

N90-22279

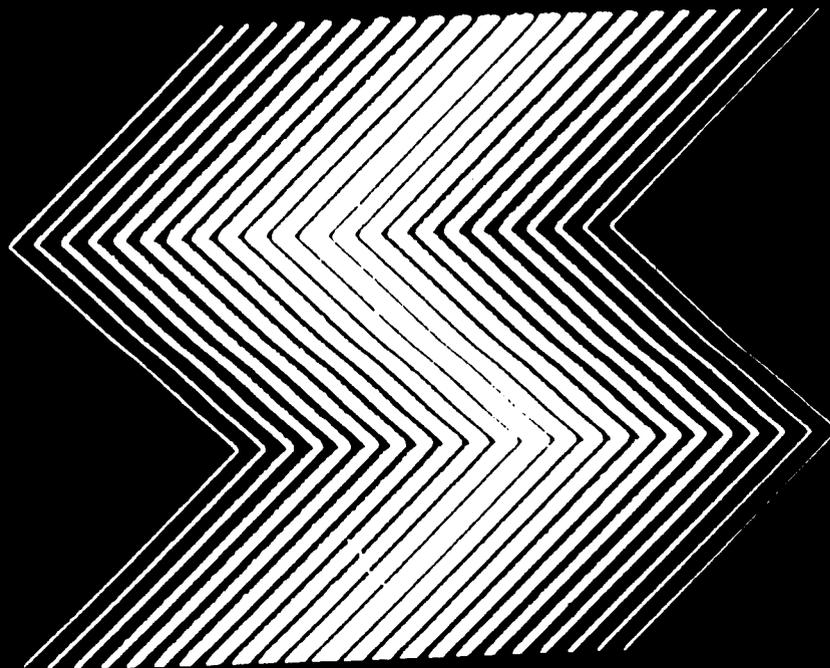
Computer Communications Review
Volume 19, Number 4
September 1989



SIGCOMM '89 SYMPOSIUM

Communications Architectures & Protocols

Austin, Texas
September 19-22, 1989



SPONSORED BY ACM SIGCOMM WITH SUPPORT GIVEN
BY THE INFORMATION SCIENCES AND TECHNOLOGY
CENTER OR SRI INTERNATIONAL.

Traffic Placement Policies for a Multi-Band Network*

*K. J. Maly E. C. Foudrial D. Game
R. Mukkamala C. M. Overstreet*
Department of Computer Science
Old Dominion University
Norfolk, Virginia 23529-0162

Abstract

Recently protocols have been introduced that enable the integration of synchronous traffic (voice or video) and asynchronous traffic (data) and extend the size of local area networks without loss in speed or capacity. One of these is DRAMA, a multiband protocol based on broadband technology. It provides dynamic allocation of bandwidth among clusters of nodes in the total network. In this paper, we propose and evaluate a number of traffic placement policies for such networks. Metrics used for performance evaluation include average network access delay, degree of fairness of access among the nodes, and network throughput. The feasibility of the DRAMA protocol is established through simulation studies. DRAMA provides effective integration of synchronous and asynchronous traffic due to its ability to separate traffic types. Under the suggested traffic placement policies, the DRAMA protocol is shown to handle diverse loads, mixes of traffic types, and numbers of nodes, as well as modifications to the network structure and momentary traffic overloads.

*This work was sponsored in part by contracts CITT-596041 from CIT and NAG1-1-908 from NASA Langley Research Center.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© 1989 ACM 089791-332-9/89/0009/0094 \$1.50

1 Introduction

Several recently introduced protocols illustrate the change in performance that results from subdividing network capacity into multiple channels. In addition, demand for integration of video and voice traffic with data traffic has resulted in protocols that allow for both synchronous and asynchronous traffic, such as DRAMA. The DRAMA protocol not only takes advantage of the multichannel efficiency but also allows for synchronous and asynchronous traffic over a large distance without significant loss in speed and capacity. This protocol is introduced in [15,16], and is based on a broadband technology, allowing for dynamic allocation of bandwidth among clusters of nodes, called local area network groups (LANGs), and dynamic allocation of synchronous/asynchronous traffic bandwidth.

Marsan and Roffinella [12] evaluate multichannel multiple-access schemes such as CSMA/CD for local networks, but standard CSMA/CD protocols do not provide for transmission of synchronous data as addressed in this paper. Chlamtac and Ganz [4] propose a multichannel design which statically allocates bandwidth capacity to avoid the simultaneous delivery of frames to a node from different channels, but does not allow for efficient use of unwasted bandwidth due to unbalanced traffic patterns. Merakos and Bisdikian [10] divide the capacity of the network into channels for intra- and inter-LAN traffic, but does not provide for dynamic allocation of the bandwidth allocated to each traffic type. Minimum delivery times for synchronous traffic cannot be guaranteed due to the use of bridges for interconnection of LANs. Wong and Yum [17] allocate channel bandwidth by a contention-based reservation protocol which provides circuit switching services for all traffic types, whereas, DRAMA only provides circuit switched services for synchronous traffic.

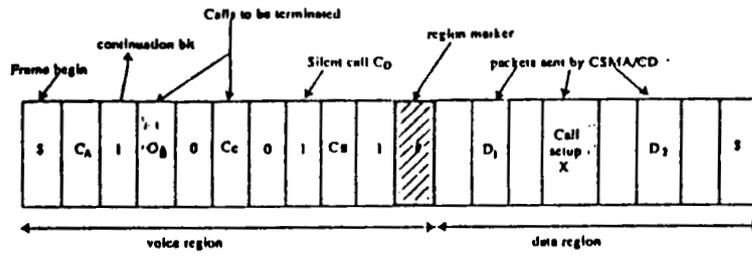


Figure 1: Sample Frame

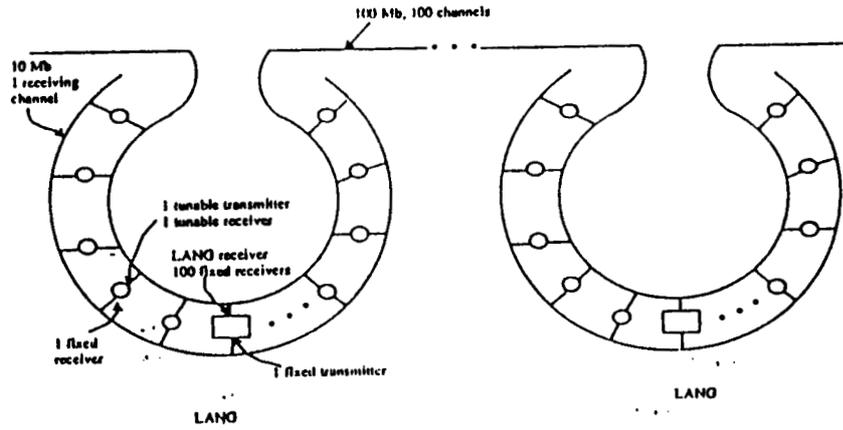


Figure 2: A possible overall network configuration

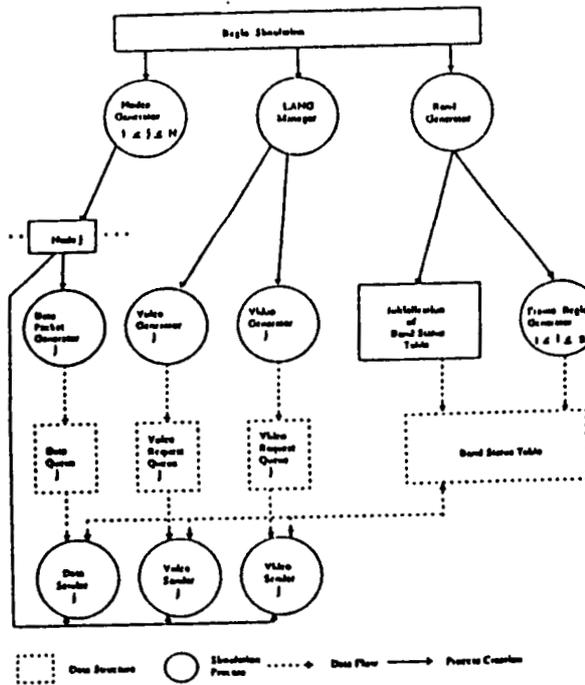


Figure 3: Structure of Simulation program

In [11], we presented the results of a collection of simulation studies of the DRAMA protocol. That paper focused on the ability of the protocol to reallocate bandwidth among LANGs in order to respond to changes in individual LANG workloads. The current paper reports on simulation studies modeling an individual LANG; in it, we analyze network performance in handling loads and evaluate how traffic placement policies affect performance. We assume a multichannel network based on CSMA/CD and wish to decide if the approach is feasible.

This paper is organized as follows. Section 2 describes the DRAMA protocol; Section 3 describes the simulation model used in these experiments; Section 4 presents the traffic placement policies studied here; and Section 5 discusses the results of these studies.

2 The DRAMA Protocol

The DRAMA protocol is designed for the dynamic sharing of bandwidth of a single broadband bus among groups of nodes in a large, integrated voice-video/data network. The amount of bandwidth available is assumed to be large, say 350-500 MHz. The nodes, each capable of transmitting all traffic types, are clustered by distance and function into LANGs. This type of clustering is typical for various locations of a company within a particular city, installations on a large ship or military base, or among different departments within a university. The cable bandwidth is frequency-divided into bands dedicated to particular LANGs and a global pool of bands that may be acquired by any of the LANGs. For each LANG, requesting, acquiring, or releasing a band depends on the current distribution and amount of traffic within that LANG relative to the current traffic within the entire network. A LANG is allowed to transmit only on those bands which have been assigned specifically to it, but is required to receive on all bands. For a more detailed discussion of DRAMA, including error recovery, the reader is referred to [15,16,11].

Basic design objectives of DRAMA are the integration of synchronous/asynchronous traffic and dynamic bandwidth allocation.

2.1 Synchronous/Asynchronous Transmission Protocol

This section briefly describes the DRAMA protocol for a LANG. In the protocol, the fraction of a band's capacity allotted to each traffic type depends on the current synchronous and asynchronous load. Time on each band is slotted into frames. Each band's frames

are delimited by "frame-begin" markers broadcast by the band's current *band-leader* and are partitioned into voice/video and data regions. The boundary between the two regions varies from frame to frame, depending on the number of voice calls in talk-spurts during that frame. Either data or voice/video may consume the entire frame if no traffic of the other type is present.

The data region is composed of data packets and call-setup requests. The bandwidth in the data region is allocated using a CSMA/CD traffic placement policy. Normally CSMA/CD is not suitable for use among nodes separated by more than several kilometers because the interval during which a collision can occur is directly proportional to the propagation delay between the most distant nodes. DRAMA circumvents this problem by restricting transmission privileges on a band to exactly one LANG at a time, while allowing all LANGs to receive the transmissions. In this way, the CSMA/CD-based protocol can be used over the entire set of LANGs with the same efficiency as in a single LANG.

The voice/video region provides a virtual circuit for each established (one-way) voice call. In the multi-channel version this means that a two-way inter-LANG call uses two bands, one for each direction of the call. One varying size slot is allocated in the voice region to each (one-way) voice circuit. The slot contains a varying size voice packet followed by control information, called the *control byte*. The slot size may differ because silence periods, which comprise roughly 60% of an average voice conversation [3], are not transmitted. The control byte informs the other nodes whether the call will terminate after this frame or will be continued. The slots for the different voice calls are contiguous and precede any data transmitted in the frame.

Figure 1 shows a frame in which five calls are ongoing (A,B,C,D,E), one of which is silent (D). Three packets are transmitted in the data region using CSMA/CD (data packets 1 and 2, and a successful call-setup for the next frame).

Since in CSMA/CD the amount of bandwidth wasted due to collisions increases dramatically with load, it is important to allocate bandwidth in a fair manner. Therefore, reducing load on heavily used channels at the cost of increasing it on lower utilized channels should reduce total collisions.

2.2 Dynamic Bandwidth Allocation

In proposed broadband systems such as [9] and [13], as well as in currently available commercial systems, the broadband frequency spectrum is statically parti-

tioned by user group and/or by traffic class. For example, CableNet [14], Sytek's LocalNet [6], and Mitre CableNet [8] each partition the bandwidth into fixed bands for particular applications of specific groups of users; some bands are permanently reserved for video channels, some are reserved for time division multiplexing for a set of closely located users, and some bands are dedicated to specific functions such as process control. However, a common characteristic of a network that supports diverse traffic classes such as voice, data, and video is that the bandwidth requirements of both an individual node and a LANG fluctuate widely over time compared with the LANG's average requirement. In such diversified systems, static partitioning according to average requirements will often waste idle bandwidth; at other times, it will be insufficient to satisfy a LANG's traffic requirements while bandwidth is available elsewhere in the network.

In the DRAMA protocol the bandwidth is frequency-divided into $M + 1$ bands. One band is reserved for a slotted band-control channel used by all the LANGs to coordinate band-sharing, and the remaining M fixed-size bands (say, 10MHz each) are available for voice, data and video transmissions. The M bands are partitioned into a set of *dedicated bands* and a set of *available bands*. Dedicated bands guarantee that no LANG "starves" and that each LANG's performance is at least that of a normal, solitary LANG using a baseband cable.

The available bands are either assigned to a LANG or are in a global pool. Bands in the pool are shared via a dynamic, fully-distributed, band-sharing policy that allows each LANG to obtain global bands based on three factors: its current needs, the current needs of other LANGs, and the current availability of global bands. When a LANG acquires a band, that LANG's nodes have exclusive transmission rights to the band; each node in the system must be able to receive all bands that might contain packets addressed to it. Because all bands can be received at all nodes, a uniform communication mechanism exists among all nodes in all LANGs; this is preferable to the use of gateways, which introduce additional delays and buffering requirements since transferred traffic must compete with local traffic when it is sent between nets. Simple wide-band repeaters may be required over a long distance network to overcome attenuation and/or signal distortion. The details of this band allocation can be found in [11].

2.3 Traffic Placement

When a node wishes to transmit, it chooses among the bands on which its LANG currently has transmission

privileges; the node seeks a band with no traffic. If the LANG owns 10 bands, a typical figure, we estimate this to take about 20 microseconds. We refer to the policy for choosing among the free transmission bands as the *traffic placement policy*.

Given the frame format in Figure 1, two primary objectives in formulating a placement policy exist. First, it is important to keep bands free of voice/video traffic if possible since a band cannot be freed if its release would interrupt synchronous traffic. Second, in order to minimize delay, data traffic is spread as evenly as possible over all the LANG's current bands. The systems and algorithms that control this traffic placement policy and its resulting performance are the subject of this paper.

2.4 DRAMA Implementation

In the original DRAMA system each node needed as many receivers as there were channels and enough tunable transmitters to service the channels assigned to its LANG. This is not economically feasible. Figure 2 presents a possible solution to this cost problem by showing how nodes in a LANG can share receivers and how the number of transmitters per node can be limited. The network sketched has 100 Mb total capacity, divided into 100 channels. The LANG illustrated has 10 bands assigned to it. Each node needs at least one tunable transmitter/receiver pair to be able to determine whether a selected band is free and to detect a possible collision on that band after it has begun transmitting. Nodes can have more than one tunable transmitter if they need to send information on more than one channel simultaneously. All incoming information is handled by the LANG central receiver system and is forwarded to the nodes in that LANG by a secondary channel that connects all nodes, or, alternatively, by twisted pair (not shown). In the system illustrated, the LANG receiver has 100 fixed receivers to listen to all channels; it filters all messages belonging to its LANG (including the ones sent from within the LANG) and channels them to the LANG net. Hence the total number of signaling devices in a LANG with n nodes is $3n + 100$.

3 Simulation Model

We used simulation to study performance of the DRAMA protocol with different traffic policies. The network model and related protocols were written in SIMSCRIPT II.5. The simulation program is highly parameterized to allow experimentation with different loads and different traffic placement policies.

Initial experimentation showed that the model achieved steady-state in 1.8 seconds of simulation time; thus for each experiment, data collection began after this point. Three techniques were used for model validation. First, traces – that is, printing sequences of significant model state changes – satisfied our concerns about correct program implementation. Second, as indicated by figures in this paper, the model replicated standard CSMA/CD behavior. Third, as we altered input parameters the model responded appropriately. We did occasionally encounter what was, at least initially, counterintuitive behavior in experiments, but further analysis showed in each case that the model had behaved appropriately; it was our intuition that failed us.

4 Model Objectives and Structure

The main objectives of our simulation experiments are:

- To create an implementation of the DRAMA protocol at the packet level as it would operate in a single LANG,
- To compare several traffic placement policies, and
- To measure performance under a variety of loads and traffic placement parameters.

Our earlier studies [11] dealt with performance issues corresponding to band allocations among the LANGS in a network. In those studies, a LANG was modeled as an abstract entity with varying needs for bands. In this study, however, we concentrate on the performance issues related to a single LANG (as opposed to a network of LANGs). Accordingly, a LANG is modeled as a group of nodes with the ability to communicate to nodes across the network. Restricting simulation to a single LANG considerably reduces the execution time required for each run. The conclusions determined for a single LANG are easily generalized to the multiple LANG case since each LANG operates independently.

In our earlier study [11], we have shown that bands can be rapidly reallocated such that the band utilization at every LANG is kept within a small percentage of the total network average. In this study, we want to hold the number of bands constant and study the effects of traffic placement on the network.

As shown in Figure 3, the simulation model includes:

- **Band Status Table:** Each node maintains a band status table for all bands. Each entry of this ta-

ble indicates the status (dedicated/global) of each band.

- **Request Queues:** Each node maintains three queues, one each for data, voice setup, and video setup. A first-in-first-out policy is used to service each queue.
- **Sender Units:** In addition to the request queues, each node maintains units to control the sending of voice, video and data blocks. These buffers contain the packet that is being currently transmitted as data.

The simulator initially generates the n nodes corresponding to a single LANG. For each of these n nodes, data traffic is created in terms of data packets as opposed to data messages of variable length. The data generator assumes Poisson arrivals of packets that are then placed in a data queue at the appropriate node. For efficiency, voice and video traffic are generated at the LANG level (described as LANG manager in Figure 3), and then randomly assigned to nodes in the LANG. Both video and voice traffic have Poisson arrivals and exponential service times. When a data packet, voice request packet, or video request packet is placed in its respective queue, the data frame controller checks with the band status table to determine whether a band is available for the packet. The time required to check all bands is called the *band choice thinking time*. A packet is not removed from its queue until after it has been placed and the collision interval has passed.

The video/voice/data traffic mix generated in a LANG is described as a percentage of the total load in the LANG. Once these percentages are chosen for a particular experiment, the average arrival rate for each traffic type at each node (λ_{DT} , λ_{VO} , λ_{VI} respectively) can be determined using the following system of equations:

$$\lambda_{DT}\mu_{DT} + \lambda_{VO}\mu_{VO} + \lambda_{VI}\mu_{VI} = \frac{C_u}{n} \quad (1)$$

$$p_{DT} + p_{VO} + p_{VI} = 100 \quad (2)$$

$$\lambda_{DT}\mu_{DT} = \frac{p_{DT}C_u}{100n} \quad (3)$$

$$\lambda_{VO}\mu_{VO} = \frac{p_{VO}C_u}{100n} \quad (4)$$

$$\lambda_{VI}\mu_{VI} = \frac{p_{VI}C_u}{100n} \quad (5)$$

where

- C_u is the assumed average channel capacity used by the LANG,

- p_{DT} , p_{VO} , and p_{VI} are the desired percentages of data, voice, and video traffic on the channel respectively, and
- μ_{DT} , μ_{VO} , and μ_{VI} are the the given average service times of data, voice, and video traffic respectively.

Some simulation experiments studied the effect of traffic mix. Equations (1) through (5) were used to generate arrival rates. For example, a traffic mix of 15% voice, 25% video and 60% data and a total network load of 60% results in individual network loads of 9%, 15% and 36% for voice, video and data respectively. The data arrival rate necessary to generate its appropriate load is further dependent upon the length of the specific packet. Data traffic is then uniformly distributed among the nodes of the network.

Other simulation experiments generated varied mixes of voice, video and data traffic in order to test the DRAMA system's ability to handle different traffic conditions. However, in many runs where other features of the protocol were being studied, a typical traffic mix of 15% voice, 25% video and 60% data was used. Standard transmission rates of 64 Kb/s for voice and 500 Kb/s for video were used for each circuit. The justification for using the higher video load is that a single video transmission occupies half a band, making it harder to place this kind of traffic than to place equivalent voice traffic. We felt that the increased video traffic would cause greater disruptions since these calls would occupy a large block on a band, making it more difficult to place data traffic and would tend to occupy more bands less effectively. The results would, in turn, have a greater tendency to show where problems in traffic integration would occur. Finally, most of the tests used a capacity of 10 Mb/s for the entire network, implemented as ten 1 Mb/s bands.

We considered the following performance metrics for system evaluation:

1. *Access Delay*: Due to the significance of network access delay in our experiments, we refer to this as *access delay*. Thus the access delay does not include the transmission and propagation delays. δ_i indicates the average access delay at node i .
2. *Fairness*: The network access delay for each of the n nodes should be independent of its position in the network and should be close to the average access delay E of the network. This is measured by the degree of fairness D_f , where

$$D_f = \sqrt{\frac{\sum_{i=1}^n (\delta_i - E)^2}{n}} \quad (6)$$

where

- δ_i is the mean access delay at the i th node,
- E is the mean access delay in the LANG,
- n is the number of nodes in the LANG

Ideally, D_f should be zero. Note that this equation is used to measure fairness within a single LANG. Fairness among LANGs was studied in [11].

3. *Throughput*: We measure the throughput of a LANG as a function of the offered load, which is measured as a percent of network capacity. In these studies offered load ranged from 0 to 200 percent. Here, throughput is the percent of the network capacity taken up with successful traffic.
4. *Recovery time*: We were interested in the response of a single LANG to impulse traffic. To this end, in some experiments a sudden pulse of traffic was generated in order to determine how long the system would take to return to within 5 percent of the nominal load. Ten percent of the nodes were given a burst of additional data traffic in order to build up the queues at these nodes.

5 Traffic Placement Policy

As previously stated, the traffic distribution is a mix of video, voice and data traffic. Our simulation has separate traffic generation procedures for each type of traffic. In an effort to minimize the number of bands carrying synchronous traffic, new synchronous traffic is assigned to the lowest numbered band carrying synchronous traffic upon which it will fit. Video and voice traffic is handled essentially as described in the DRAMA protocol. Because of the short length of the simulation runs relative to the length of video and voice transmissions, these traffic demands are fairly static. We paid more attention to the placement of data traffic because its dynamic nature has greater effect on the network. We also studied the response to the load changes that would be induced by starting or stopping a video transmission.

By examining a number of placement strategies, we intended to determine the sensitivity of the network's performance to various methods of placing data traffic on the network.

5.1 Traffic Placement Strategies

Traffic placement policies are concerned with techniques to compete for transmission time and to handle

unsuccessful transmissions. More specifically, a traffic placement policy must address the following situations:

1. A node would like to transmit but finds all bands busy. The basic question is how long the node should pause before reexamining the bands. Strategies vary from no delay to waiting a large amount of time under the assumption that the network is currently saturated. The time required to scan all bands, which we call *scan time*, is the minimum possible delay. In some policies an additional delay is added in consideration of other nodes' attempting to transmit. Having a multiband network with no available bands is a much better indication of the state of the network than finding a single-band Ethernet network busy. We believed that this information might be used to devise better placement policies.
2. A node attempts to transmit but a collision occurs. A collision does not necessarily indicate that the network is busy, only that a band is busy. When a node attempts to (re)transmit a packet, it examines all bands to determine which ones are not busy. With light traffic it is better to attempt immediate retransmission. Under high loads, a delay may reduce collisions.
3. A node transmits successfully and has additional packets in its queue to be sent. This state is similar to the condition for busy bands, but has no information about the state of the network.

We studied four basic strategies. The simplest, *fixed*, is a static policy that delays a constant amount of time before retransmitting. We used a minimum of 20 microseconds for this figure. Since the policy assumes no knowledge of network load, it provides a baseline for comparison with other methods that use local or global knowledge in order to improve performance. Each of the three situations above employs this delay.

Speedup bases the delay between attempts upon the node's own queue size, using only local knowledge. The delay is inversely proportional to the queue length. In this approach the node transmits packets more frequently as its queue grows and incorporates an opposite philosophy of *backoff*. Under light network loading, a node will empty its queue more quickly by transmitting more often. We assumed that this method would reduce delays at low loads but that performance might suffer at high loads. *Speedup* uses this calculation for each of three situations.

Backoff employs the binary backoff approach, in which the delay time increases as more collisions oc-

cur. This method has been shown to be very successful on CSMA/CD networks during heavy loads, and would give us a point of comparison for alternative strategies under similar adverse conditions. In non-collision situations *backoff* waits for a time not less than *scan time* and not more than twice the value of *scan time*.

Given our intuition that *speedup* would be preferable in light loads and *backoff* in heavy loads, we also investigated a combination of the two, which we called *tempered backoff*. If a node is experiencing few collisions, *tempered backoff* will approximate *speedup*. As the number of collisions increases, the delay incorporated into the formula for *backoff* will quickly dominate the term for *speedup* and will give the characteristics of a *backoff* approach.

6 Results

Figure 4 illustrates DRAMA's performance compared with equivalent single-channel CSMA/CD protocols. The most important aspect shown in Figure 4 is that the access delay of the multiband DRAMA system is considerably lower than that of equivalent or higher bandwidth single-band CSMA/CD systems. Figure 5 substantiates this point and demonstrates that as the number of bands in the LANG increases from 1 to 20, the access delay for integrated voice/video/data traffic decreases significantly. One can conclude that multiple-band local area network systems are able to successfully handle integrated traffic using a CSMA/CD-type protocol. Also, DRAMA's performance is for networks that can cover up to 100 km whereas single-band higher frequency CSMA/CD protocols work only for a few km.

Specifically, the curves of Figure 4 depict the average access delay of a data package. As expected, the delay is high even for low loads in a 50 Mb/s channel without synchronous traffic. Here, the collision slot time is a large percentage of the time it takes to send a packet so collisions become very costly. The 10 Mb/s channel with synchronous traffic uses a framed Ethernet, where framing is necessary in order to provide guaranteed access for the synchronous information. The cost of incorporating synchronous traffic is significant, since the framing creates a point, located immediately after the voice/video frame terminates, where the probability of collisions is high.

In contrast, DRAMA effectively separates the synchronous effects caused by framing so that data packets generally have immediate access to at least some bands. In addition, DRAMA provides for efficient recovery of that portion of the frame which is not used

DATA DELAY
IN MICROSECONDS

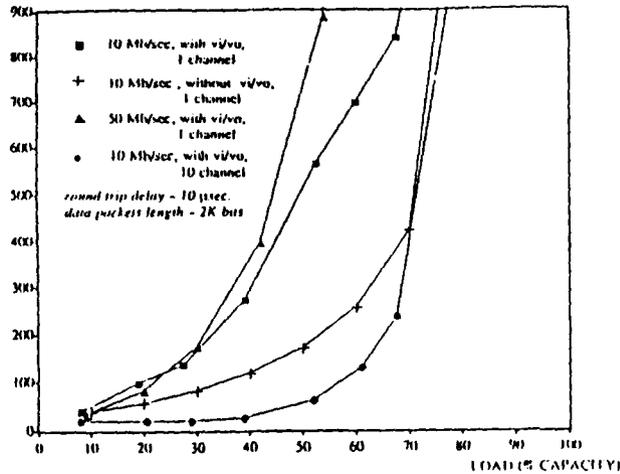


FIGURE 4: One channel DRAMA with and without video/voice compared with multichannel DRAMA

DATA DELAY
IN MICROSECONDS

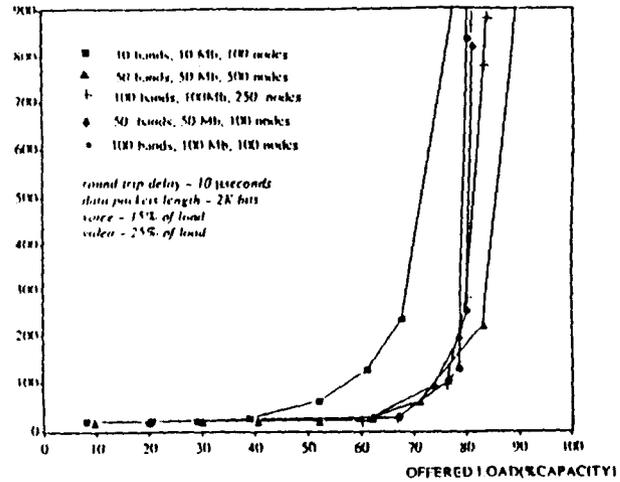


FIGURE 7: Scaling of network by bandwidth, number of bands and nodes

DATA DELAY
IN MICROSECONDS

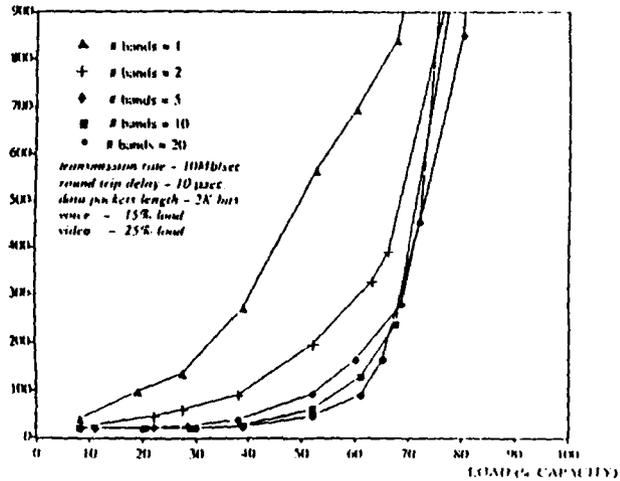


FIGURE 5: Multiple channel versus single channel

TOTAL NUMBER OF
PACKETS IN QUEUES

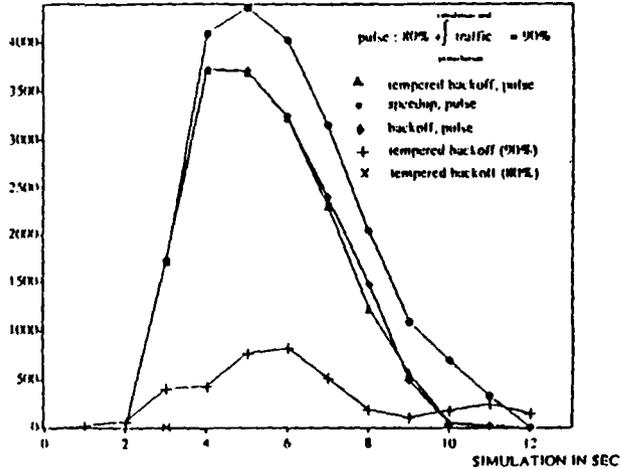


FIGURE 8: Reactions of various algorithms to strong overload at 10% of nodes

NETWORK THROUGHPUT
(% CAPACITY)

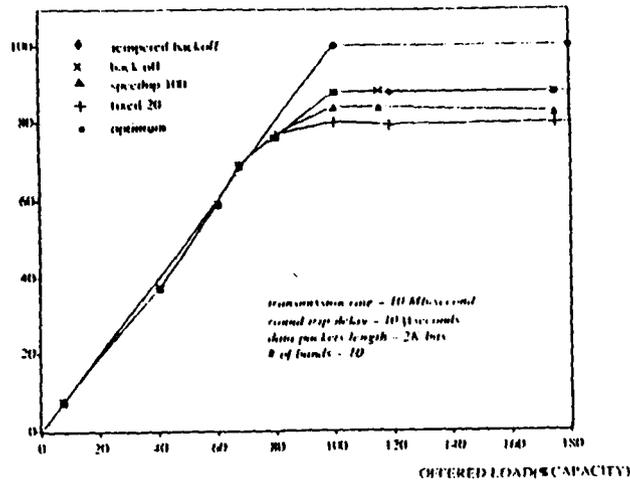


FIGURE 6: Network throughput for various traffic placement algorithms

for synchronous messages.

Figure 5 shows the effect of changing the number of bands assigned to a LAN. For instance, with two channels of 5 Mb/s each (of which 2 Mb/s of load are due to voice and video), we have approximately the same delay as on one 10 Mb/s channel without synchronous traffic (see Figure 4). For the LAN with 100 nodes, ten 1 Mb/s bands appear to be a good compromise, since delay does not improve greatly after that point. There is clearly an optimal number of bands; as the bandwidth per band gets smaller, the negative impact of both framing, as we shall see below, and delivery delay become more evident.

Note that we are comparing access delays only. In the overall performance of a network, additional transport delay ($5 \mu\text{sec}/\text{km}$) and bandwidth delay (bits/sec) affect the arrival of the complete packet at its destination. Thus, a 2k packet at 10 Mb/s would take $205 \mu\text{sec}$ to arrive at a station 1 km from the source, while at 1 Mb/s it would take $2050 \mu\text{sec}$. In fact, one can conclude that single-band CSMA/CD is always superior in overall delay to multi-band CSMA/CD if all traffic is asynchronous. Hence, the price one pays in order to have an integrated CSMA/CD that can handle voice, video and data effectively is to accept this increase in delivery delay. However, a 2 msec transmission delay for a 2K packet, while it may be serious in some high-performance distributed systems, is still considerably less than the software delays that occur in many higher-level communications protocols [2].

Figure 6 illustrates the stability of the DRAMA protocol to handle high-load conditions without choking. It compares the traffic placement policies described in Section 4 and demonstrates that all the policies are effective when traffic overloads occur on the network; however, *backoff* and *tempered backoff* work somewhat better at high loads. Up to about 75% of capacity, practically all traffic on the net is message traffic with a negligible (<1%) amount of noise. At about a 90% load the message traffic levels off and stays constant, independent of the amount of traffic offered. The last data point we measured is with an offered load of 185% of capacity. The amount of noise, i.e. collision, is about 5% and the additional capacity wasted is between 5% and 15% depending upon the placement policy.

Table 1 presents the data shown in Figure 6 in order to compare the traffic placement policies. Up to 80% load we see that all policies provide approximately the same throughput. Above that point, the policies with *backoff* show a slight improvement over *speedup* and *fixed*. Delays are a little better for the *backoff* policies starting about 60%, and show more improvement as

load increases. Thus, policies with *backoff* should be used for traffic placement in the DRAMA protocol.

An interesting - and not totally unexpected - property is revealed in Figure 7. As we scale up from 10 bands with 100 nodes to 50 and 500, respectively, the performance improves. The increase in available bands allows further separation between synchronous and asynchronous traffic so that the latter has more bands to choose among. At the upper node and band count, little deterioration in access delay occurs up to 80% offered load. Most other protocols fail to reach this level of performance and most get worse as the node count increases.

One important question discussed in Section 3 is how DRAMA handles a combination of localized, bursty overload traffic. Figure 8 shows how quickly the system reacts under the various traffic placement algorithms in order to clear the nodes' queues. The total number of packets in the system is shown at various simulation times for three situations: (1) with 80% offered traffic for *tempered backoff*; (2) with 90% offered traffic for *tempered backoff*; and (3) at 2 seconds in the simulation of an 80% run we offer an overload to 10% of the nodes. The overload consists of as many packets as would have been needed to raise the total offering to 90% over the entire simulation period. The network response to this severe transient is good in that packet queues return to previous levels within 6 to 10 seconds.

The DRAMA protocol shows a marked difference from normal CSMA/CD behavior in the effect of round trip delay (potential collision slot time). Here, performance is not the same for different values of transmission time and round trip delay even though the ratio of the two is the same. Figure 9 shows that only at higher values (approximately 20 km length) does the round trip delay decisively increase the access delay. Two factors in DRAMA that do not exist in single-band CSMA/CD help to reduce the effect of round trip delay and hence improve the LAN spanning distance. First, with the reduced bandwidth the effect of a collision is considerably lessened, since the collision slot time is a much smaller percentage of the packet transmission time. Second, the probability of collisions is further reduced since, with multiple bands, a collision occurs only if two packets arrive within the slot and select the same empty band. Thus, besides extreme flexibility in configuring a network as noted above, DRAMA provides the additional feature that LANs can span both much greater and widely differing distances.

In our experiments, we have seen a major impact on performance by virtue of DRAMA's frame structure. In these runs the frame structure has been enforced so

	7.5%	40%	60%	67.5%	80%	100%
<i>Speedup</i>	20	30	135	315	1749	1167949
<i>Backoff</i>	20	28	123	284	1909	713018
<i>Fixed</i>	20	28	127	311	2277	1668377
<i>Tempered Backoff</i>	20	28	126	305	2050	897411

Table 1: Average Delay versus Load

DATA DELAY
IN MICROSECONDS

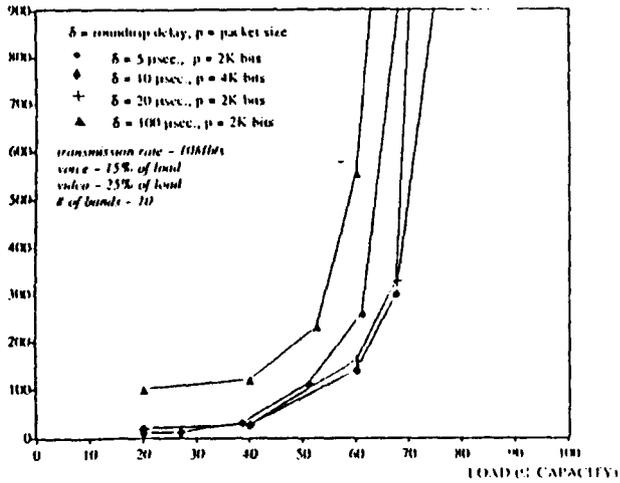


FIGURE 9: Data transmission delay versus load with varied normalized to 2K packet size and round trip delay

DATA DELAY
IN MICROSECONDS/
PACKET SIZE IN KILOBITS

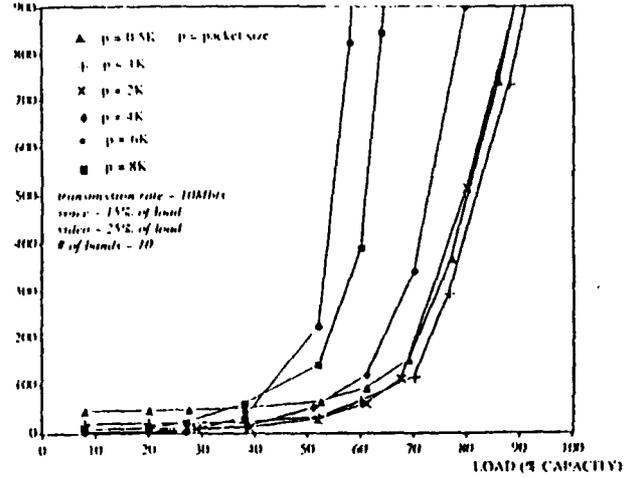


FIGURE 11: Data transmission delay versus load with normalized packet size

DATA DELAY
IN MICROSECONDS

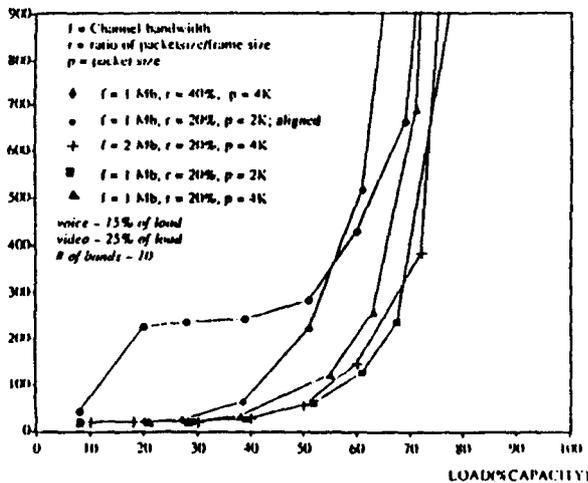


FIGURE 10: Data transmission delay versus load with varied frame length and transmission rate

DATA DELAY
IN MICROSECONDS

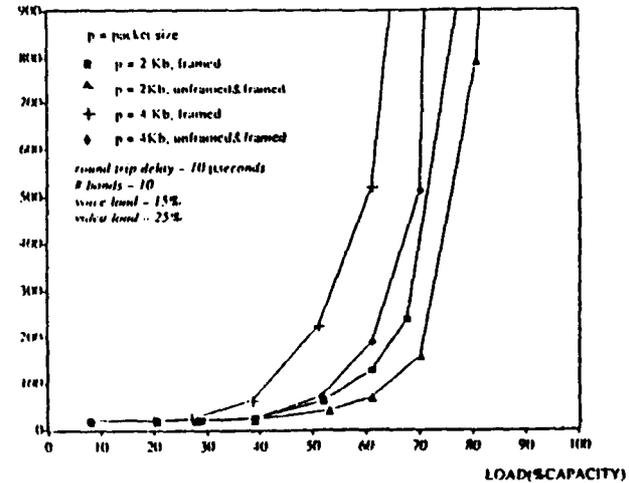


FIGURE 12: Data bands without packet framing

that packets which are too large to fit at the end of a frame are denied access until after the frame mark has occurred. This can cause additional delays especially as load increases. Figures 10-12 examine the relation between frame size, packet size, transmission rate, and frame begin alignment. In order for packets to fit at the end of a frame they should be small in comparison with the frame length. Figure 10 shows that if we keep packet to frame ratios to 20% then the delay performance is relatively stable under changes in band structure or frame size. However, if we increase this ratio to 40%, delay at higher loads is noticeably increased. Further, if we force all packets to be aligned on frame boundaries performance even at low loads is much poorer. Figure 11 illustrates the problem by change in packet length from .5K to 8K bits. In this figure, the delay has been normalized to a 2K packet delay since one is required to send more packets at lower packet size in order to transmit the same information. Again, we see that for ratios 20% or less the performance is relatively constant, but above that delay increases appreciably. Note that at 6K and 8K packet sizes it is possible to fit only one packet into a frame.

Finally, Figure 12 shows that performance can be increased significantly by changing the protocol algorithm slightly. In this modification each band leader elides the frame begin marker whenever no voice/video calls are assigned to the band. Essentially, this makes the protocol unframed CSMA/CD for data until voice/video calls are placed on the band. Dynamic framing with overflow and timing recovery on those bands which have synchronous traffic assigned, similar to the FDDI overflow system [1,5], would provide further recovery of wasted capacity.

These studies also addressed a concern left from our earlier study [11]. That study focused on reallocation of bands among LANGs and assumed that the allocation could be "easily" done; this will not be true if synchronous traffic tends to spread across most bands allocated to a LANG. This study shows that, using the DRAMA protocol, data remains reasonably well concentrated on a small number of bands. Even under high loads (e.g. 100%), in each experiment we ran, some channels never carried any synchronous data throughout the run. This allows the LANG to release some bands if network conditions require reallocation.

Fairness is another metric for a local area network. In the DRAMA system there are two considerations of fairness: (1) that across the whole network all LANG have equal access delay, and (2) that within a LANG each node have equal access delay. The first fairness factor, across the network was reported in [11] and showed equal LANG access. The second fairness fac-

tor was measured in this simulation study for the individual nodes. Statistics were computed to observe mean delay and standard deviation of delay about the mean for each node in the network. In 90% of all cases observed, the means tended to be equal for all nodes and the standard deviation of the delay was always less than the mean. In the remaining 10% of the cases, the standard deviation never exceeded 125% of the mean. We note that beyond 100% load these statistics become meaningless since the queues grow unbounded. This indicates that nodes are being fairly treated and that no node is being denied access excessively due to some quirk in the protocol or node location.

7 Conclusions

The analysis of the DRAMA protocol reported here and in [11] indicates that the protocol has several important features. The most impressive characteristic is its extreme flexibility. A single metropolitan area network, using the DRAMA protocol can support:

- LANGs with very different numbers of nodes,
- LANGs spread across a wide geographic area,
- Dramatic fluctuations of load, and
- Widely varying mixtures of traffic types.

The protocol provides this flexibility since:

- A network can quickly rebalance loads by reallocating bands among the LANGs (taking on the order of 30 to 300 ms),
- The protocol effectively integrates voice, video and data on a CSMA/CD network,
- The traffic placement policies work well with dynamic resource allocation since the number of bands with synchronous traffic is kept to a minimum, and
- The network is stable even at very heavy loads and with momentary overloads at some nodes.

These studies also suggest that, even with a limited bandwidth (say 10 Mb/s), use of a multiband network rather than a single band (say ten 1 Mb/s bands versus one 10 Mb/s band) can significantly lower average access delay (though transmission time for large packets on a 1 Mb band will be longer). In addition the multiband approach allows integration of synchronous and asynchronous traffic.

As a result of these studies we have additional encouraging performance information on the DRAMA protocol; we know that both *backoff* and *tempered backoff* provide satisfactory traffic placement.

An area that needs further research with the DRAMA protocol is the problem of providing tunable transmitter hardware for each node so that a node is able to transmit at various bands and enough receivers so that the node can receive voice/video and data packets sent to it in a random fashion. Further study is underway to assess the feasibility of hardware to accomplish the reception function without having rapidly tunable receivers to detect when information is arriving for a particular node. We are also studying buffering mechanisms to handle the problem of a single node receiving a large influx of information on several bands simultaneously.

Certainly the DRAMA protocol provides one alternative for supporting the multitransic-type flexible gigabit networks which are becoming progressively more important. We are now in the process of developing a simulation testbed with the capability of generating comparative performance data for protocols such as DRAMA, FDDI-II, QPSX, CSMA/RN [7], and tree-structured MAN protocols.

References

- [1] FDDI token ring media access control (MAC). *ANSI Standards Document ASC X3T9.5 Rev. 10*, February 28, 1986.
- [2] T. Mueller, B. Bhargava, and J. Riedl. Experimental analysis of layered Ethernet software. *Proceedings of the ACM-IEEE Computer Society 1987 Fall Joint Computer Conference*, 559-568, October 1987.
- [3] P.T. Brady. A technique for investigating on-off patterns of speech. *Bell System Technical Journal*, 44, January 1965.
- [4] Imrich Chlamtac and Aura Ganz. A multi-bus train communication (AMTRAC) architecture for high-speed fiber optic networks. *IEEE Transactions on Communications*, 6(6):903-912, July 1988.
- [5] Doug Dykeman and Werner Bux. Analysis and tuning of the fddi media access control protocol. *IEEE Transactions on Communications*, 6(6):997-1010, July 1988.
- [6] G. Ennis and P. Filice. Overview of a broad-band local area network protocol architecture. *IEEE Journal on Selected Areas in Communications*, SAC-1, No.5:832-841, November 1983.
- [7] E.C. Foudriat, D. Game, L. Founds, K.J. Maly, R. Mulkamala, and C.M. Overstreet. A carrier sensed multiple access ring for gigabit networks. *Technical Report TR-89-9, Dept. of Computer Science, Old Dominion University*, February 1989.
- [8] T.B. Fowler and S.F. Holmgren. A wideband cable bus local area network. *COMPCON*, 405-414, Fall 1982.
- [9] M. Hatamian and E.G. Bowen. Homenet: a broadband voice/data/video network on a CATV system. *AT&T Technical Journal*, 64:347-367, February 1985.
- [10] G. Exley, L. Merakos, and C. Bisdikian. Interconnection of CSMA local area networks: the frequency division approach. *IEEE Transactions on Communications*, 730-738, July 1987.
- [11] Kurt Maly, C. Michael Overstreet, Xia-Ping Qiu, and Deqing Tang. Dynamic resource allocation in a metropolitan area network. *Proceedings, SIGCOMM Symposium*, 13-24, August 1988.
- [12] M.A. Marsan and D. Roffinella. Multichannel local area network protocols. *IEEE J-SAIC*, 885-897, Nov. 1983.
- [13] A.N. Netravali and Z.L. Budrikis. A broadband local area network. *Bell System Technical Journal*, 64, no. 10:2449-2465, December 1985.
- [14] M.W. Rahn. Cablenet: a user perspective. *COMPCON*, 337-342, Fall 1982.
- [15] S. Sharrock, K. Maly, S. Ghanta, and H. Du. A broadband integrated voice/data/video network of multiple LANs with dynamic bandwidth partitioning. *Proceedings, INFOCOM '87*, 417-425, March 1987.
- [16] S. Sharrock, K. Maly, S. Ghanta, and H. Du. A framed, movable-boundary protocol for integrated voice/data in a LAN. *Proceedings, SIGCOM '86*, 9 Pages, August 1986.
- [17] P.C. Wong and T.S. Yum. Design and analysis of a contention-based lookahead reservation protocol on a multichannel local area network. *IEEE Transactions on Communications*, 36, no.2:234-238, Feb. 1988.