

## A Broadband Local Area Network

By A. N. NETRAVALI\* and Z. L. BUDRIKIS†

(Manuscript received April 18, 1985)

The IEEE 802 standard for local area network based on Carrier Sense Multiple-Access with Collision Detection (CSMA/CD) operates at a peak rate of 10 Mb/s on a cable of maximum length 2.5 km using baseband signaling. In many situations, larger channel rates are required over a much larger area. However, the efficiency of the CSMA/CD access method decreases rapidly if either the length of the cable is increased for a fixed bit rate or if the bit rate is increased for a fixed cable length. In this paper, we propose a broadband network for computer communications containing several CSMA/CD-type systems, each operating in a different frequency band. In addition, in order to have a wide area access, while minimizing the loss of performance associated with large collision delays, terminals in a small given geographical area are given one of the frequency bands for transmission. Two access protocols are developed. Using these schemes, it is possible to increase the channel throughput and the access area and to reduce the collision delay. We present a simplified analysis to quantify the improvement in performance using our schemes.

### I. INTRODUCTION

Local Area Networks (LANs) share computing and other resources among many users and, if properly designed, increase the reliability by reducing the dependence of a user on one processing unit or a peripheral. Unlike long-haul networks, where channel utilization has to be optimized owing to high cost of communication over long

---

\* AT&T Bell Laboratories.

† Visiting from the Department of Electrical and Electronic Engineering, University of Western Australia, Nedlands, Australia.

---

Copyright © 1985 AT&T. Photo reproduction for noncommercial use is permitted without payment of royalty provided that each reproduction is done without alteration and that the Journal reference and copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free by computer-based and other information-service systems without further permission. Permission to reproduce or republish any other portion of this paper must be obtained from the Editor.

distance, local networks use bandwidth somewhat extravagantly to reduce the switching costs. Several network topologies, such as rings, buses, and trees, have been proposed along with access methods such as carrier sense, token passing, etc.<sup>1,2</sup>

The IEEE 802 standard for local area networks uses CSMA/CD (Carrier Sense Multiple-Access with Collision Detection) as one of its access methods.<sup>3</sup> It uses baseband transmission on coaxial cables (although other media are possible) at a peak rate of 10 Mb/s. For a variety of reasons, length of the cable (and therefore length of each segment of the network) is limited to 2.5 km. Within the limitations of the above parameters, the CSMA/CD-based access method provides an efficient means of computer communication for low loads on the channel. However, if the channel loading is increased, or if the requirements dictate either higher bit rates or longer cable lengths—for example, to serve a metropolitan area—there is considerable loss of efficiency. Much of this inefficiency comes from the use of the CSMA/CD protocol. In CSMA/CD, a source transmits a packet when the channel is sensed as idle, but this injection of the packet can be known to the other sources only after it has propagated throughout the length of the cable, during which time another source may attempt to transmit on the channel. Thus, the number of bits wasted due to collision is proportional to the propagation delay and the peak bit rate. Also, the need to detect collisions makes it necessary that each packet have a duration equal to at least the round-trip delay. With very large nets and high bit rates, that may represent an unreasonably large minimum number of bits.

Baseband CSMA/CD has been extended to broadband CSMA/CD by several CSMA/CD networks, each in a different frequency band put on the same cable (see, for example, Ref. 4). However, each of these networks operates almost independently, connected usually by a signaling channel. Also, the cable length limitation still applies, making it difficult to use for a metropolitan area. In this paper, we propose schemes that extend the capabilities of both the baseband and the broadband CSMA/CD networks by allowing higher bit rates on a cable, larger cable segments, and at the same time smaller collision delays. We do this by dividing the available bandwidth of the cable into several frequency bands and operating a network (or channel) in each frequency band. Since coaxial cables can easily carry up to 400 MHz, several networks can be accommodated on one cable rather easily. By modulating the baseband data from devices connected to the network to a high-frequency band, total channel bit rates of higher than 50 Mb/s can be obtained easily. However, since the bit rate of each of the nets is kept low, channel inefficiency due to the use of CSMA/CD protocol is not increased. To increase the length of the

cable segment, and at the same time limit the collision delay, we divide the users into communities based on their location and give each user community a network (i.e., one band of frequencies) to transmit most of the time. Thus the "effective" end-to-end delay is reduced although the cable length is increased. The principal characteristics and advantages of our system are the following:

1. Larger channel throughput by using multiple frequency bands.
2. Larger cable length but smaller collision delay by dividing the cable into several parts and operating a network in a given frequency band for each part to be used by a user community, while retaining the listening ability on the full cable length.
3. Complete connection of users between any network.
4. Different grade of service, depending on the complexity of network interface.
5. Restricting most of the high-speed processing to the analog domain and baseband processing to digital domain. Thus, although the network may have channel throughput over 50 Mb/s, individual networks may carry at much lower bit rates.

We describe the system in more detail in the next section, and we develop two protocols for access. A crude analysis is presented at the end to bring out some of the trade-offs for our system.

## II. SYSTEM DESCRIPTION

In this section, we describe one possible implementation of our system. A block diagram is shown in Fig. 1. Each terminal has a frequency-agile Radio Frequency (RF) modem that can modulate binary data for transmission on the cable and demodulate the signal from the cable to extract the transmitted binary data. Unlike long distance transmission, since the intent is not to maximize the data transmission rate, simple inexpensive modulation schemes can be chosen with enough separation between the various frequency bands to keep the filtering simple and to reduce the crosstalk. As an example, if modems based on Frequency Shift Keying (FSK) with 1/4 bit/Hz are used, then a cable of bandwidth 300 MHz can support six CSMA/CD networks of peak rate 10 Mb/s each with a guard band of 10 MHz to separate each of them.

Our block diagram in Fig. 1 shows bidirectional transmission, that is, signals injected on the cable at each tap travel in both directions and the amplifiers are bidirectional. It is necessary for the taps to be bidirectional so that they can receive signals from either direction. Although this is a straightforward extension of the baseband CSMA/CD network, bidirectional amplification and taps may present engineering difficulties, particularly at high frequencies. Alternative de-

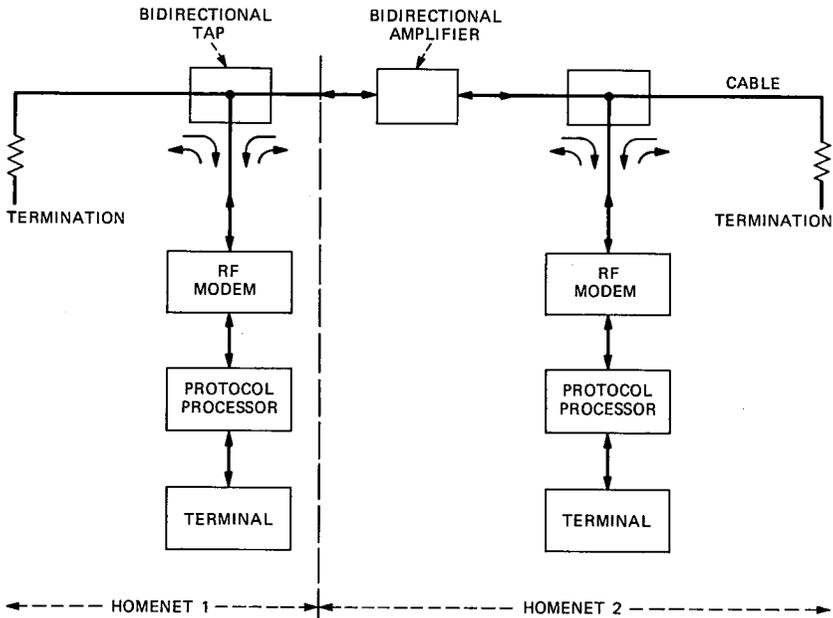


Fig. 1—Single cable system for multiple CSMA/CD networks.

signs using cable-television-type systems and technology are described in Ref. 5.

In Fig. 1, there are two networks (called Homenets) and therefore two frequency bands. The transmitter can, by the agility of the modem, transmit on any of the two frequency bands, and the receiver can receive and demodulate data from both the frequency bands. Terminals attached to homenet 1 transmit mostly on frequency band f1 and those attached to homenet 2 transmit mostly on frequency band f2. If several simultaneous conversations with terminals on different homenets are required (as in the case of a host computer), then a terminal may need multiple transmitters and receivers. Details of the protocols for the access are given in the next section.

### III. THE ACCESS PROTOCOLS

Below we give two types of access protocols; the first does not require synchronization of the different terminals, whereas the second does. Some desirable characteristics of any protocol should be noted first. The access delay should be decreased by scheduling the transmission on a net that is either least busy or has the least chance of collision. The load on the different networks should be distributed such that a situation does not arise in which many terminals are trying to transmit on a network and are unable to do so, while the rest of the

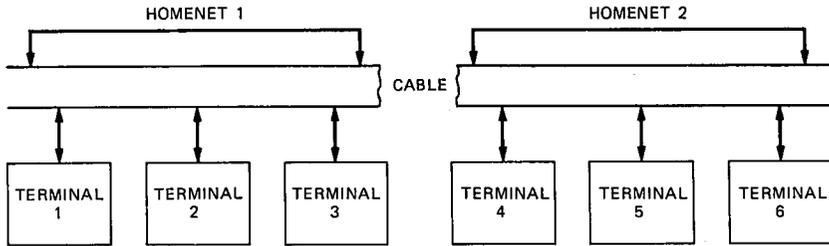


Fig. 2—Homenet assignments based on spatial distance.

networks are carrying a very light load. In carrier sense multiple-access systems, collisions are a result of the terminals knowing the transmission by other terminals only after the propagation delay. Therefore, by requiring all the terminals that are close to each other to initiate their transmission on a particular network, the collision delay can be reduced considerably. Thus, in both the protocols, each terminal is assigned to a particular network. This network is called the Homenet of the terminal. Homenet assignment is included as a part of the address of the terminal. Each terminal maintains a list of the homenet assignments of the other terminals. The homenet assignment, although made primarily by geographical location, may also take into account the desired connectivity, traffic patterns, etc.

Figure 2 shows a system in which there are two homenets and three terminals per homenet. The transmission initiated by any of terminals 1, 2, and 3 is mostly on homenet 1, whereas the transmission by terminals 4, 5, and 6 is on net 2. Since net 1 and net 2 are on two different bands of frequency, the collisions are now localized. That is, data from terminal 1 can only collide with data from terminals 2 and 3. Since the distance between the taps on the cable of terminals 1, 2 and 3 is much shorter compared with distance between the taps of terminals 1 and 6, the probability of collision and, therefore, of data wasted due to collision is significantly reduced. This increases the channel utilization and decreases the delay. Of course, the protocols must and do allow communication between the terminals on different homenets.

### 3.1 Protocol 1

The flow diagram for this protocol is shown in Fig. 3. The details follow.

1. Every terminal has at least two receivers and, therefore, is capable of listening to at least two networks. One of these receivers always listens to the homenet. The other receiver is free to listen to any net.
2. An inactive terminal has one of its receivers listening to its



homenet. The other receiver of that terminal becomes active only after the first receiver starts listening to a network other than its homenet.

3. Any terminal A desiring to transmit to terminal B goes through the following sequence.

a. Determine the homenet of B, say it is  $net_b$ .

b. Listen to  $net_b$  before transmitting.

c. If  $net_b$  is idle (i.e., absence of carrier), transmit  $net_b$  carrier for a period  $T$ , the two-way propagation delay through the total network. This amounts to a priority preempt on  $net_b$ .<sup>6\*</sup>

d. If during the second half of the period  $T$  there is collision on  $net_b$ , then it implies a preemptive transmission from another terminal, not on  $net_b$ . In that case, terminal A backs off and attempts a transmission on  $net_b$  with reduced probability at the next time slot of  $T$ . If there is no collision, then terminal A follows its preempt with a message to terminal B.

e. Terminal B always has one receiver listening to  $net_b$ ; therefore, it receives information from every collision-free packet on  $net_b$ .

f. If packet communications is to be continued, then terminal A starts listening on  $net_b$ , terminal B on  $net_a$ , and both terminals transmit on their homenets. Thus, if a message has several packets, only the first packet may be transmitted on a homenet different from the homenet of the source; all the subsequent packets are transmitted on their own homenet with standard CSMA/CD protocol with retrial period equal to round-trip delay of the homenet.

g. If at step e terminal B is already in communication with some other terminal on a different network, then it still has a receiver on  $net_b$ . If terminal B's transmitter is on  $net_b$  (as it normally is, except when it is trying to set up an initial connection with a terminal on a network other than  $net_b$ ), even if it is in communication with some other terminal, it can send an acknowledgment back to terminal A on  $net_b$ . If, however, Terminal B's transmitter is transmitting on a different channel, there may be delay in sending the acknowledgment.<sup>†</sup>

h. If after successful connection there is no communication for a given amount of time and if the receiver on the homenet receives a message for communication from another source, then the other receiver of both the home terminals go back to their respective homenets.

---

\* Alternatively, there could be a signaling channel in a different frequency band accessed by all the terminals, and the first packet could be transmitted on the signaling channel. This alternative is attractive for a large network, since it confines collision to a common homenet and, hence, does not constrain minimum packet length. With the utilization very low and the network large, the most appropriate protocol on the signaling channel would be ALOHA.

<sup>†</sup> If several simultaneous conversations are required, then a terminal may need multiple transmitters and receivers.

- i. Many different broadcast modes are possible. If the broadcast to only the terminals in the particular homenet is desired, then data is transmitted on that network only. However, if broadcast is required to all the terminals, then the transmitter has to successfully transmit on each network.

### **3.2 Protocol 2**

The above protocol is reasonable in that it reduces collisions and works well when the traffic is quite bursty, with many terminals trying to transmit messages containing large numbers of small packets frequently. However, when there are large file transfers, use of the homenet by a terminal prevents other terminals with the same homenet from using the channel even though the other networks may be idle. Thus, a reasonable protocol is needed that will share the channels more evenly in the presence of large file transfers by one of the terminals. Protocol 2 attempts to accomplish this at the cost of slightly larger average delay in establishing a connection. In this protocol, networks on which a given group of terminals begin transmitting are switched on a periodic basis. The period is of the order of several packets long (or tens of milliseconds). It thus requires a clock at every terminal, which may be provided from a central clock on a different band of frequencies. Details of the protocol are

1. As in protocol 1, terminals in a given geographical area are grouped together. This grouping is made known to all the terminals (similar to homenets).

2. A group has a homenet and an assigned transmission network. The homenet is fixed, whereas the transmission net changes cyclically. A terminal may, at any time, initiate a transmission only on the transmission network to which its group is then assigned. Once initiated, the transmission may spill beyond this fixed interval, since the packet size is not fixed.\* As an example, Fig. 4 shows the case of three groups and three sectors of time.

3. A terminal has at least two receivers. When the terminal is idle, both of them listen to the homenet. After establishing a connection, however, one of the receivers switches to the network on which it has established connection with the other terminal (and, therefore, the network to which it listens changes cyclically) and the other receiver remains at the homenet.

---

\* In CSMA/CD the entire packet must be received and CRC checked before the destination address is verified. Thus, if the entire packet is not received before the period ends, the receiver may miss it. To overcome this problem, a separate CRC is provided for the header information and a source terminal starts a transmission sufficiently before the end of a period such that the destination is able to receive the header information before the period ends.

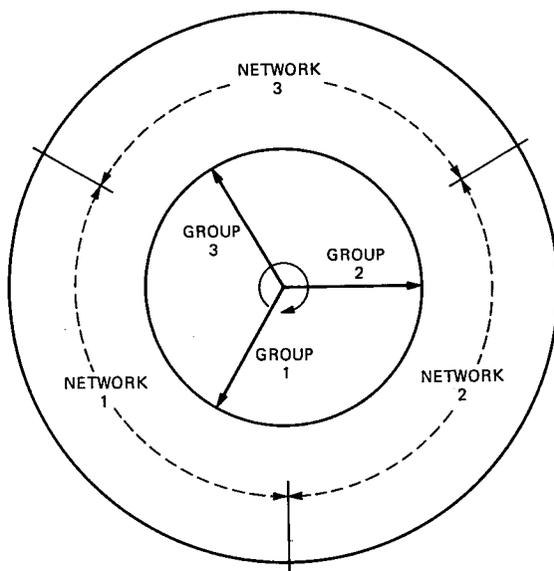


Fig. 4—Protocol of Section 3.2. The assignment of group  $j$  terminals to  $net_k$  is made by rotating the inner circle at a given speed. The terminals in group  $j$  may initiate transmission on  $net_k$  if the group  $j$  pointer is in the sector corresponding to  $net_k$ .

4. Terminal A, desiring to transmit to terminal B, goes through the following sequence:

- a. Determine the homenet of terminal B; let it be  $net_b$ .
- b. Wait for the time period when terminal A is allowed to transmit on  $net_b$ .
- c. At that time, listen to  $net_b$  before transmitting.
- d. If  $net_b$  is idle (during the assigned time slot), then transmit to terminal B on  $net_b$ .
- e. If there is no collision, a packet is assumed to have been received by its intended receiver. If there is a collision, terminal A ceases transmission immediately and tries again, using a standard retry strategy, but the additional constraint that its starting time must be when terminal A is allowed to transmit on  $net_b$ .
- f. If communication is to be continued, then terminal B switches one of its receivers to the net on which terminal A will be transmitting (this will periodically switch), and terminal A will set its receiver on the net on which terminal B will be transmitting.
- g. If at step f terminal B is already in communication, then its acknowledgment to terminal A will so indicate.
- h. If after a successful connection there is no communication for a given amount of time, receivers of both the terminals go back to their respective homenets.

## IV. ANALYSIS

In this section, we present a simplified analysis of the access protocol described in Section 3.1. The assumptions follow those made by Metcalfe and Boggs<sup>7</sup> for a single CSMA/CD network. The symbols used are

$P$  = Number of bits in a packet.

$C$  = Peak channel capacity.

$T$  = Round-trip propagation delay.

$Q$  = Number of stations continuously queued to transmit a packet; represents the total offered load.

$N$  = Number of networks on the frequency-multiplexed cable.

$\Pi_m$  = Probability that a message from a source contains  $m$  packets.

The analysis is used first to calculate the delay versus the channel throughput characteristics of a single CSMA/CD network operating at a peak channel capacity of  $C$ . This uses the formulas given by Metcalfe and Boggs. The extension is then made to the case of multiple networks whose total capacity is  $(NC)/t$ , but is divided equally into  $N$  networks. Number  $t$  is assumed to be larger than 1. If  $t = N$ , then the total capacity  $C$  is divided equally in  $N$  networks, each with capacity  $C/N$ . If  $t = N/2$ , then each net has capacity of  $2C/N$ , adding to a total of  $2C$ .

### 4.1 Single network

The average delay in sending a packet (including transmission time) when  $Q$  stations are continuously queued to transmit a packet, is given by

$$D_1 = \frac{P}{C} + T[(1 - 1/Q)^{1-Q} - 1]. \quad (1)$$

This assumes an optimum retry strategy. Since the Offered Load (OL) is  $Q$  packets, in terms of bits it is given by

$$OL = PQ. \quad (2)$$

### 4.2 Multiple networks

Here if a message from a source contains  $m$  packets, then the first packet may be transmitted on a different network, but the subsequent  $(m - 1)$  packets will be transmitted on the homenet of the source. Thus, the total transmission time is divided into two parts: time to transmit the first packet and time to transmit the remaining packets. It is assumed for simplicity that the length of each homenet is the same and it is  $1/N$  times the total cable length. The following analysis can be easily modified for other configurations.

#### 4.2.1 Time to transmit first packet

If the first packet of a source on homenet  $j$  is transmitted on homenet  $k$ , then

$$\text{Time for a packet} = \frac{tP}{C} + \text{contention time.} \quad (3)$$

Assume that a source on the  $j$ th homenet has probability  $\rho_{jk}$  of wanting to communicate with a station on the  $k$ th homenet. Further, assume that the number of stations is the same for all nets; distribution,  $\Pi_m$ , of packets is uniform for all messages; and  $\{\rho_{jk}\}$  are a constant\* independent of  $j, k$ . Then the total traffic offered to the  $k$ th homenet is

$$q_k = \frac{Q}{N}, \quad (4)$$

of which the offered load from out-of-net is

$$q_k^1 = \frac{N-1}{N} Q \Pi_1 \quad (5)$$

and from within homenet

$$q_k^r = \frac{Q}{N} \left( 1 - \frac{N-1}{N} \Pi_1 \right). \quad (6)$$

Therefore the time per out-of-net packet is

$$D_{21} = \frac{tP}{C} + T[(1 - 1/q_k)^{1-q_k} - 1]. \quad (7)$$

#### 4.2.2 Time for subsequent packets

This transmission is on the homenet itself. Therefore, the time for each packet is simply obtained by

$$= \frac{tP}{C} + \frac{T}{N} [(1 - q_k^r)^{1-q_k^r} - 1]. \quad (8)$$

This neglects the traffic generated by first packets of terminals from other homenets. It is assumed that the first packet is a small fraction of the total message and does not result in any significant traffic.

#### 4.2.3 Total time

Since the probability that a message contains  $k$  packets is  $\Pi_k$ , the average time per packet is given by

---

\* We have made no measurements of traffic on real systems to justify this assumption. It is made only so that a closed-form expression can be derived for the delay. If other values of  $\rho_{jk}$  are more realistic, they can be substituted easily in the equations that follow.

$$\begin{aligned}
 D_2 = & \frac{N}{Q} (q_k^1 D_{21} + q_k^r D_{2r}) = \frac{tP}{C} \\
 & + (N - 1) \Pi_1 T [(1 - 1/q_k^1)^{1-q_k^1} - 1] \\
 & + \left(1 - \frac{N - 1}{N} \Pi_1\right) \frac{T}{N} [(1 - 1/q_k^r)^{1-q_k^r} - 1]. \quad (9)
 \end{aligned}$$

Thus the delay versus the offered load characteristic will be given by  $D_2$  versus OL.

#### 4.2.4 Total time with Protocol 2

With Protocol 2, all packets are sent on homenet, and the delay for optimum strategy is given as

$$D_3 = \frac{tP}{C} + \frac{T}{N} \left[ \left(1 - \frac{N}{Q}\right)^{1-Q/N} - 1 \right].$$

When  $\Pi_1$  is small, that is, messages consist on average of many packets, then  $D_3$  and  $D_2$  differ very little from each other.

#### 4.3 Optimum $N$ for multiple networks

It is possible to compute the optimum number of nets based on the above expressions for average delay per packet. This can be done for the case when  $Q$  (and  $Q/N$ ) is large and the messages contain a large number of packets, implying that the average delay per packet is dominated not by the first packet, but by the subsequent packets. From eq. (1), for single network, the delay is given by

$$D_1 = \frac{P}{C} + T[(1 - 1/Q)^{1-Q} - 1]$$

since

$$\begin{aligned}
 \lim_{k \rightarrow \infty} (1 + x/k)^k &= e^x \lim_{Q \rightarrow \infty} [D_1] = \frac{P}{C} + T \left[ \frac{1}{e^{-1}} - 1 \right] \\
 &= \frac{P}{C} + T(e - 1). \quad (10)
 \end{aligned}$$

For multiple networks, the delay is approximated by

$$D_2 = \frac{tP}{C} + \frac{T}{N} [(1 - N/Q)^{1-Q/N} - 1]. \quad (11)$$

If  $t = N$ , then the total capacity  $C$  is divided among  $N$  nets equally. However, in most cases, the individual  $N$  nets may have a capacity such that the capacities add up to more than  $C$ . Thus let  $t = N/s$ . In

this case, each net has a capacity  $(sC)/N$  and the total capacity is given by  $sC$ . The delay then becomes

$$D_2 = \frac{NP}{sC} + \frac{T}{N} [(1 - N/Q)^{1-Q/N} - 1]. \quad (12)$$

For large  $(Q/N)$  (or small  $N/Q$ ) we can expand  $D_2$  as a function of  $N/Q$  in Taylor's series

$$D_2 \approx \frac{NP}{sC} + \frac{T}{N} \left[ e - \frac{e}{2} \frac{N}{Q} - 1 \right] = \frac{NP}{sC} + \frac{T}{N} (e - 1) - \frac{Te}{2Q}. \quad (13)$$

Minimum  $D_2$  with respect to  $N$  is achieved when

$$N = \sqrt{\frac{sCT(e - 1)}{P}}. \quad (14)$$

For some typical cases, the optimum  $N$  can be worked out as follows:

$$s = 1, C = 10 \text{ Mb/s}, T = \text{propagation delay in seconds on a cable of length 2.5 km} \cong \frac{(2.5 \times 10^{-5})}{2} \text{ sec.}$$

$$P = 1000 \text{ bits;}$$

then

$$N \leq 1 \rightarrow N = 1.$$

Thus with standard Ethernet\* parameters, from the point of view of average delay per packet,  $N$  should be 1.

$$s = 1, C = 10 \text{ Mb/s}, T = \text{propagation delay in seconds on a cable of length 20 km} = (20 \times 10^{-5})/2 \text{ sec.}$$

$$P = 500 \text{ bits;}$$

then

$$N \cong 2.$$

This implies that as the length of the cable increases, more networks are required.

---

\* Trademark of Xerox Corporation.

$$s = 5, C = 10 \text{ Mb/s}, T = \text{propagation delay in seconds on a cable of length 20 km} = \frac{(20 \times 10^{-5})}{2} \text{ sec.}$$

$$P = 500 \text{ bits;}$$

then

$$N \cong 4.$$

This implies that if each net is operated at 12.5 Mb/s, adding up to a total capacity of  $12.5 \times N$  Mb/s, average delay is minimized when  $N = 4$ .

#### 4.4 Delay versus offered load plots

The average time per packet derived in Section 4.1 was evaluated for a variety of cases and is plotted in Figs. 5 and 6. In all the cases, a

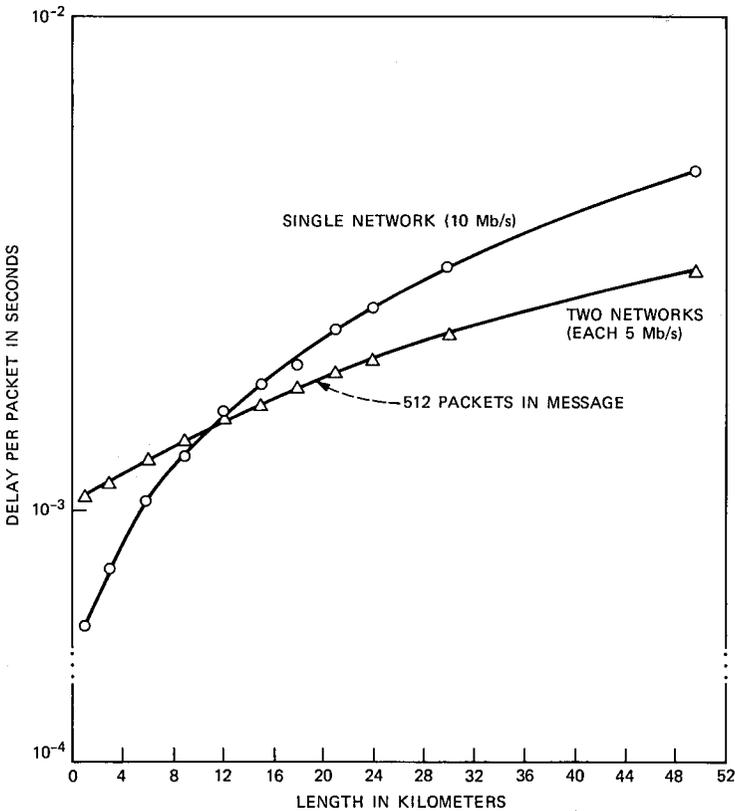


Fig. 5—Average delay per packet versus length of the cable. Performance comparisons are made between two networks on a cable and a single net.

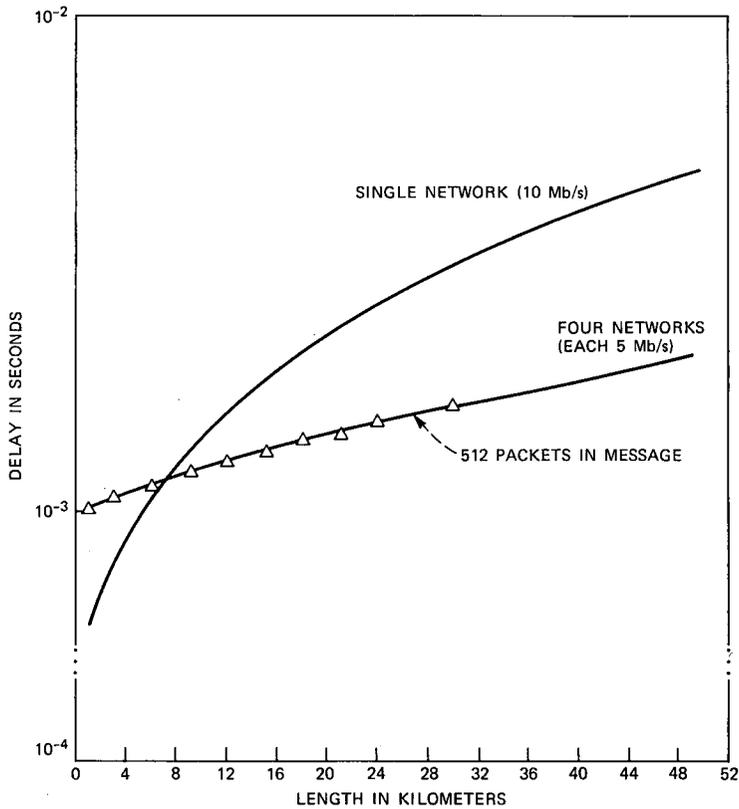


Fig. 6—Average delay per packet versus length of cable. Comparison is made between four networks on a cable and single network.

packet size of 500 bits was used. The curves in Figs. 5 and 6 are for  $Q = 200$ , that is, 200 packets are continuously queued. It is assumed that a message has 512 packets, that is,

$$\Pi_1 = \Pi_2 = \dots, = \Pi_{512} = 1/512; \quad \Pi_k = 0, \quad k = 513, \dots$$

As expected, the delay increases with length, and depending upon the other parameters of the network, the delay corresponding to multiple network becomes smaller than that corresponding to single network if the length is increased beyond a certain value. Figure 7 shows the variation of the average delay with respect to number of networks. The capacity of each of the nets is equal and is such that the total capacity of all the nets adds up to capacity of the single network. We find that, as expected, for the parameters chosen in Fig. 7, the average delay does show a minimum around  $N = 4$ . This verifies our approximations of the previous section.

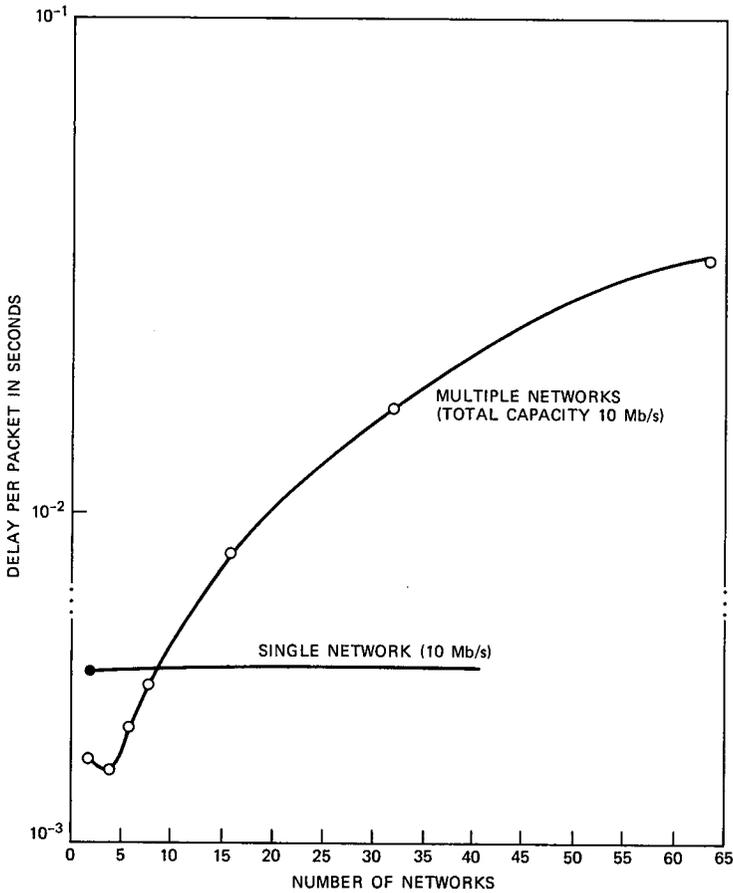


Fig. 7—Average delay per packet versus number of networks. Each network has a capacity of  $C/N$ , where  $N$  is the number of nets and  $C$  is the capacity of the single network.

## V. CONCLUSIONS

We have described a broadband local area computer network. It consists of several local area networks whose data is frequency multiplexed on a single cable. The entire cable length is divided into parts; each part is assigned a network and a frequency band for transmission. We have also described two protocols that overcome some of the limitations of the present baseband as well as broadband CSMA/CD networks. Using our schemes, it is possible to increase the channel throughput and the length of the cable network, reduce the delay due to collisions, and at the same time allow complete connectivity among all the terminals and devices logged into any network. Approximate analytical results are also presented to substantiate these claims.

## VI. ACKNOWLEDGMENTS

We are grateful to Jerry Foschini and Nick Maxemchuck for many helpful discussions.

## REFERENCES

1. A. S. Tanenbaum, *Computer Networks*, New York: Prentice-Hall, 1981, Chap. 7.
2. W. R. Franta and I. Chlamtac, *Local Networks*, Lexington, MA: Lexington Books, 1981.
3. The IEEE Project 802, Local Area Network Standards, CSMA/CD Access Method and Physical Layer Specification, IEEE P802-3-82/0.1 10, December 1982.
4. M. Stahlman, "Inside Wang's Local Net Architecture," *Data Commun.*, (January 1982), pp. 85-90.
5. N. F. Maxemchuck and A. N. Netravali, "A Multifrequency Multiaccess System for Local Access," *Proc. ICC 83*, Boston.
6. N. F. Maxemchuck, "A Variation on CSMA/CD That Yields Movable TDM Slots in Integrated Voice/Data Local Networks," *B.S.T.J.*, 61, No. 7 (September 1982), pp. 1527-50.
7. R. M. Metcalfe and D. R. Boggs, "Ethernet: Distributed Packet Switching for Local Computer Networks," *Commun. ACM*, 19, No. 7 (July 1976), pp. 395-404.

## AUTHORS

**Zigmantas L. Budrikis**, B.Sc., 1955, and B.E. (Hons I, Electrical Engineering), 1957, University of Sydney; Ph.D., 1970, University of Western Australia; P.M.G. (now Telecom Australia) Research Laboratories, 1958-1960; Aeronautical Research Laboratories, Fishermen's Bend, 1961; Electrical Engineering Faculty at University of Western Australia, 1962—. Mr. Budrikis has had a number of visiting appointments: University of California at Berkeley, 1968; AT&T Bell Laboratories, 1972, 1973, 1981, 1983, 1984; TU Munich, 1977. He is interested in problems in communications, man-machine interfaces, and foundations of electromagnetism. Fellow, IE Australia; member, IEEE, Optical Society of America, New York Academy of Science.

**Arun N. Netravali**, B. Tech. (Honors), 1967, Indian Institute of Technology, Bombay, India; M.S., 1969, Ph.D. (Electrical Engineering), 1970, Rice University; Optimal Data Corporation, 1970-1972; AT&T Bell Laboratories, 1972—. Mr. Netravali has worked on problems related to filtering, guidance, and control for the space shuttle. At AT&T Bell Laboratories, he has worked on various aspects of digital processing and computing. He was a Visiting Professor in the Department of Electrical Engineering at Rutgers University and the City College. He is presently Director of the Computer Technology Research Laboratory. Mr. Netravali holds over 20 patents and has had more than 60 papers published. He was the recipient of the Donald Fink Prize Award for the best review paper published in the Proceedings of the IEEE and the journal award for the best paper from the SMPTE. Editorial board, Proceedings of the IEEE; Editor, IEEE Transactions on Communications; senior member, IEEE; member, Tau Beta Pi, Sigma Xi.