promise even greater reductions in chip counts and increases in processing power and speed. With machines such as these, bit-by-bit processing can be readily incorporated into software, thus further eliminating the need for external interfacing hardware and simultaneously providing greater flexibility in the implementation of additional functions. A more detailed discussion of communications microprocessors is given in a companion paper in these proceedings.⁹

Size and Power

In the earlier versions of the TCU smaller size and power drain of the unit were not considered major design objectives. The first units were designed for ease of access and hardware modifications to these TCU's were made on a fairly casual basis. As more and more of the ALOHANET came into use, however, small size, portability and lower power drain became desirable.

Of particular interest is the possibility of designing low power battery operated portable PCU's for mobile units in the ALOHANET. Since the transmitter power need only be on for a short burst corresponding to the period of the data burst, the average power of the transmitter can be a small percentage of the peak power. Since low power and small size were not original design objectives, it appears that the construction of low power portable PCU's will involve redesign of several subsections of the PCU and some new design efforts. Of particular importance is selection of a microprocessor unit which provides a minimum power-drain computer architecture consistent with functional requirements. The modem should be redesigned to use MOS devices to minimize power drain, and the transceiver designed for minimum complexity.

CONCLUSIONS

As the system has been modified during the past several years it has become apparent that packet broadcasting architecture is remarkably flexible in its tolerance of hardware, system and protocol modifications. This flexibility follows from the packet verification algorithms which lie at the basis of packet broadcasting. The only packets accepted by a remote unit or by the MENEHUNE are packets which meet all the tests expected by

the potential acceptor; and the only system resource consumed by an unaccepted packet is the capacity of the channel during the short burst of the packet duration. Thus it is perfectly feasible in a packet broadcasting network to introduce a new form of packet (new in format, new in packet length, or even new in modulation technique) without disturbing any unit operating with the existing scheme. Only the units designed to look for the new packets will accept these packets and all other units will simply discard them.

We plan to employ this property of packet switched channels to switch the polynomial used for error control in the present packet format. The new polynomial is available in a single IC chip and will allow the possibility of error correction as well as error detection in some cases. As remote units with new packet formats are put into operation we can continue to operate the existing remote units without modification as long as we have a single unit capable of accepting the new packet format at the MENEHUNE. As a side benefit of the introduction of this modification we also note that we have effectively doubled the number of user addresses in the system. An address in use with the old packet format may be reused with the new, since each is effectively invisible to the other.

Another result of our ALOHANET experience, current technology, and recent theoretical work on ALOHA channels, is that a *single-channel* network configuration appears preferable to the two channels used in our present system. The major reason why this is so has to do with the broadcast property of the single-channel system, in which all nodes can (for a given geographic range) hear the transmission of all other nodes in the net.

A number of desirable properties result from this broadcast feature. First, each node can determine if the channel is free before transmitting, greatly reducing the number of packet conflicts -- Kleinrock and Tobagi¹⁵ have shown analytically that this can increase the throughput of a random access channel by a factor of three to five for reasonable user delays, depending on the propagation times between nodes. Second, the problem of sending acknowledgements from user nodes is resolved in a simple manner. Third, system bandwidth can be optimally allocated to both directions of traffic by simple time-sharing of the channel. Fourth, single channel repeaters require only half the radio hardware of two-channel repeaters, and, in fact, the radio transceivers at all nodes need be only half duplex. Finally, a single-channel system constitutes a fully-connected network allowing direct communication between all nodes. A star configuration can still be imposed by protocol to direct all user traffic through a central node, but is no longer required.

It is important to note that many of the above properties are made feasible by the availability of PCU's at a reasonable cost through microcomputer technology. This raises a related issue: the desirability of distributing presently centralized protocol functions such as flow control among the user nodes. Since we have just begun to gain experience with PCU's in a packet broadcast network, we must leave this as an open question.

REFERENCES

- 1 L G ROBERTS B WESSLER
The ARPA computer network
In Computer-Communication Networks edited by Abramson and Kuo
Prentice-Hall 1973
- 2 R E KAHN
The organization of computer resources into a packet radio network
NCC 1975 In this issue
- 3 N ABRAMSON
Packet switching with satellites
NCC 1973 pp 695-702
- 4 R BINDER W S LAI M WILSON
The ALOHANET MENEHUNE - version II
ALOHA SYSTEM Technical Report B74-6 University of Hawaii September 1974
- 5 F E HEART R E KAHN S M ORNSTEIN W R CROWTHER D C WALDEN
The interface message processor for the ARPA computer network
SJCC 1970 pp 543-567
- 6 S M ORNSTEIN F E HEART W R CROWTHER H K RISING S B RUSSELL A MICHEL
The terminal IMP for the ARPA computer network
SJCC 1972 pp 243-254
- 7 P E JACKON C D STUBBS
A study of multi-access computer communications
SJCC 1969 pp 491-504
- 8 N ABRAMSON
The ALOHA System
In Computer-Communication Networks edited by Abramson and Kuo
Prentice-Hall 1973
- 9 S C FRALICK D H BRANDIN F F KUO C HARRISON
Digital terminals for packet broadcasting
NCC 1975 In this issue
- 10 R M METCALFE
Packet communication
Project MAC Technical Report MAC TR-114 MIT July 1973
- 11 S S LAM
Packet switching in a multi-access broadcast channel with application to satellite communication in a computer network
UCLA-ENG-7429 Report University of California at Los Angeles Computer Science Dept April 1974

- 12 J DAVIDSON
An echoing strategy for satellite links
NIC Document 10599 (RFC 357) Stanford Research Institute June 1972
- 13 H FRANK R M VAN SLYKE I GITMAN
Packet radio network design - system considerations
NCC 1975 In this issue
- 14 Y OKIMURA E OHMORI T KAWANO K FUKUDA
Field strength and its variability in UHF and VHF land-mobile radio service
Review of the Electrical Communication Laboratory Vol 16 No 9-10
Sept-Oct 1968
- 15 L KLEINROCK F TOBAGI
Packet switching in radio channels: carrier sense multiple-access modes and their throughput delay characteristics
to be published in IEEE Transactions on Communications