

Voice and Data on a CATV Network

NICHOLAS F. MAXEMCHUK, MEMBER, IEEE, AND ARUN N. NETRAVALI, SENIOR MEMBER, IEEE

Abstract—A technique for transmitting voice and data on a conventional subsplit CATV network is described. The technique uses a variation on the carrier sense multiple access/collision detection protocol used in local area networks. A variation on this protocol, called movable slot TDM allows periodic voice sources to share this network without distorting the voice signal. A scheme for sharing frequency spectrum and space is outlined that removes the distance and transmission rate constraints associated with such networks. A transmission strategy is described which allows these protocols to be used on a subsplit CATV system. The system is thus a candidate for local/metropolitan area networks which support digital voice and data services along with analog (mostly one-way) video.

I. INTRODUCTION

A system is described for transmitting and switching voice and data signals on a CATV system. The voice signals are digitized and transmitted in packets. Voice and data packets contend for the broadcast CATV channel by using a variation on the carrier sense multiple access/collision detection (CSMA/CD) protocol designed for local area networks [1]. The packets can be received anywhere in the network, and connections are established by recognizing packet addresses. Therefore, switching in this network is distributed.

In this system, voice packets contend for the network using a variation on CSMA/CD called movable slot TDM (MSTDM) [2], [3]. This protocol guarantees that voice samples are not lost because of the variation in delay encountered in CSMA/CD networks. A technique called Homenets [4] is used to remove the distance and transmission rate constraints which have restricted the CSMA/CD protocol to small local area networks. Finally, a transmission strategy is described which allows these protocols, which are designed for networks in which signals propagate in all directions, to be applied to subsplit CATV networks, which are highly directional.

II. DESCRIPTION OF SYSTEM

A. MSTDM

In the CSMA/CD protocol a source listens to the channel before transmitting a packet of data. If it detects another source transmitting, it tries again later. While a source is transmitting it also listens to the channel. If it detects a collision with another source it stops transmitting

and tries again later. When the network is heavily utilized, a source may have to try many times before acquiring the channel. If the source is transferring packets of periodic voice samples to a receiver which is reconstructing the voice waveform, and must wait a long time before acquiring the channel, the receiver will run out of samples. This results in distortion of the speech waveform.

MSTDM is a variation of the CSMA/CD protocol. It guarantees that voice samples are not lost because of delays incurred accessing the network. Furthermore, if the voice utilization of the network goes to one, the network degenerates to a TDM network and continues to function.

Data sources and the first packet from idle voice sources access the network using the conventional CSMA/CD protocol with any retry strategy. Successive packets from a voice source listen before transmitting, but do not detect collisions. If the channel is busy they wait for the channel to become idle, then transmit immediately. The packet format and the periodic nature of voice sources guarantee that active voice sources win collisions with data packets or new voice sources and that active voice sources never collide.

The voice packet format is shown in Fig. 1. It contains


- 1) a preempt interval long enough for a source which is performing collision detection to detect a collision and stop transmitting before the voice source starts transmitting useful information;
- 2) a data area long enough to transmit all of the voice samples acquired in T_p s, the nominal time between voice packets; and
- 3) an overflow area long enough to transmit all voice samples acquired during a voice packet transmission time T_v .

The only constraint placed on data sources is that the data packets are shorter than voice packets.

The protocol for voice sources is a subset of the CSMA/CD protocol, and sources do not have to collect any additional information about the state of the system. It is simple because there is only one periodic rate supported by the system. Higher rate channels must be obtained as multiple voice channels. A formal proof of the operation of MSTDM is given in [2]. In this proof the propagation delay between terminals is taken into account, and it is shown that timing discrepancies between voice sources can be tolerated. These intricacies are disregarded in the following explanation of the protocol in order to provide some insight into its operation.

Voice sources always plan to transmit their next packet T_p s after they successfully acquire the channel.

Manuscript received February 15, 1984; revised September 5, 1984.
The authors are with AT&T Bell Laboratories, Murray Hill, NJ 07974.



VOICE PACKET
Fig. 1. Voice packet format.




Fig. 2. Channel utilization in an MSTDM system. (a) Three voice channels delayed by a data packet. (b) Three voice channels delayed by a new voice channel, then a data packet. (c) A nearly full system when a new voice source gains access.

1) When a new voice source successfully acquires the channel any samples which cannot be transmitted in the packet are discarded, and in T_p s this source has enough samples for the next packet.

2) If a packet from an active voice source is delayed, it transmits the samples which arrive during the delay in the overflow area, and in T_p s this source has enough samples for the next packet.

The constant scheduling interval guarantees that packets from an active voice source are scheduled at least a packet transmission time, T_v s, apart. Active voice sources are only delayed by sources which are already transmitting. Since all packets take T_v s or less to transmit, an active voice source is delayed less than T_v s, and two active sources can never collide. This is shown in Fig. 2(a) and (b) where voice sources 1, 2, and 3 are delayed by a data packet D and a new voice source A . The first line in the figure shows where the sources originally plan to transmit, and the second line shows where they plan to transmit after they are delayed. The third line of Fig. 2(b) shows that the new voice can be delayed without causing a collision. All of the samples which arrive during a voice packet's delay can be transmitted during the overflow interval, and no voice samples are lost.

In most CSMA/CD networks, when the channel utilization goes to one the delay goes to infinity. In an MSTDM system, when the voice source utilization goes to one the system degenerates to a TDM system. Assume that active voice sources have almost completely used the channel capacity. As long as there is spare capacity in the system, a




Fig. 3. Voice transmit and receive buffer operation in an MSTDM system.

new voice source may acquire the channel. If the spare capacity is insufficient to transmit a complete packet, the next voice source will be delayed. The amount the next active voice source is delayed is always less than or equal to the delay of the preceding active voice source. Therefore, the active sources never collide, and the extra samples can always be transmitted in the overflow area. In effect, there is always an active voice source waiting to acquire the channel, so that new voice sources or data sources cannot transmit, and the active voice sources are always delayed, so that the overflow area in a voice packet is always being used. This is shown in Fig. 2(c) where a new voice source 4 acquires a channel in a nearly full system.

The operation of the buffers for voice sources is shown in Fig. 3. The first voice packet is placed in the output buffer so that the last sample will be used $T_p + T_v$ s after it arrives. When an active voice source acquires the channel, it transmits all of the samples acquired since the last packet was transmitted. The number of samples used at the receiver between the arrival of two packets is the same as the number of samples acquired at the transmitter between the transmission of the two packets. Therefore, after a packet arrives, there are enough samples at the receiver to last $T_p + T_v$ s. The next packet is scheduled to arrive in T_p s, and the maximum it can be delayed is less than T_v s. Therefore, the receiver never runs out of samples. A more detailed description of an MSTDM receiver is given in [3].

In an MSTDM system, voice is given a higher priority than data. However, this does not necessarily increase the delay experienced by data packets. When a data packet and a voice packet collide, the voice packet completes its transmission, therefore the probability of collision when the data packet retries is reduced. There are effects which counteract this advantage. The channel utilization increases because the voice packets are longer by the preempt and overflow areas, and the data source finds the channel busy




Fig. 4. Regional and frequency division of homenets.

if it retries within a packet transmission time. The exact effect of this protocol on data delay can not be determined analytically; however, extensive simulations have been conducted and are reported in [2].

B. Homenets

In a CSMA/CD network, the ratio of the propagation time between terminals to the transmission time of the packet determines the fraction of the transmission which is susceptible to collisions. The ratio also determines the relative duration of packets which have collided and packets which have been successfully transmitted. CSMA/CD networks are most efficient when the propagation delay between sources is a small fraction of the time it takes to transmit the packet. In an MSTDM system, the preempt area equals the round-trip propagation delay plus the collision detection time. Therefore, the ratio of the propagation delay to the transmission rate determines the fraction of the packet which carries useful data.

CSMA/CD and MSTDM protocols work well on short networks with a high transmission rate or long networks with a low transmission rate. They are designed for local area networks where the distance between users is typically less than 1 or 2 km and the total bit rate of the system is on the order of 10 Mbits/s. When these protocols are applied to CATV networks, with an order of magnitude larger distances, the bit rate should be decreased by an order of magnitude in order to maintain the same efficiency. However, this is incongruous with providing data and voice communications to a large community of users. Homenets allow these protocols to be applied to networks which span a larger area and have a higher total bit rate than local area networks.

Homenets are created by dividing the users in a network into regional groups, as shown in Fig. 4. The communications channel is also divided, and a fraction of the capacity is assigned to each regional group. Users in a region only transmit in the fraction of the capacity assigned to their




Fig. 5. Homenets on a CATV system.

group. This decreases the ratio of the propagation delay between contending terminals to the packet transmission in two ways.

- 1) The maximum distance between contending users is reduced from the maximum distance between users in the entire network to the maximum distance between users in a region.

- 2) The time it takes to transmit a packet increases because only a fraction of the network capacity is available in a region.

The connectivity between users on the network is complete. When a source transmits a packet it can address it to any user. Once a source successfully contends for the channel it does not make any difference how long it takes for the packet to be received. Therefore, a terminal in any region can receive packets addressed to it by receiving the packets on the source's fraction of the channel.

In the CATV implementation of homenets, the channel is partitioned into frequency bands corresponding to TV channels. This allows homenets to coexist on a network being used for conventional CATV functions and makes it possible to implement homenets with readily available CATV components.

Partitioning the channel in the frequency domain allows a system with a high total bit rate to be constructed without high bit rate digital components. For instance, if there are 80 channels on the network, and each channel has a 10 Mbit/s CSMA/CD network, the total bit rate of the system is 800 Mbits/s. However, a receiver first selects a channel, then searches for its address. Therefore, a receiver must only look at 10 Mbits/s. Telling a receiver which channel to receive corresponds to setting up a call. Three techniques for doing this are described in Section IV.

C. Transmission Strategy

In order to apply the homenet strategy to a conventional tree-configured CATV network the network should be partitioned as shown in Fig. 5. However, CATV taps and amplifiers are directional. The taps are designed to receive

signals from, and transmit signals toward, the head end of the system. The amplifiers are designed so that signals in the TV bands propagate from the head end toward the leaves of the tree. In subsplit and midsplit systems there are also frequency bands, which are amplified and propagate upstream, towards the head end.

Completely connected data networks have been implemented on CATV networks [5]. In these networks all of the sources transmit toward the head end of the system on one of the upstream channels. At the head end the upstream channel is received and translated into one of the downstream channels. All of the receivers receive data on the downstream channel. A CSMA/CD system is implemented by listening to the downstream channel to determine if the channel is busy, transmitting on the upstream channel, and then listening to the downstream channel to determine if any other signal has arrived at the reflection point at the same time as the signal transmitted.

It is easier to implement collision detection on this type of a system than on a system in which signals propagate in all directions from a source. If the signal propagates in all directions, the signal from a distant source has been attenuated. When a collision occurs, that signal must be detected in the presence of a much stronger signal being transmitted by the local source. In order to detect this weak signal an analog comparison must be performed on the signal. However, in a system with a single reflection point, the relative amplitude of two colliding signals is the same at all of the receivers. If a source decodes the received signal, performs a comparison in the digital domain, and determines that the same data are received as were transmitted, then every other receiver can receive these data and the source can assume that a collision has not occurred.

When the CATV tree is partitioned into homenet, each homenet has only one upstream connection to the rest of the network, as shown in Fig. 5. This point is the head end of the homenet. Signals transmitted by sources within a homenet are reflected back at the head end of the homenet, instead of the head end of the entire network. Sources only contend with other sources in the same homenet for the homenet receive frequency. Therefore, the propagation distance between two sources which may collide is the sum of the distances between the sources and the head end of the homenet, rather than the sources and the head end of the network. As the network is divided into more and more homenets the propagation distance between terminals which may collide decreases and the efficiency of the CSMA/CD networks increases.

The upstream channel is only used to communicate between sources in the homenet and the head end of the homenet. Therefore, notch filters can be inserted at the boundaries of the homenets to prevent the upstream signal from propagating between homenets. This allows the same upstream frequency band to be used in every homenet and reduces the number of upstream channels required by the network. Since every homenet uses the same upstream channel the transmitter in the sources is independent of the source's homenet. The notch filters also reduce the noise




Fig. 6. Frequency assignments for homenet 2.

sources affecting the transmitted signal to those in the homenet, rather than those in the network.

Each homenet has its own receive channel, but there is only one transmit channel. It is possible to take advantage of this asymmetry and to trade transmit channel bandwidth for simpler source transmitters. If the bandwidth of the transmit channel is increased, every source transmits fewer bits per hertz. At the reflection point the signal must be demodulated and remodulated into a narrower bandwidth receive channel. However, there are many more sources than reflection points, which makes this tradeoff worthwhile.

Once the transmit channel is translated to the homenet receive channel, any terminal downstream from the reflection point can receive data from a source in the homenet by listening to the homenet receive channel. However, not all terminals are downstream from the reflection point. In order to let all terminals receive data from terminals in a homenet the reflection point for the homenet is made to appear as the root of the tree for the entire network for the homenet receive band.

An example of how this is done is shown in Fig. 6 for a terminal in Homenet 2. All sources transmit in the frequency band F0. Notch filters at the head ends of Homenets 2 and 3 prevent signals transmitted in band F0 on Homenets 1, 2, and 3 from colliding. At the head end of Homenet 2 the signal received in F0 is retransmitted in frequency bands F2 and FD2. Two sources in Homenet 2 detect collisions by receiving the signal in frequency band F2 and comparing it to what they transmitted. The propagation distance between these sources is the sum of their propagation distances to the head end of Homenet 2. A receiver in Homenet 3 receives signals from a source in Homenet 2 by receiving the signal in F2 and looking for its




Fig. 7. Frequency utilization, translators, and filters in a 4 homenet system.

address. Frequency band FD2 is a point-to-point link between the reflection point for Homenet 2 and the head end of the entire system. This link does not have to be accessed at any other point in the system. Therefore, it can be put on the CATV cable if capacity is available, or it can be put on any other available channel. At the head end of the network the signal from FD2 is retransmitted on Channel F2. Receivers in Homenets 1 and 4 can now also receive signals from a source in Homenet 2 by receiving the signal in band F2. In order to prevent signals in F2 generated at the head end of the network from colliding with signals in F2 generated in Homenet 2, F2 must be notch filtered at the head end of Homenet 2. Fig. 7 shows the frequency translators and filters which are required in a system with 4 homenets.

III. ANALYSIS

A. Scaling with the Number of Homenets

Increasing the number of homenets decreases the propagation time between users and the transmission rate within a homenet. Both of these effects increase the efficiency of

the CSMA/CD and MSTDM protocols. However, exactly how the propagation delay decreases with the number of homenets depends upon the distribution of sources and the geometry of the system.

There are special distributions of sources for which the improvements provided by homenets are much greater than for uniformly distributed sources. Some of these distributions correspond to practical networks and are therefore of particular interest. Two such distributions are:

- 1) a small number of clusters of sources with large areas between them; and
- 2) a small number of large point sources in the midst of a large number of small sources.

The first type of distribution occurs in metropolitan areas in which a company has a small number of large buildings. By defining homenets which encompass each of the buildings, the propagation distance between contending sources is reduced from the entire metropolitan area to a building. In effect, the network is reduced to a small number of local area networks with the homenet transmission strategy replacing the gateways between the local networks. The second type of distribution occurs in both data and voice networks. In time-shared computer networks there is much more traffic from the central computer than from any of

the terminals accessing the computer. The central computer is a large point source. In voice networks the access point between the metropolitan area network and the long-haul network is the point source. There is much more traffic from the long-haul network than from any particular phone. By defining one of the homenet to encompass just the large point source, this portion of the network can be made extremely efficient. It should be noted that the strategy of isolating a single large source from a large number of small sources has been successfully employed in the ALOHA network [6].

In the remainder of this section the effect of dividing uniformly distributed sources into N homenets is investigated. The area of the resultant homenets is inversely proportional to N . It will be assumed that the resultant homenets are roughly circles or squares. Therefore, the maximum distance between terminals in a homenet is inversely proportional to $N^{1/2}$. If the maximum signal propagation time between sources in the original system is p s, then the maximum propagation time between two sources in the same homenet is $p_N = p/N^{1/2}$. In an actual system the resultant propagation time is affected by the geometry of the system. For instance, if the original system is a long rectangle, approaching a straight line, a greater reduction in the propagation time is obtained. The maximum distance between terminals is inversely proportional to N . Whereas if the original system is a circle that is divided into two semicircles, no reduction in the maximum propagation delay is obtained.

It will be assumed that as a system is divided into homenets the following hold.

- 1) The total bandwidth is constant.
- 2) The modems operate at constant number of bits per hertz.
- 3) The transmit and receive filters are frequency scaled versions of the same RLC filters.

Because of the last assumption the ratio of the stopband between homenet frequency bands to the passband in a homenet frequency band is constant. Therefore, the fraction of the total system bandwidth which can be used to transmit data is constant as the number of homenets is increased, and the bandwidth per homenet is inversely proportional to N . Because of the second assumption the transmission rate in a homenet is $R_N = R/N$ where R is the transmission rate of a system which has not been subdivided.

In a broad-band CSMA/CD network, signals propagate from a source to a reflection point, then back to the receiver. Signals from different sources collide if they arrive at the reflection point at the same time. If K bit packets are transmitted on the network, and $K/R < p/2$, sources at the extremities of the network cannot effectively use a CSMA/CD protocol. By the time a source at the extremity of the network detects a transmission any bits which it transmits will not arrive at the reflection point before the transmission is complete. The maximum propagation delay to the reflection point is $p = d/c_m$ where d is the maximum distance to the reflection point and c_m is the speed of

light in the medium. Therefore, $d * R = 2K/c_m$ is the maximum size rate product for which carrier sense systems can be used. When the system is divided into N homenets, collisions only occur between sources within the homenet, and the size rate product for the entire network increases to $(d * R)_N = N^{3/2}(d * R)$. Therefore, by dividing a system into homenets, CSMA/CD protocols can be applied to larger or faster networks.

B. Data Delay

As the system is divided into more homenets than required to support a CSMA/CD protocol, the efficiency of the protocol increases. In the limit, the system is divided so that there is one source per homenet. The system becomes a frequency division multiplexed system, and there are no collisions or access delays. However, as the network is divided, it takes longer to transmit a packet. The average delay is the sum of the average access delay and the transmission time. There is an optimum number of homenets a system can be divided into to minimize the average delay.

The average delay in a CSMA/CD network has been determined in a number of papers. Here, the results obtained by Lam [7] and modified by Bux [8] are used. Lam's analysis assumes a slotted system, an infinite number of sources, and an adaptive retry strategy which guarantees an approximately constant probability of collision. Bux modifies this result to apply to unslotted systems. These analyses assume an arbitrary distribution of packet lengths. By constraining these results to constant length packets the following normalized average delay within a homenet is obtained:

$$D_N = \frac{\lambda_N(1 + (4e + 2)a_N + 5a_N^2 + 4e(2e - 1)a_N^2)}{2(1 - \lambda_N(1 + (2e + 1)a_N))} + 1 + (2e + 1/2)a_N - \frac{(1 - e^{-2a_N\lambda_N})\left(\frac{2}{\lambda_N} + \left(\frac{2}{e} - 6\right)a_N\right)}{e^{-\lambda_N - a_N\lambda_N - 1} - 1 + 2e^{-2a_N\lambda_N}} \quad (1)$$

where

- D_N average delay in packet transmission times;
- a_N ratio of propagation delay to the packet transmission time;
- λ_N average number of packet arrivals per packet transmission time.

To investigate the effect of dividing a system into homenets the above expression for the delay in a homenet can be expressed in terms of a and λ in the undivided system and normalized with respect to the packet transmission time in the undivided system. Since the propagation delay is reduced by $1/N^{1/2}$ and the packet transmission time is increased by N , $a_N = a/N^{3/2}$. Since the message arrivals within a homenet are reduced by $1/N$ and the packet transmission time is increased by N , $\lambda_N = \lambda$. The delay in a system with N homenets, normalized with respect to the




Fig. 8. Normalized average delay as a function of the number of homenets for system utilizations of 0.2, 0.4, 0.6, and 0.8.

packet transmission in the undivided system is

$$D_N = N * \left[\frac{\lambda \left(1 + \frac{(4e+2)a}{N^{3/2}} + \frac{(5-4e+8e^2)a^2}{N^3} \right)}{2 \left[1 - \lambda \left(1 + \frac{(2e+1)a}{N^{3/2}} \right) \right]} + 1 + \frac{(2e+1/2)a}{N^{3/2}} \right] - \frac{\left[1 - e^{-2a\lambda/N^{3/2}} \right] \left[\frac{2}{\lambda} + \left(\frac{2}{e} - 6 \right) \frac{a}{N^{3/2}} \right]}{e^{-\lambda - a\lambda/N^{3/2}} - 1 + 2e^{-2a\lambda/N^{3/2}}} \quad (2)$$

In Fig. 8 the normalized delay is plotted versus the number of partitions for $a = 0.5$ and utilizations λ of 0.2, 0.4, 0.6, and 0.8. The delay in the equation has a pole in the first term. When $1 - \lambda(1 + (2e+1)a_N) \leq 0$, packets arrive faster than they can be serviced, the number of backlogged users goes to infinity, and the delay becomes infinite. In Fig. 8 it is seen that by increasing the number of homenets N , a_N decreases, and the utilization that can




Fig. 9. The number of partitions which minimize the average delay as a function of the system utilization when the ratio of the maximum propagation time between sources in the entire network to the packet transmission time in an undivided system is between 0.01 and 100.




Fig. 10. The ratio of the propagation time in a homenet to the packet transmission in a homenet which results in the minimum delay when a system is partitioned as a function of the system utilization.

be supported by a network with finite delay increases. It is also seen in this figure that as the number of partitions increases the delay achieves a minimum, then starts to increase. The increase occurs because the time it takes to transmit a packet increases.

In Fig. 9 the number of partitions which result in the minimum delay is plotted as a function of the utilization for values of "a" between 0.01 and 100. It was found that if the network was divided the effective value of "a"

within the homenet a_N was approximately the same at a particular utilization for all of the a 's investigated. This effective value of " a " is plotted as a function of the utilization in Fig. 10.

C. Voice Capacity

In an MSTDM system a voice packet is a fixed size packet consisting of a data area that contains voice and overhead bits, an overflow area, and a preempt area. The time to transmit the voice samples in the data area is

$$\tau_D = \frac{D}{T_s} * \frac{b_s}{R_N}$$

where D is the time between packets, T_s is the time between voice samples, b_s is the number of bits per sample, and R_N is the transmission rate in the homenet. Defining the relative rate of the voice source to the channel transmission rate as $X_N = b_s/T_s * 1/R_N$, $\tau_D = D * X_N$. The overhead area of a packet consists of a fixed number of bits H , used for addressing and synchronization. The duration of the preempt interval τ_{PH} is a function of the maximum propagation delay in the homenet P_N and the number of bits C_D required to detect a collision. Defining $a'_N = P_N/\tau_D$

$$\tau_{PH} = 2 * a'_N \tau_D + \frac{C_D}{R_N}.$$

H and C_D are fixed numbers of bits included in packets independent of the transmission rate, propagation delay, or the delay between voice packets. Defining Y_D as the ratio of the fixed overhead bits to the data bits $(H + C_D)/((b_s/T_s) * D)$, the time to transmit the fixed segment of the packet is $Y_D \tau_D$. The time to transmit a packet is

$$T_x = \tau_D * (1 + 2a'_N + Y_D) + \tau_o$$

where τ_o is the time to transmit the overflow area. The overflow area is long enough to transmit all voice samples which arrive during a packet transmission time. Therefore,

$$\tau_o = \frac{b_s}{R} * \left\lceil \frac{T_x}{T_s} \right\rceil$$

where $\lceil x \rceil$ rounds x up to the next integer. Approximating $\lceil T_x/T_s \rceil$ by T_x/T_s , $\tau_o = X_N T_x$, and

$$T_x = \frac{X_N}{1 - X_N} D * (1 + 2a'_N + Y_D).$$

In addition to the time to transmit a packet the time between packets restricts the number of sources that can share the channel. In a completely full MSTDM system the next voice terminal acquires the channel as soon as it realizes that the previous terminal has completed its transmission. In a system with fixed size packets and a single reflection point a terminal can determine how long it takes its signal to propagate to the reflection point and can

arrange for its packet to arrive at the reflection point as soon as the previous transmission is complete. However, in this analysis it is assumed that a terminal does not start to transmit until it detects the end of the previous transmission. This creates an idle period between packets equal to the average propagation delay between terminals. For the remainder of this analysis, this idle period will be set equal to the maximum propagation delay between terminals, p_N and the effective transmission time for a packet is $T_x + a'_N X_N D$.

In an MSTDM system another voice connection can enter the system as long as it is possible to enter the system without a collision occurring. The number of simultaneous channels that can be supported is

$$C = \left\lfloor \frac{D}{T_x + a'_N X_N D} \right\rfloor = \left\lfloor \frac{1}{\frac{X_N}{1 - X_N} * (1 + 2a'_N + Y_D) + a'_N X_N} \right\rfloor.$$

In order to calculate the number of voice connections which can be supported by an MSTDM system, it will be assumed that the voice sources operate in a TASI mode [9] and only acquire the channel during active talk spurts. In a TASI system, the fraction of voice which is frozen out of the system due to all of the lines being occupied is [10]

$$\phi = \frac{1}{np} \sum_{k=C+1}^n (k - C) \binom{n}{k} p^k (1 - p)^{n-k}$$

where n is the number of voice sources and p is the probability that a voice source is in the active mode. This calculation for the freeze-out fraction assumes that if there is a vacant channel a voice source will obtain the channel.

The voice carrying capacity of an MSTDM network has been determined for systems with address and transmission parameters associated with commercially available networks [3] and with these parameters adjusted for voice operation [11]. Here, the effects of homenets on the efficiency of MSTDM systems is presented. The efficiency of the system, η , is defined as the number of voice sources supported by this system over the number of voice sources which would be supported if a separate frequency division multiplexed channel were dedicated to each digital voice source. The calculations have been performed assuming

- 1) 32 kbit speech coders;
- 2) the speed of light in the cable is 0.77 of the speed of light in a vacuum;
- 3) a freeze-out fraction of $\phi = 0.005$, as in [9];
- 4) a probability that a speaker is active of $p = 0.443$, as in [12]; and
- 5) $H + C_D = 64$ bits.

In Figs. 11 and 12 the efficiency is plotted as a function of the number of homenets for voice packets containing 10, 30, and 100 ms of voice samples. Fig. 11 is for a large network with 32 Mbits/s of capacity and a 20 km diame-




Fig. 11. The ratio of the number of voice connections supported by a speech activated MSTDM system to the number of voice connections supported by a dedicated channel FDM system as a function of the number of homenets for 10, 30, and 100 ms long voice packets for a 32 Mbit/s 20 km system.

ter, and Fig. 12 is for a smaller network with 10 Mbit/s of capacity and a 1 km diameter. The large network corresponds to the size and rate that may occur in a CATV system, and the smaller network corresponds to what is found in typical local area networks. These figures show the relative applicability of the homenet approach to CATV systems and local area networks. The delay D in these figures is the time to accumulate the packet. As the delay increases there are more voice samples per packet, and the system becomes more efficient.

In the large network the system must be partitioned before the MSTDM system can accommodate as many voice connections as a system with dedicated channels. The number of homenets required to achieve equity depends upon the delay that can be tolerated in the voice path. If only 10 ms of delay can be tolerated, at least 9 homenets are required, whereas if 100 ms of delay can be tolerated, only 2 homenets are required. These systems are more efficient than a TDM system without any overhead because they use a TASI approach to eliminate silent intervals and share the channel. As the number of homenets increases, the ability to share the channel decreases. Eventually, when the bandwidth of the homenet is reduced to one voice source, there is no sharing, and the overhead bits must make this system less efficient than TDM.

In the small network, the ability of a packet addressed system to transmit packets only when there are significant




Fig. 12. The ratio of the number of voice connections supported by a speech activated MSTDM system to the number of voice connections supported by a dedicated channel FDM system as a function of the number of homenets for 10, 30, and 100 ms long voice packets for a 10 Mbit/s 1 km system.

voice samples always results in a more efficient system. In this network there is little to be gained by dividing the system into smaller homenets or by delaying the voice by more than 30 ms. Therefore, it is reasonable to design a CATV system based upon homenets with a 1 km diameter, 10 Mbit/s channels, and 30 ms of delay in the voice path, as suggested in [10].

Thus far, the effects of homenets on voice and data systems have been analyzed separately. The effects of voice channels on the data delay and the effects of data contention on the time to acquire a voice slot have not been analyzed. These effects cannot be determined by the simple models used here. In [2], simulation results are reported which indicate what these effects are in an MSTDM network.

IV. CALL CONNECTION

Each source transmits on the downstream channel and receives signals in its homenet frequency band to detect collisions. A destination receives signals in the source's homenet frequency band. If a destination receives signals in every homenet frequency band, any source may address

packets to this destination. However, this causes the receiver complexity to increase as the number of homenet increases. By switching a destination's receiver to a specific homenet frequency band, one connection at a time can be maintained, and the receiver complexity is relatively independent of the number of homenets. However, a strategy must be developed to tell the destination which band to receive signals in. In this section three strategies are described, one which is completely distributed among the destinations, one which is distributed among the homenets, and one which is centralized at the head end of the CATV system. In each of the strategies the utilization of the homenets is monitored to determine if new connections can be established.

In the completely distributed strategy separate transmit and receive signaling channels are used to establish connections. When a destination is not busy it listens to the signaling receive channel instead of one of the homenet channels. When a station wants to establish a new connection it transmits a call packet to the appropriate destination in the signaling transmit band. The packet identifies the calling station and its homenet frequency. If the destination is not busy it receives the call packet and transmits a response, including its homenet frequency. This completes the connection, and both stations start communicating on their respective homenet frequencies.

The signaling channel is a contention channel with the reflection point at the head end of the CATV system. The propagation distance in this channel is too large to use CSMA/CD. A random access strategy similar to that in the ALOHA system [6] is used on this channel. In the ALOHA system a source knows that it has successfully acquired the channel and stops retransmitting the packet when it receives an acknowledgment from the destination. In this system the destination does not always listen to this channel; therefore the source cannot depend upon receiving an acknowledgment. However, there is one reflection point in this system, and all sources receive the same reflected signal. This enables a source to tell if it has successfully contended for the channel by listening to the receive channel for its packet. It is possible that a successfully transmitted packet is received with errors at the destination. Therefore, a source must retransmit successful packets until it is reasonably certain that the receiver is not listening to the signaling channel.

If the probability of having to retransmit a packet is P_r , the probability of stopping after the i th attempt is $(1 - P_r)P_r^{i-1}$, the average number of attempts is $1/(1 - P_r)$, and the probability of not completing the transaction by the K th attempt is P_r^K . The call packet must be repeated until the destination successfully receives this packet and the calling station successfully receives the response. If the probability of error is P_e , $P_r = 1 - (1 - P_e)^2$, the average number of times the call is transmitted is $1/((1 - P_e)^2) \approx 1 + 2P_e$, for $P_e \ll 1$, and the probability of not completing in K attempts is $P_e^K(2 + P_e)^K$. For the response packet $P_r = P_e$, and the average number of transmissions is $1/(1 - P_e) \approx 1 + P_e$. Therefore, the average number of

packets per successful connection is $\bar{n}_c \approx 2 + 3P_e$. Unsuccessful connections occur when the destination is busy and not listening to the signaling channel. An unsuccessful connection requires K transmissions where $P_e(K)$ is sufficiently small that connections are usually completed if the destination is not busy. If the probability that a destination is busy is P_B the average number of packets transmitted on the signaling channel per successful connection is $\bar{n}_s \approx 2 + 3P_e + P_B/(1 - P_B)K$.

The average number of packets transmitted on a homenet per successful connection is $\bar{n}_{HN} = P_v S_v + (1 - P_v) S_d$ where P_v is the probability that a connection is a voice connection, S_v is the average number of packets transmitted in a voice connection, and S_d is the average number of packets transmitted in a data connection. The utilization of a homenet must be less than 1, therefore the utilization of the signaling channel by a homenet is less than $\tau_s \bar{n}_s / \tau_{HN} \bar{n}_{HN}$ where τ_s and τ_{HN} are the packet transmission times on the signaling channel and on the homenet. The total utilization of the signaling channel must be less than $1/2e$, therefore the number of homenets which can be supported by a signaling channel is

$$N = \frac{1}{2e} \frac{\tau_{HN}}{\tau_s} \frac{P_v S_v + (1 - P_v) S_d}{2 + 3P_e \frac{P_B}{1 - P_B} K}.$$

To obtain a conservative estimate of the number of homenets which can be supported by a single signaling channel the following will be assumed.

- 1) The probability of error $P_e = 10^{-3}$.
- 2) The number of retries before quitting is $K = 3$. This results in fewer than 1 in 10^8 calls which can be completed not being completed.
- 3) During the busy period half the stations called are busy ($P_B = 0.5$). This results in an average of 5.003 packets being transmitted on the signaling channel per connection.
- 4) A voice call lasts 3 min, a packet is transmitted every 30 ms, and each of the two speakers is active 0.443 of the time. An average of 5300 packets is transmitted on the homenets per voice connection.
- 5) Half the connections are voice connections, and there are many fewer packets in a data connection than in a voice connection, so that an average of 2650 packets is transmitted on a homenet per connection.

- 6) It takes as long to transmit the control packets used on the signaling channel as it does to transmit the data packets on the homenet.

For these parameters, a signaling channel can be used by at least 98 homenets. This should be adequate for most systems.

This distributed strategy has the advantage that stations can be moved between homenets without telling anyone because the same signaling channel is received everywhere in the network. It has the following disadvantages.

- 1) The transmitter must be able to transmit in both the communications band and the signaling band.

2) Two additional channels must be allocated for signaling.

3) A station cannot receive an indication that another station is trying to call while it is busy.

In addition, each station must monitor its homenet to determine if there is sufficient capacity to establish another connection or must communicate with a monitor on its homenet before establishing a connection. These disadvantages are removed by the centralized strategies.

In the partially distributed strategy, each station transmits call requests and receives the response as data packets in its own homenet. On each homenet there is a monitor, and the monitor has a receiver on every homenet. The monitor measures the utilization of each homenet to determine which can support additional connections. It also receives all of the connection packets for terminals in its homenet and forwards them in the homenet frequency band. Since each transmitter must listen to its homenet to detect collisions, these packets can be received even if a call is in progress. When a terminal moves to a new homenet, it must notify the monitor so that connection packets can be forwarded.

In the centralized strategy there is only one monitor, at the head end of the CATV network. At the head end of the network, signals from each homenet must be received and retransmitted in the homenet frequency band. In this strategy these signals are monitored to determine the utilizations of the homenet and to receive connect packets. In addition, before each homenet frequency is retransmitted it is delayed by a packet transmission time. When a connect packet is received the monitor determines which homenet the destination is on, waits until there is an unused packet in that homenet frequency band, and inserts the connect packet, as in a register insertion loop [13]. At the reflection point of each homenet the homenet frequency band from the head end is received before it is notch filtered. Connect packets are received, and are forwarded in the homenet frequency band.

V. ALTERNATIVE NETWORKS

The system which has been described will operate on an existing CATV network. However, the reliability and maintainability of this system can be improved by modifying the equipment and topology of the CATV network. Two possible ways of doing this are code operated switches [14] and hub configured networks [15], [16]. In addition, noise problems in the network can be reduced by regenerating, rather than just amplifying the digital transmission channels.

Code operated switches block signals from a small portion of the network and are remotely controlled. Two sources of failure in this type of network are transmitters which misbehave and transmit continuously and local noise sources which make the entire transmit channel unuseable. The notch filters on the transmit band isolate this type of a problem to a homenet. If this type of a problem is detected

within a homenet, code operated switches can be used to isolate the fraction of the homenet causing the problem. This allows the remainder of the homenet to continue operating and reduces the area which must be examined for the problem.

In a hub configured network, trees, consisting of a small number of users, are rooted at hubs. The hubs are interconnected to form a network. In this system the hubs are the reflection points of the homenet and are interconnected by multiple paths on a backbone network. Since a number of small trees are rooted at each hub, a problem can be isolated to one of the trees. At the hub the signals received from the other homenets are combined with the signal for this homenet and transmitted on the trees emanating from the hub and all of the outgoing links. Since the backbone network is multiply connected, it can be designed to survive link or hub failures without disconnecting hubs which are still operating. Hub failures can be isolated so that they only effect the terminals on the hub's homenet.

VI. CONCLUSION

A system has been described which uses a random access broadcast protocol to transmit digitized voice and data on a CATV network. The system extends a local area network access protocol, CSMA/CD, so that it can be applied to a larger metropolitan area, and so that it can handle periodic voice traffic which has a severe delay constraint. It also provides a mechanism for interconnecting several local area networks without using store-and-forward gateways between the networks.

Switching in this system is completely distributed. Packets of data are addressed to a specific destination and appear everywhere in the network. Distributed switching allows the switch to grow linearly with the number of users and does not require a large amount of equipment for a small number of users. Since a packet appears everywhere in the network, anyone who knows the destination address can receive the packet, and a privacy problem exists. However, this problem can be solved by conventional data encryption techniques.

Voice signals in this system are digitized and assembled into packets. Since the packets are addressed, and the destination is always searching for the address, packet transmission can be discontinued during inactive voice intervals without increasing the complexity of the receiver. This allows a simple TASI system to be implemented. In a fixed channel system the channels would have to be dynamically allocated by a central controller, the destination would have to be told which channel each talk spurt is transmitted on, and the destination would have to switch to that channel. This system inserts a delay in the voice path while packets are assembled. However, the analyses indicate that a small delay of 30 ms or less is adequate.

Data signals in this system can obtain any throughput they can successfully contend for, up to the capacity of the homenet. They are not restricted by the capacity of a fixed

channel. A low data rate device, possibly performing a monitoring function, does not have to tie up a fixed channel designed for a higher rate device. It can transmit packets as infrequently as it needs to, and when it is not transmitting packets the capacity is available for other devices. And, a high rate device can obtain more capacity when it is available.

The transmission strategy used in this system allows all of the sources to transmit in one and the same frequency band, so that frequency agile transmitters are not required, and the same transmitter is used in all of the homenet. The receiver must be able to select a frequency band. In a system which has been implemented the homenet frequency bands were selected to correspond to standard TV channels. The frequency selection was performed by the same device that is used to select channels in standard CATV networks, and the frequency translators and filters were conventional CATV equipment.


REFERENCES

- [1] R. M. Metcalf and D. R. Boggs, "Ethernet: Distributed packet switching for local computer networks," *Commun. ACM*, vol. 19, no. 7, pp. 395-404, July 1976.
- [2] N. F. Maxemchuk, "A variation on CSMA/CD that yields movable TDM slots in integrated voice/data local networks," *Bell Syst. Tech. J.*, pp. 1527-1551, Sept. 1982.
- [3] —, "Some characteristics of movable slot TDM," in *Proc. 8th Conf. Local Comput. Networks*, Minneapolis, MN, Oct. 1983.
- [4] A. N. Netravali and Z. L. Budrikis, "A broad-band system for wide area local access," to be published.
- [5] I. Kong, "Cablenet: A local area network," in *Proc. ICC '82*, June 1982, pp. 6C.2.1-6C.2.5.
- [6] N. Abramson, "The Aloha system—Another alternative for computer communications," in *Fall Joint Comput. Conf. AFIPS Conf. Proc.*, vol. 37, 1970, pp. 281-285.
- [7] S. S. Lam, "A carrier sense multiple access protocol for local networks," in *Computer Networks 4*. Amsterdam, The Netherlands: North-Holland, 1980, pp. 21-32.
- [8] W. Bux, "Local-area subnetworks: A performance comparison," *IEEE Trans. Commun.*, vol. COM-29, pp. 1465-1473, Oct. 1981.
- [9] K. Bullington and J. M. Fraser, "Engineering aspects of TASI," *Bell Syst. Tech. J.*, vol. 38, pp. 353-364, Mar. 1959.
- [10] C. J. Weinstein, "Fractional speech loss and talker activity model for TASI and for packet-switched speech," *IEEE Trans. Commun.*, vol. COM-26, pp. 1253-1257, Aug. 1978.
- [11] N. F. Maxemchuk and A. N. Netravali, "A multifrequency multi-access system for local access," in *Proc. ICC 83*.
- [12] P. T. Brady, "A technique for investigating the on-off patterns of speech," *Bell Syst. Tech. J.*, vol. 44, pp. 1-22, Jan. 1965.
- [13] M. T. Liu, "Distributed loop computer networks," in *Advances in Computers*, M. C. Yovits, Ed. New York: Academic, 1978, pp. 163-221.
- [14] J. B. Wright, M. P. Block, and D. S. McVoy, "An evolutionary approach to the development of two-way cable technology communication," vol. CATV-2, no. 1, pp. 52-61, Jan. 1977.
- [15] E. J. Callahan, "The design, implementation and operation of bidirectional wideband CATV networks for two-way video audio and data services," in *Two-Way Cable Television: Experience with Pilot Projects in North America, Japan, and Europe*, W. Kaiser, H. Marko, and E. Witte, Eds. Berlin-Heidelberg: Springer-Verlag, 1977, pp. 75-82.
- [16] K. H. Voge, "The two-way CATV laboratory project of the Heinrich-Hertz-Institute Berlin," in *Two-Way Cable Television: Experience with Pilot Projects in North America, Japan, and Europe*, W. Kaiser, H. Marko, and E. Witte, Eds. Berlin-Heidelberg: Springer-Verlag, 1977, pp. 267-278.



Nicholas F. Maxemchuk (M'72-M'82) has been at Bell Laboratories since 1976 and is currently Head of the Distributed Systems Research Department. He was an Adjunct Faculty Member at the University of Pennsylvania, Philadelphia, in 1981-1982 where he taught a course on computer communications networks. Prior to joining Bell Laboratories he spent 8 years at the RCA David Sarnoff Research Center.

Mr. Maxemchuk is Editor for Data Communication Systems of the IEEE TRANSACTIONS ON COMMUNICATIONS.



Arun N. Netravali (S'67-SM'78) received the B.Tech. (Honors) degree from the Indian Institute of Technology, Bombay, India, in 1967, and the M.S. and Ph.D. degrees from Rice University, Houston, TX, in 1969 and 1970, respectively, all in electrical engineering.

He worked from 1970 to 1972 on problems related to filtering, guidance, and control for the Space Shuttle. Since 1972 he has been at Bell Laboratories, Holmdel, NJ, working on various aspects of signal processing and computing. He

was Head of the Visual Communication Research Department from 1978-1983. Currently, he is Director of the Computer Technology Research Laboratory. He was a Visiting Professor in the Department of Electrical Engineering at Rutgers University, New Brunswick, NJ, and at the City College of New York, New York, NY.

Dr. Netravali serves on the Editorial Board of the PROCEEDINGS OF THE IEEE and is the Editor for Signal Processing and Communication Electronics of the IEEE TRANSACTIONS ON COMMUNICATIONS. He is a member of Tau Beta Pi and Sigma Xi.