A HYBRID MEDIA ACCESS PROTOCOL FOR FIBER OPTIC RING NETWORKS

A. Bhargava, J. Kurose, D. Towsley

COINS Technical Report 87-116

October 29, 1987

A HYBRID MEDIA ACCESS PROTOCOL FOR FIBER OPTIC RING NETWORKS

A. Bhargava, J. Kurose, D. Towsley

COINS Technical Report 87-116

October 29, 1987

A Hybrid Media Access Protocol for Fiber Optic Ring Networks[1]

AMIT BHARGAVA[2], JAMES F. KUROSE[3], DON TOWSLEY[3]

University of Massachusetts
Amherst, Mass. 01003

Abstract

A new hybrid protocol is proposed for high bandwidth rings in which the round trip propagation delay is much larger than the packet transmission time. Features of random access protocols and conflict-free protocols such as token passing are combined to achieve superior performance. The scheme permits simultaneous use of the channel by many packets, is fair to all stations, and is completely distributed. The performance results presented in this paper indicate that the system remains stable for throughputs up to a maximum of 1, and that the delay characteristics are better than those of other related access protocols. The protocol additionally provides for reservation of bandwidth on demand and bounded delays for real-time applications.

[1]This work was supported by the National Science Foundation grant ECS 83-10771 and the Office of Naval Research grant N00014-87-K-0304

[2]Department of Electrical and Computer Engineering. This author is now with Codex Corporation, Canton, MA 02021.

[3]Department of Computer and Information Science

1 Introduction

New applications which require the integrated transmission of video, graphics, voice and data are placing increasing demands on local area network (LAN) communication facilities. While optical fiber provides a transmission medium which can satisfy the raw bandwidth requirements of these applications, the effectiveness of any fiber-optic LAN configuration will be strongly influenced by the choice of *topology* and *access protocol.*

In the case of optical fiber, the unidirectional nature of the medium and the nature of the coupling taps, as well as the need to ensure fairness among all stations, argue in favor of the *ring topology*. Access mechanisms that have been proposed for fiber-optic LANs fall into one of two broad classes: the *random access* CSMA-type protocols and the *controlled access* token-passing-type protocols. A critical network parameter influencing the performance of such protocols is the round trip propagation delay of a ring, α (measured in units of packet transmission time). In the case of a high speed network, α can be quite large. For example, for a 5 km ring with a transmission rate of 1 Gbit/sec and 1000 bit packets, α is approximately 20, as compared to a value of 0.1 in a conventional low speed ring. It is well known that the performance of CSMA-type protocols exhibits a severe degradation in throughput with rising values of α [4]. On the other hand, when α becomes large, token passing schemes have very poor delay characteristics at low loads. Therefore, new protocols are needed for the high speed case, and *hybrid* protocols, which combine features of CSMA and token passing are a possible solution.

In this paper we propose a hybrid protocol which exploits the fact that multiple simultaneous transmissions are possible when α becomes large. At low loads, a station can transmit as soon as it senses the channel idle and several packet *trains* [8,9,7] can simultaneously propagate around the ring. As the load rises, these multiple trains increase in size and number and may collide with each other, thus destroying packets. In such cases, the hybrid protocol switches to a more restricted mode of operation in which a single packet train propagates around the ring. As we shall see, the asynchronous hybrid access scheme described in this paper is fully distributed, does not suffer from problems such as lost/duplicated tokens and endlessly circulating packets, and achieves superior performance over other related access schemes for a wide range of network parameter values.

In the following section, a description of the hybrid protocol is presented. The performance evaluation of the scheme and representative numerical results are presented in section 3. Section 4 describes how the basic protocol can be used to bound delays for real-time applications. In section 5 the protocol is compared to other related access schemes and section 6 contains the conclusions of the paper.

2 Description of the Protocol

This section describes the operation of the hybrid access protocol in detail. The main characteristics of the protocol are as follows:

1. The topology used is a unidirectional ring.

2. *Active* taps are used to couple stations to the ring. In general, the taps could be either passive or active. Passive taps require the removal or insertion of a signal into the cable through optical couplers, whereas active taps remove the entire signal from the incoming link and retransmit it on the outgoing link. Difficulties in using passive taps [3] and falling costs of interface electronics indicate that *active taps* are the better option [1]. We thus assume through the remainder of this paper that active taps are used, with the receive tap placed upstream of the transmit tap. We also require that the distance between the taps (in units of number of bits) be at least the length of the packet header (described in detail below).

3. If a station is transmitting a message and detects another message passing its upstream (receive) tap, it aborts its own transmission in deference to the upstream transmission, *except* in special circumstances (described in the following sections).

4. There are two modes of operation: *Multiple Train Mode (MTM),* when the load is *estimated* to be low, and many trains of packets can propagate around the ring simultaneously, and *Single Train Mode (STM),* at higher estimated loads when only *one* train of packets is allowed to form.

5. The various fields in the header of each packet are shown in figure 1. The FLAG field contains a standard bit pattern which is used to delineate the packet. The SOURCE and DEST fields identify the transmitting and intended receiving stations, respectively. A receiving station sets the COPIED bit on each passing packet for which it is the intended destination and copies the body of the packet from the ring. As we shall see, the setting of the COPIED bit *alone* does not imply that the packet has been received completely at the DEST station, since partially transmitted (incomplete) packets may be transmitted by a station as a result of the access protocol.

   The MODE field is 2 bits wide and can take the following values:

   MTM – if the source of the packet transmitted it in the Multiple Train Mode
   STM – if the source transmitted it in the Single Train Mode
   STM:1st – if the source is attempting to switch modes from MTM to STM.

   As discussed in the following section, the PRIORITY field is used to ensure the eventual emergence of a single train when the mode is changed from MTM to STM.

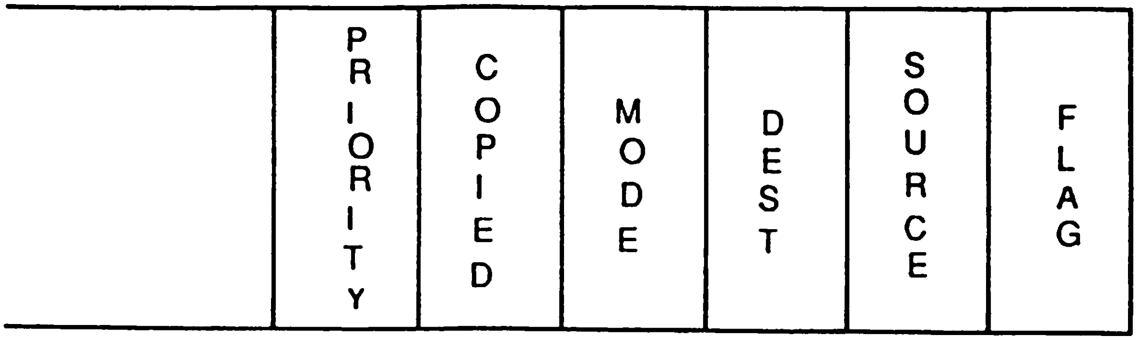

Figure 1: The packet header.

2.1 Space-time diagram representation

The operation of the protocol can be best understood using the space-time diagrams [19,20]. In this representation, the ring is cut at an arbitrary reference point and transformed into a line segment of length α. A packet transmission sweeps out a diagonal strip of area in the space-time diagram as it propagates around the ring and until it is removed from the channel. In figure 2, for example, a station located at x_1 starts transmitting at time t_1, and sweeps out an area of width 1; note that the packet 'wraps around' the space time diagram, as indicated in figure 2. The transmission which begins at (x_3,t_2), however, does not complete, as the station at x_3 defers to the upstream transmission by the station located at x_2.

Use of the space-time diagrams permits the spatial distribution of the stations to be considered when describing the state of the channel. *Areas,* rather than periods of time, now indicate whether the channel is busy or idle. This is a particularly valuable descriptive approach for the protocol under study, since multiple non-interfering message transmissions can exist simultaneously on the channel in a high speed network.

2.2 Protocol Operation

As previously indicated, the protocol has two modes of operation. We now describe these modes of operation in detail and discuss the mechanisms for changing from one mode to another.

A. Operation in the Multiple Train Mode (MTM)

In MTM, a station senses the channel when it first receives a packet to transmit from the higher layer protocol. If it finds the channel idle, it starts transmitting the packet. If the channel is sensed busy, the station waits until the channel is sensed idle again, and then *immediately* begins message transmission. That is, the station senses the channel to determine whether a train of packets is

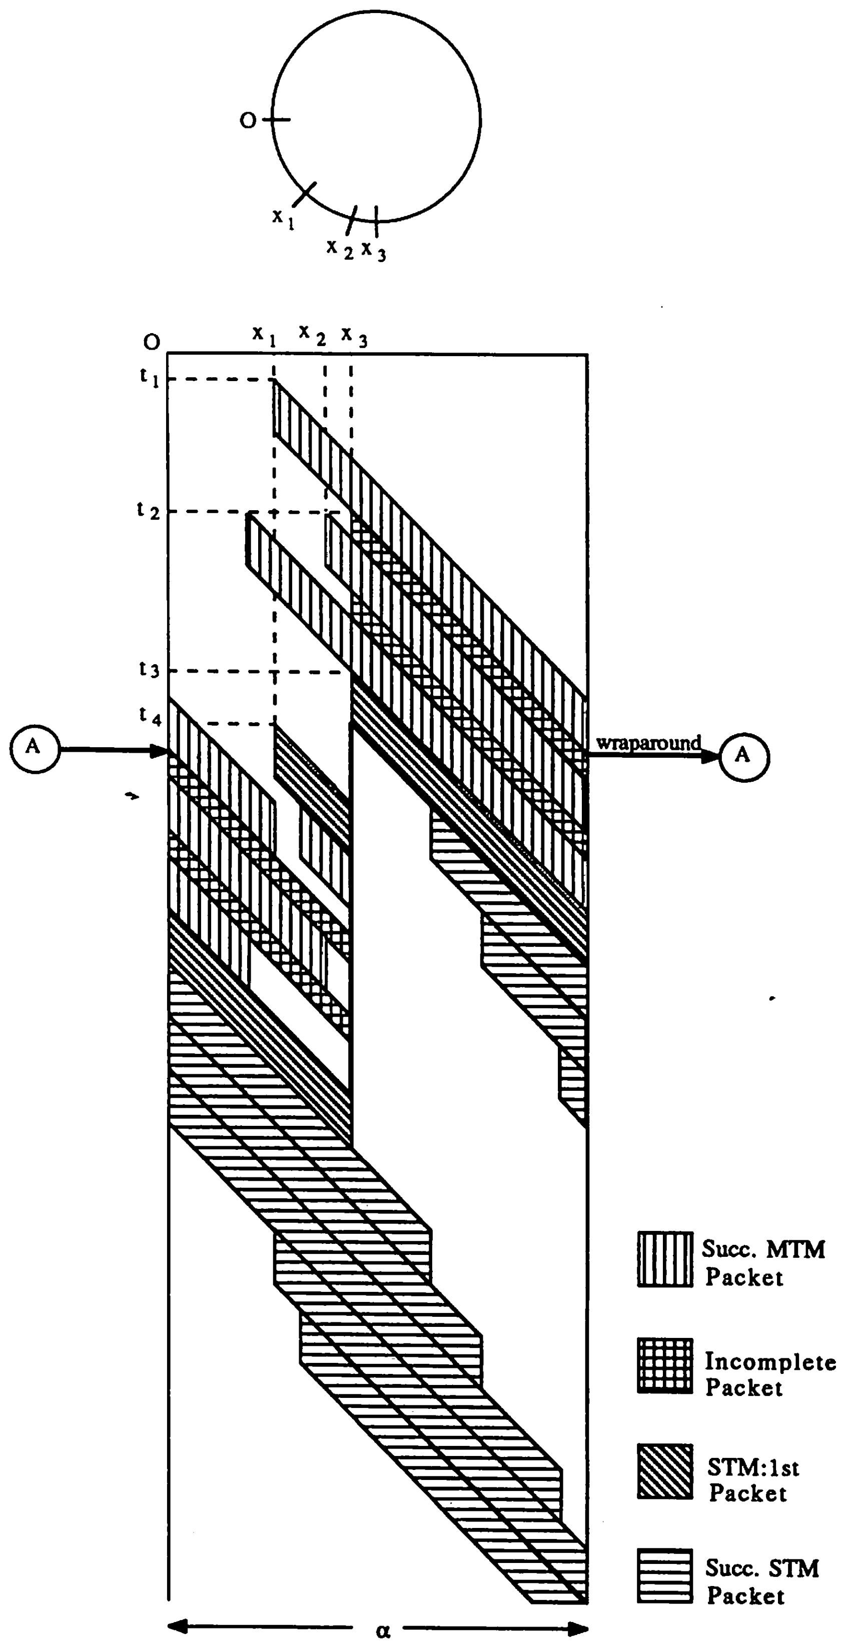

Figure 2: Operation in MTM and Switch to STM

going by - if so, it adds its packet to the end of the train [8,9,7]; otherwise, it starts a new train.

If a station senses the channel busy at its receive tap (which is upstream with respect to the transmit tap) while it is transmitting in this mode, the station aborts the current transmission, and the upstream packet is allowed to pass through. Note that as a result of this deferral process, a packet which begins transmission successfully may not be transmitted in its entirety. At the receiving station, the COPIED bit in the packet header is set as soon as the receiver detects that it is the intended destination for the passing packet, regardless of whether or not the packet is received in its entirety. As a result of the above considerations, a packet is considered successfully transmitted only when:

1. It is transmitted completely.

2. It is received back completely at the source node with the COPIED bit set.

After a packet (complete or incomplete) has propagated around the ring, it must be removed from the ring by the transmitting station. This can be done if the SOURCE field on the packet header is not destroyed. As we will see, if the header is destroyed, the incomplete packet will be removed from the ring when the mode switchover occurs. Thus, such packets will not circulate around the ring indefinitely.

Finally, we note that it is not possible to operate the system in the Multiple Train Mode alone, as this mode is unstable without a control mechanism such as STM. Figure 3 shows an example of this unstable behavior for a ring with α=4. In this example, stations located at $x_1, x_2, \ldots, x_8$ transmit at the same time t_1. Due to the attempt and defer mechanism x_2 defers in favor of x_1, x_3 defers in favor of x_2, and so on until x_1 itself defers in favor of x_8. Therefore *none* of these transmissions succeeds and since the mode of operation is MTM, *each* station retransmits at time t_2, and the exact same pattern is then repeated. This event is but one of several events which can cause this behavior. Since these events occur with a non-zero probability, it is clear that MTM alone is inherently unstable.

B. Operation in the Single Train Mode (STM)

When a station is in STM and has a packet to transmit, it does so when the end-of-train signal (EOT) is sensed. The end-of-train is recognized when the channel is sensed idle *immediately after* it was sensed busy. Since the receive tap is *upstream* of the transmit tap, a station can add its packet to the end of the train [8,9,7], as shown in figure 2. We note that this scheme is essentially the same as conventional token passing protocols, with the token now being represented by the end-of-train.

It is important to note that since packet removal is performed by the source and since the train moves in one direction, packet addition is always at the end of the train and packet removal (by

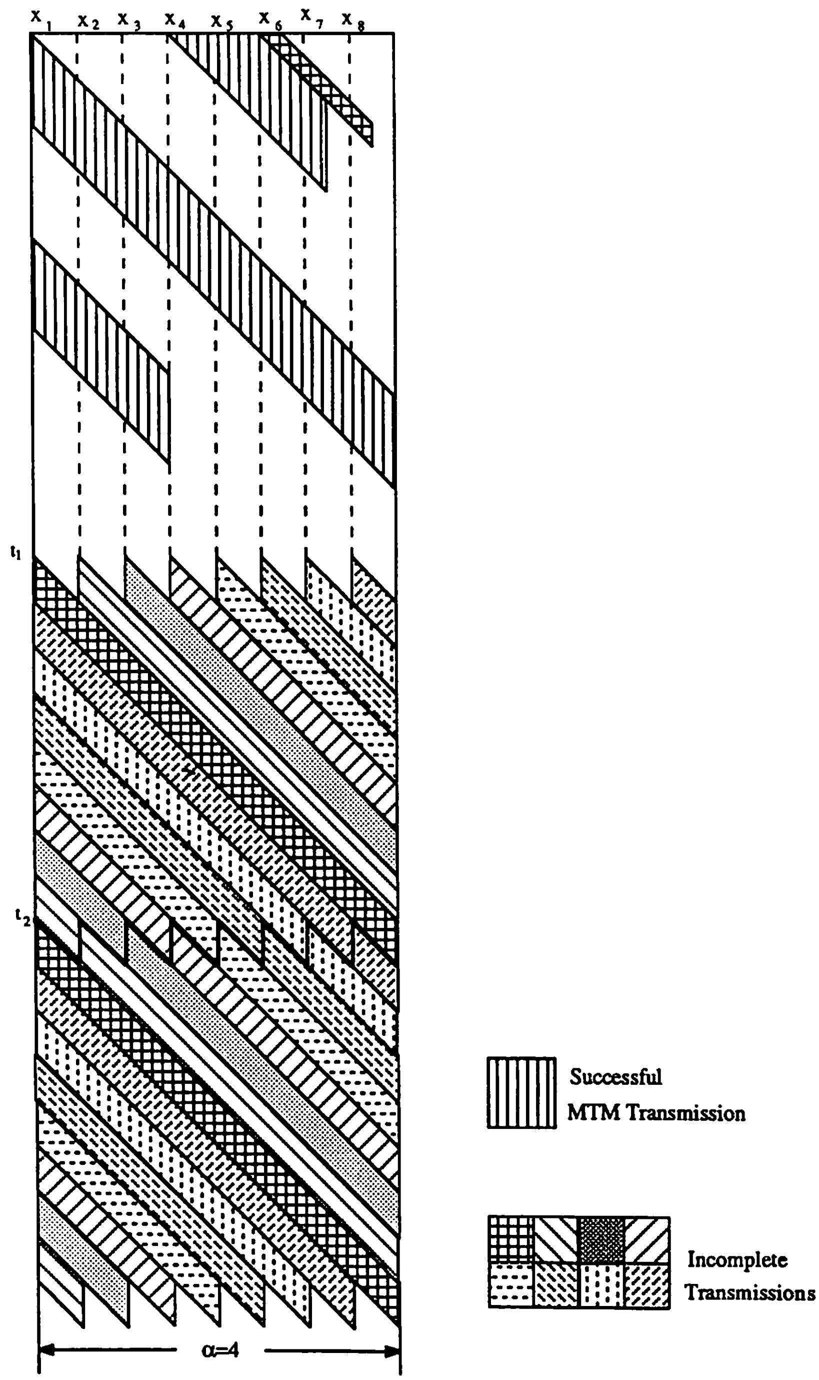

Figure 3: Event causing unstable behavior in MTM

the original transmitting station) is always from the head of the train. Therefore, unless explicitly changed, once the single train has formed, it will *remain* a single train and access to the channel will be collision-free.

C. Switchover from the Single Train Mode to the Multiple Train Mode

A station operating in STM can switch to MTM in the following ways:

1. Each station estimates the load over a window of time (L) by evaluating the fraction of time during which the channel was busy. If the estimated load falls below a threshold (k_2), it changes from STM to MTM. A station can affect this change either by setting the MODE bit to MTM on a packet that it transmits, or by changing the mode field on the next packet that passes through its tap.

2. On receiving a packet with the MODE bit set to MTM, a station switches to MTM if it has not already done so.

3. If no activity is sensed for α units of time, a station changes mode to MTM. This is required since it is possible that no station has a packet to send and yet the entire ring is in Single Train Mode, in which case the train will cease to exist (α time units after the last transmission).

Figure 4 shows the system operating in STM initially. At time t_1, the station at x_1 switches mode by setting MODE=MTM on its packet. Each subsequent station which receives this packet then switches to MTM mode.

D. Switchover from the Multiple Train Mode to the Single Train Mode

This switchover can occur in either of two ways:

1. While in MTM, each station again estimates the load on the channel by calculating the fraction of time for which the channel has been busy over some window of time of length L. If the estimated load exceeds some threshold (k_1), and the station has a packet to send which has been partially transmitted (i.e., a packet whose transmission was previously interrupted in deference to an arriving upstream packet) *at least once*, the station transmits the packet with the MODE field set to STM:1st.

   While a STM:1st packet is on the ring, the station that transmitted that packet removes any upstream packet that it receives, except in the case that the received packet is also a STM:1st packet and its PRIORITY is higher. This ensures that there is a single train on the ring and that it is *unique*. On receiving a STM:1st packet with higher priority, the station switches to STM as described in 2 below, otherwise it switches to STM on receiving back its own STM:1st back. Any simple scheme which insures that each station has a unique priority (e.g., using a unique node ID) can be used to determine priority.

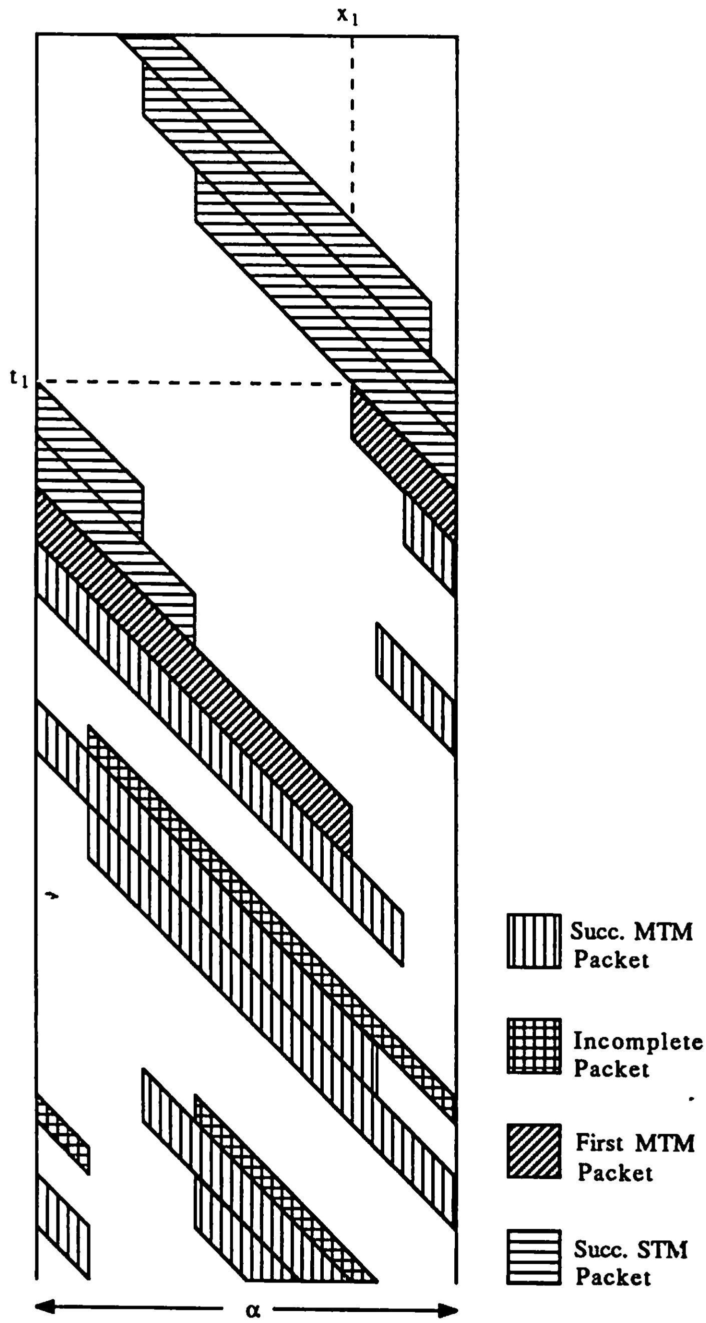

Figure 4: Operation in STM and Switch to MTM

2. α time units after the last packet with the mode bits set to STM:1st (i.e. starter of a single train) is received, a station switches to STM. During this interval (of length α) the station *can* transmit packets (with MODE=STM) but cannot initiate a change of mode back to MTM (which it can do while operating in STM).

In figure 2, the system is initially operating in MTM. At time t_3, the station at x_3 decides to switch modes and thus transmits a STM:1st packet. The station then removes all packets that are received until its original STM:1st packet is received back (after time α). Note that the STM:1st packet which starts at (x_1,t_4) is thus removed, since in this example we assume that the priority of the station at x_4 is lower than that of the station at x_3.

3 Performance Analysis

In this section performance models are developed to provide an estimate of the average access delay (i.e., the average amount of time between the arrival of a packet and the start of its successful, and complete, transmission) as a function of throughput (the average number of packets transmitted successfully per unit time).

In our models, message arrivals to the network are assumed to constitute a two dimensional Poisson process described by:

$$Pr[k \ arrivals \ in \ area \ A] = \frac{(\lambda A/\alpha)^k}{k!} e^{-(\lambda A/\alpha)}$$

This is equivalent to assuming that the network-wide message inter-arrival times are exponentially distributed with parameter λ and that the location of the station at which an arrival occurs is uniformly distributed between 0 and α.

We first consider the Single Train Mode in isolation; an analytic model is developed and verified by simulation. This mode of operation is shown to provide good performance at high loads and offer round robin service to active stations. The goal of developing this model for STM in isolation is to establish a baseline for the purposes of comparison with the hybrid protocol. We then study the performance of the hybrid protocol under combined STM/MTM through simulation.

3.1 Analysis of STM

Figure 5 shows the system operating in STM. The *cycle time* (t_c in the figure 5) is defined to be the time between two consecutive end-of-trains, sensed at a reference station. $\overline{N}$ is defined to be the average number of packets transmitted during an average cycle of length $\overline{t_c}$.

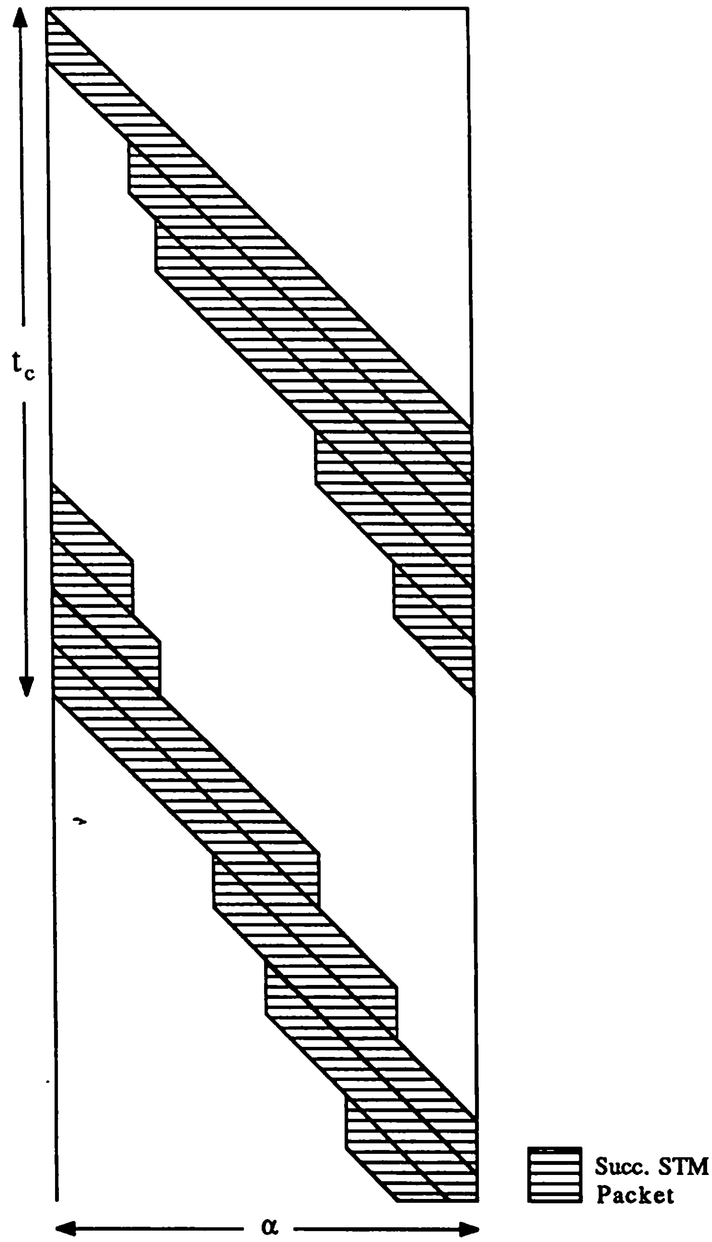

Figure 5: Cycle in STM

With reference to figure 5, note that the time $\overline{t_c}$ consists of a fixed interval of length α and the time required to transmit $\overline{N}$ packets. Therefore,

$$\overline{t_c} = \alpha + \overline{N} \times 1.$$

By flow conservation we have:

$$\overline{N} = \lambda \overline{t_c}.$$

and thus:

$$\overline{t_c} = \alpha/(1 - \lambda).$$

The operation in STM is similar to a polling system in which the polling message is represented by the end-of-train. We may thus use the analytic results for polling systems developed by Konheim and Meister [22]. For constant packet lengths and a fixed number of stations (M) with infinite buffers, the expression obtained in [22] is,

$$E[Delay] = \frac{\overline{t_c}}{2}(1 - \lambda_{stn}) + \frac{M\lambda_{stn}}{2(1 - M\lambda_{stn})}$$

where λ_{stn} is the arrival rate at a *single* station. In order to obtain an expression for the case of an infinite number of uniformly distributed stations with single buffers, we take the limit of the above expression as $M \to \infty$, $\lambda_{stn} \to 0$ such that $M\lambda_{stn} \to \lambda$. This gives:

$$E[Delay] = \frac{\alpha + \lambda}{2(1 - \lambda)}$$

The results of the analysis are verified by simulation and are plotted for α=10 in figure 6. Note that the expression for the mean delay indicates that stable operating throughputs arbitrarily close to 1 can be achieved.

3.2 Simulation Results for the Hybrid Protocol

The added complications of mode switchovers make the hybrid protocol too complex to study analytically and we therefore study its performance through simulation. We assume an infinite population of users with single buffers, and constant packet lengths in these simulations. Results are obtained for different values of the following parameters:

1. k_1: the estimated load threshold for switching from MTM to STM
2. k_2: the estimated load threshold for switching from STM to MTM
3. L: the window over which the load is estimated, measured in units of the round trip propagation delay α.

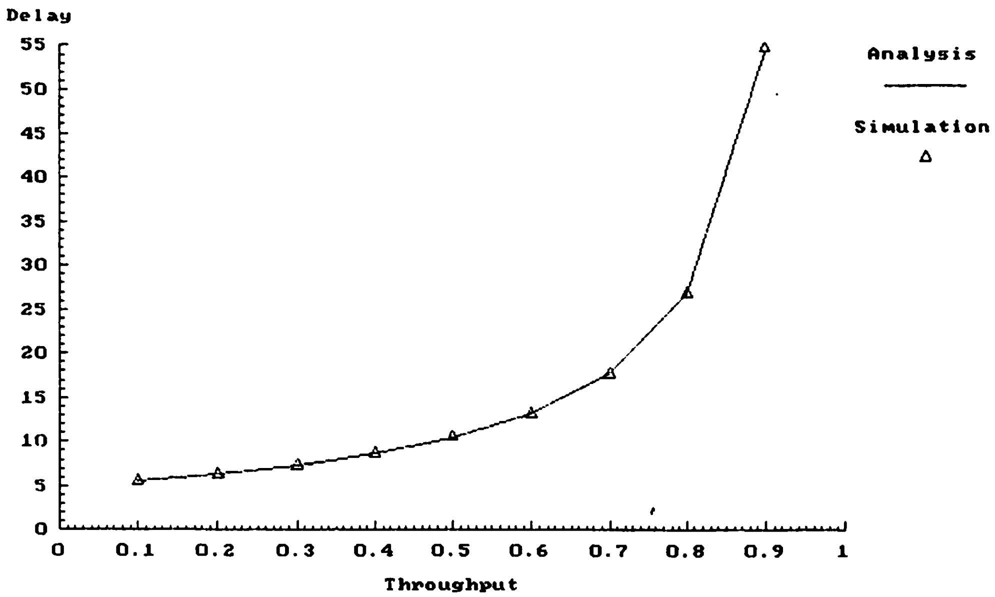

Figure 6: Comparison of Analysis and Simulation in STM, α=10

α is taken to be 40 for these studies, based on a bit rate of 2 Gbits/sec., 1000 bit packets and a ring of length 5 km.

Figure 7 shows the delay-throughput curves for different values of k_1 and k_2 with L fixed at 2. The results indicate that it is desirable to switch modes to STM only for high values of estimated load. This suggests that the operation should be MTM until a significant amount of congestion occurs. A change to STM should then be made to clear the backlog. Figure 8 shows similar curves for L=1, and though the rise in delays occurs much sooner, the general conclusions are the same: k_1 should still be a relatively large value. Clearly, the choice of k_1 and k_2 is critical, so much so that the performance of the hybrid protocol can be *worse* than an exclusively STM protocol with poorly chosen threshold values.

In figure 9, we study the effects of varying the value of k_2, i.e., effectively keeping the system in STM for different amounts of time after the mode switchover has occurred. Note that the difference between the performance curves is not very significant, suggesting that the estimated load falls rapidly as backlogged packets are transmitted after switching to STM.

Finally, in figure 10, the window size (L) is varied, keeping the values of k_1 and k_2 constant. Note that for small values of L, the performance improves as L increases. However, increasing L beyond a certain limit results in worsening performance, as can be seen for the case of L=16. This is to be expected since as $L \to \infty$, the mode of operation will be exclusively STM or MTM, each

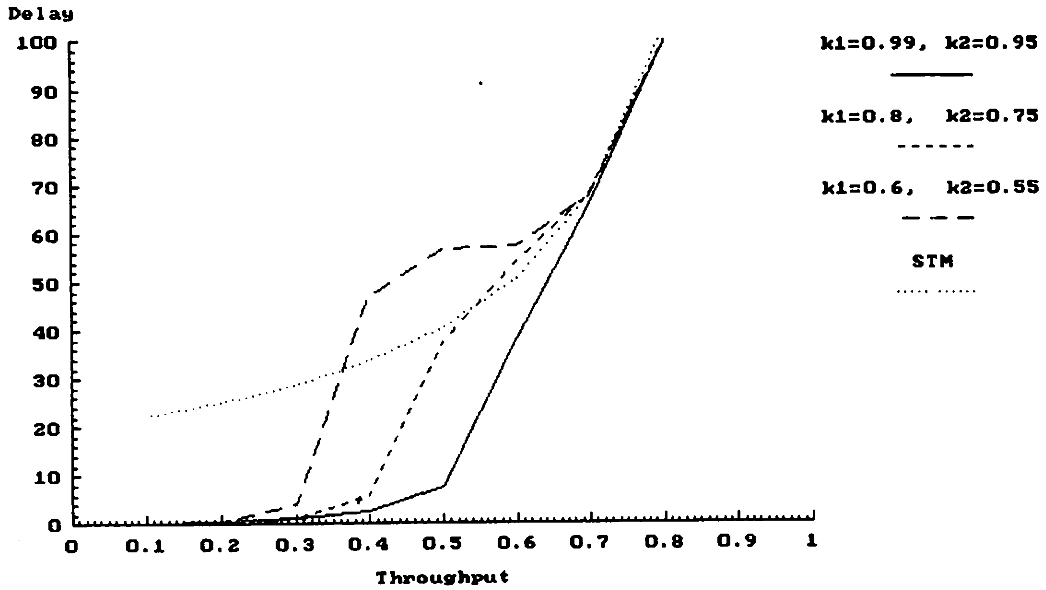

Figure 7: Delays vs. Throughput for α=40, L=2

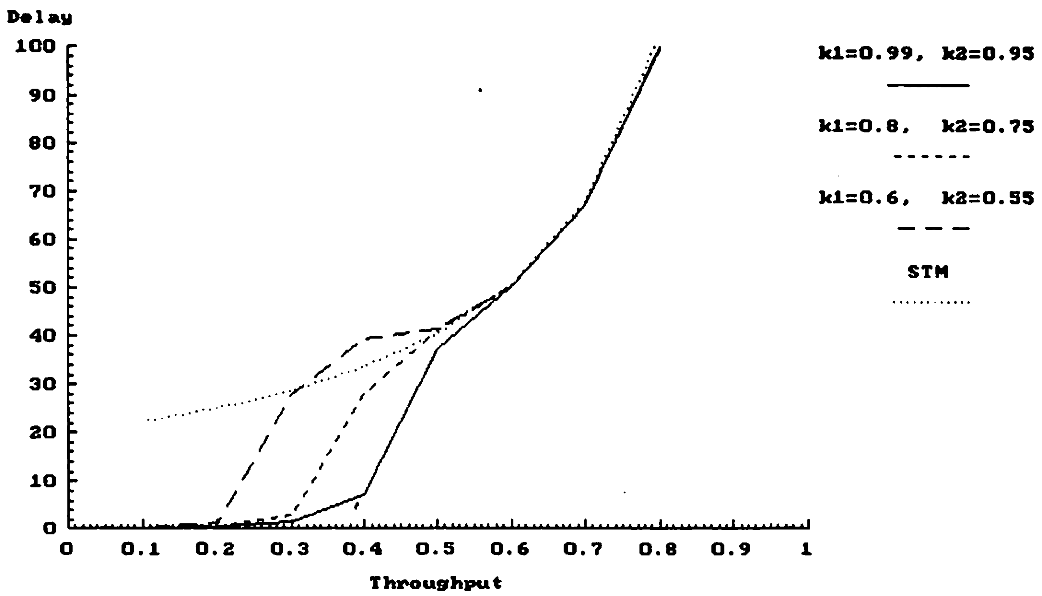

Figure 8: Delays vs. Throughput for α=40, L=1

of which have performance inferior to that of the hybrid protocol.

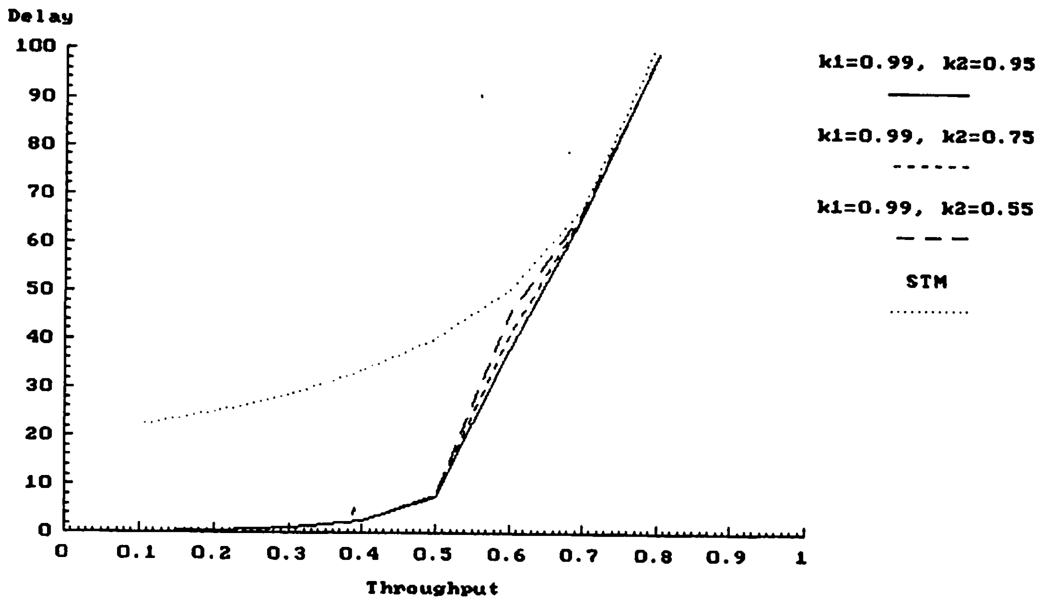

Figure 9: Delays vs. Throughput for α=40, L=2, k_1=0.99

4 Bounding Delays for Real-Time Traffic

With the hybrid protocol, a station can bound delays for real-time traffic and can 'reserve bandwidth' by transmitting a packet once every α time units after it first transmits a packet successfully. Figure 11 illustrates this mechanism. At time t_1, the station at x_1 transmits a packet which returns at time $t_2 = t_1 + \alpha$. While removing this packet off the ring at its receive tap, the station can simultaneously transmit another packet on its transmit tap. Since no other station can interfere with its transmissions, a station can continue to do this indefinitely. This approach thus provides a simple mechanism for providing a station with the reserved bandwidth required for circuit-switched types of applications. Note that it is also possible for a source/destination pair to remove each other's packet from the ring while transmitting its own, thus effectively establishing an efficient two-way connection. In figure 11, stations at x_2 and x_3 operate in this manner.

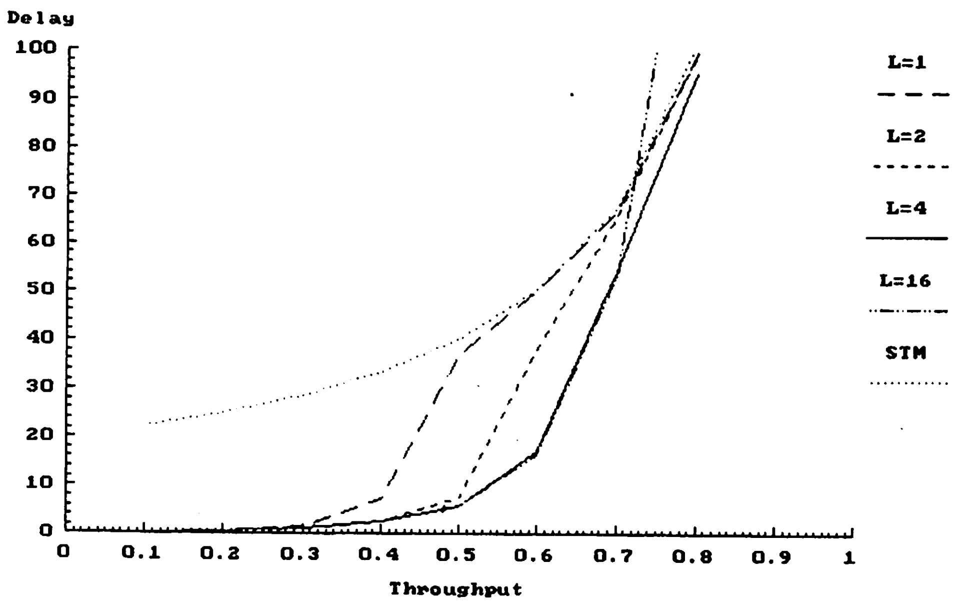

Figure 10: Delays vs. Throughput for α=40, k_1=0.99, k_2=0.95

5 Comparison with Other Schemes

5.1 General Considerations

Numerous other schemes have been proposed for fiber-optic-based LANS. Topologies and representative networks based on these topologies include:

1. U shaped - U-Net [5], D-Net [6] and others [7].
2. S shaped - Expressnet [8].
3. Spiral - Expressnet [8].
4. 2 Unidirectional Buses - Fasnet [9], Buzznet [10].
5. Unidirectional Ring - Contention Ring [11], [12].

Note that with the exception of unidirectional ring, each of these topologies introduces an implicit ordering among stations. This may result in unfair access by the various network stations and different stations may thus exhibit different throughput-delay characteristics. The unidirectional ring topology clearly does not suffer from this problem.

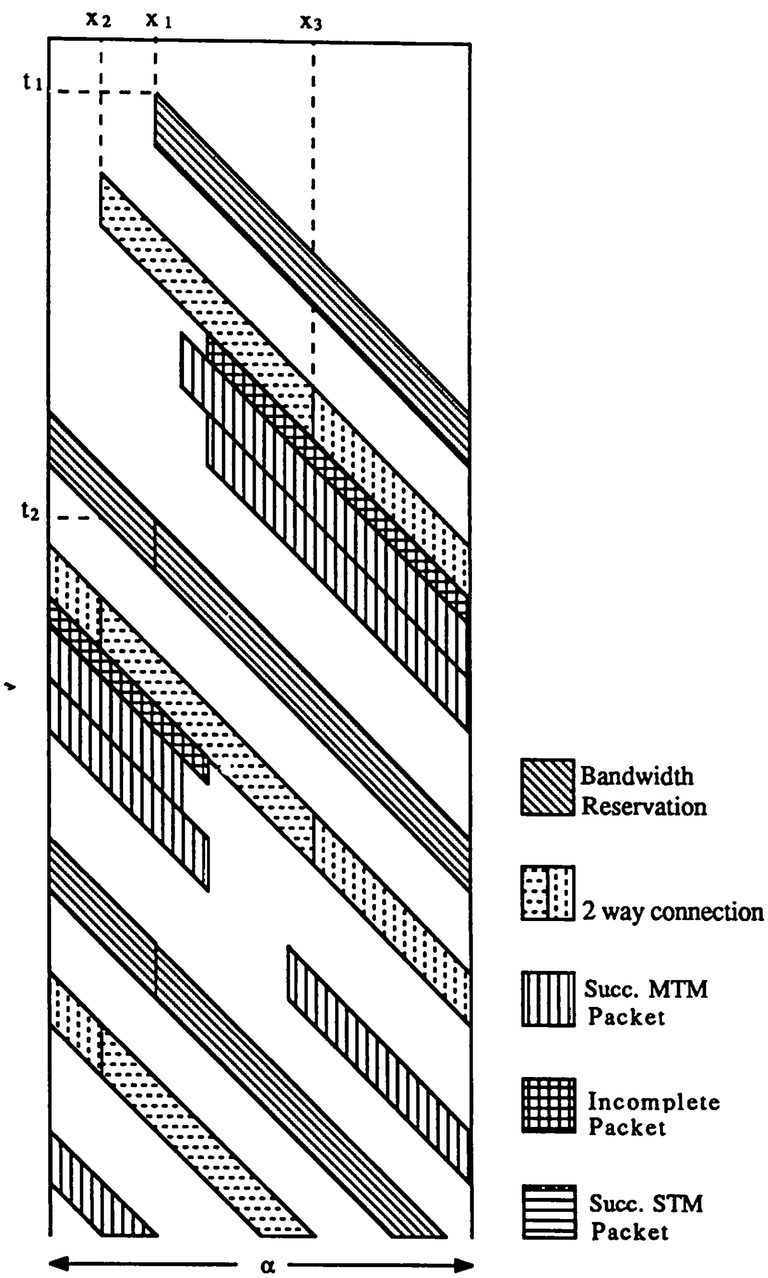

Figure 11: Bounding Delays for Real Time Traffic

The access mechanism employed on most of these topologies falls in the category of Demand Assignment Multiple Access (DAMA) schemes [13]. Prominent among these are:

1. The scheduling delay access mechanism (SDA) – after a successful transmission, access rights are given to users on the basis of a unique index number. Examples are Silentnet [14], L-Expressnet [15].
2. Reservation access mechanism (RA) – a control wire (or channel) is used to place a reservation for stations to transmit on the main channel. An example is UBS-RR [16].
3. Attempt and defer mechanism (AD) – a station transmits if the channel is idle and defers to upstream packets in the case of a collision. This mechanism is used by Expressnet [8], D-Net [6], Fasnet [9], U-Net [5], Buzznet [10].

Clearly, SDA has the disadvantages of high latency at low loads and possible unfairness. Moreover, high throughput can be achieved only for a large number of active stations. The RA mechanism wastes bandwidth since two channels are required. We believe the AD mechanism to be the most suitable for unidirectional channels and the hybrid protocol studied in this paper belongs in this general category.

Hybrid protocols attempt to combine the best of contention-based and collision-free access schemes to achieve better performance. Examples of such schemes are HYMAP[27,17], the Contention Ring [11,12] and Buzznet[10,25]. HYMAP permits CSMA/CD at low loads and switches to token passing at high loads, but requires a bi-directional bus topology (unsuitable for optical fiber). It also requires the presence of a synchronizing station to switch modes and is therefore not fully distributed. Its performance also degrades as α increases – as expected, since CSMA/CD has a poor performance for large propagation delays and also because the mode switchover in this scheme takes much longer.

Finally, we note that two other ring protocols will achieve performance similar to that of our proposed protocol, but have disadvantages associated with their implementation. The Slotted Ring [23] protocol allows many stations to access the channel simultaneously, though it requires strict synchronization to be maintained between all stations (recall that our hybrid protocol is asynchronous). The FDDI specification [24] recommends the use of explicit token passing to determine channel access. In this case, tracking of lost and endlessly circulating (busy) multiple tokens is an overhead which must therefore be considered. In addition, FDDI supports circuit-switching by allotting FDM channels to such applications using a *centralized* scheme. Our hybrid protocol, however, requires no token maintenance and uses a more robust distributed scheme to allot bandwidth to circuit-switched traffic.

5.2 Comparison of Delay Performance

In this subsection we compare the delay versus throughput performance of our hybrid protocol with other access schemes. It is difficult to make a uniform comparison among all protocols since the previous analyses of these protocols have been made under differing assumptions about the number of stations, number of buffers at each node, packet size, etc. Therefore, we make several different comparisons with various sets of protocols – each under the assumptions introduced by their designers.

Figure 12 compares the protocol with Buzznet [25,10], the Contention Ring (C-Ring) [11] and the Token Ring. These curves assume an infinite number of stations with single buffers, and *exponential* packet lengths. The propagation delay, α, is set at 5 in this example. The Contention Ring permits CSMA on a ring, but unlike the AD mechanism, *both* packets involved in a collision are destroyed. For values of α less than 1, the switchover to the implicit token passing mode is automatic in this protocol. However, switchover does not occur when α is greater than 1 and consequently the throughput is severely restricted. A possible solution to this problem is to use frequency division multiplexing (or wave division multiplexing in the case of optical fiber) to obtain several subchannels, each with a smaller bandwidth and hence a smaller value of α. Note that as α decreases, the number of collisions and retransmissions also decreases but the packet transmission time (measured relative to the transmission time on the undivided channel) *increases*, causing larger delays due to a smaller effective service rate. Therefore, it is reasonable to divide the main channel into sub-channels only up to some limiting (optimal) number. As shown in figure 12, however, the performance of our proposed hybrid protocol is still better than the multi-channel C-Ring, even assuming an *optimal* number of channels in the C-ring protocol. This difference is explained by noting that the C-Ring uses a non-persistent scheme for rescheduling incomplete transmissions, whereas the proposed protocol uses a 1-persistent scheme with a control mechanism.

Figure 12 also shows the delay characteristics of Buzznet[1], which shows poorer performance than the hybrid protocol over a large range of throughput values. (Buzznet has a dual bus topology and for purposes of comparison, α was taken to be the *end-to-end* propagation delay, i.e. the time taken for the signal to propagate from the most upstream to the most downstream user). These differences result from the fact that in Buzznet, a station transmits in the random access mode *only* if the buses are sensed idle when a packet arrives; otherwise it transmits in the controlled access mode. Our proposed hybrid protocol, however, uses the controlled access mode only when the estimated load is high, thus exhibiting a better relative performance.

Figure 12 also indicates that the Token Ring suffers from relatively higher delays at low throughput values (the hybrid operates mainly in MTM in this range), while at higher throughput values the performance of the Token Ring is essentially the same as the hybrid protocol (which operates mainly in the implicit token passing, single train mode for high throughput values). We additionally

[1]Based on Figure 7 in [25]

note that our empirical studies have shown that with exponential packet lengths, the thresholds k_1 and k_2 should generally be chosen to be smaller than in the case of constant packet lengths. In the case of constant packet lengths, a newly transmitted packet may fit exactly into an open "slot" created by a recently removed packet, i.e., a natural "pipelining" of messages can occur. In the case of non-constant message lengths, however, this pipelining does not occur. We conjecture that as a result, the onset of congestion occurs at lower throughput values in the case of non-constant packet lengths and hence an earlier mode switchover to STM is required.

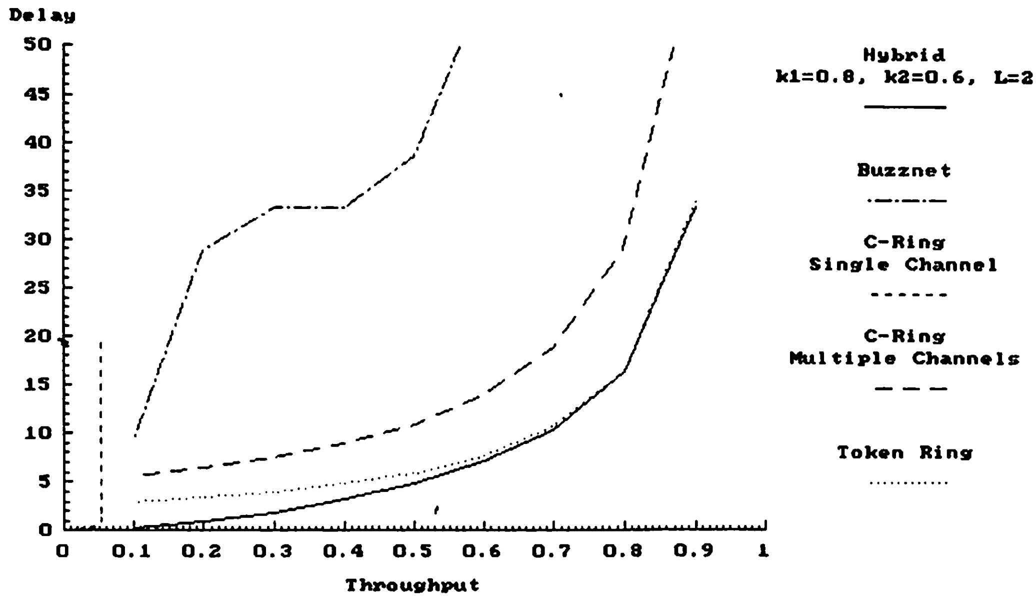

Figure 12: Comparison with Buzznet, C-Ring, Token Ring

Figure 13 compares the performance of the proposed hybrid protocol with Fasnet [26] (using the Most Upstream First Service (MUFS) and the Gated Sequential Service discipline (GSS) disciplines) [2] and Expressnet [26] (using the Nongated Sequential Service (NGSS) discipline). The number of stations (M) is taken to be 20 and 50. Figure 13 indicates that the hybrid protocol has a better delay performance in all cases except for MUFS when $M = 50$, in which case MUFS has lower delays over a small range of throughputs (though the hybrid provides much higher throughputs than the other protocols.) This behavior can be explained as follows: NGSS and GSS have round robin scheduling in which a station transmits only once a round, whereas MUFS allows the most upstream user to access the channel whenever it has a packet to send. As discussed in the following subsection, the hybrid protocol achieves its higher throughput due to a lower overhead per cycle.

[2] Based on Figure 9 in [26]

Finally, from an implementation standpoint, we note that Fasnet is a *bit synchronous* protocol which relies on one station to synchronize the entire network; it is therefore not fully distributed. Also, the MUFS service discipline is *unfair* in that all stations do not exhibit the same delay characteristics and in the case of large α and large M, the performance difference for the most upstream and most downstream station can become large.

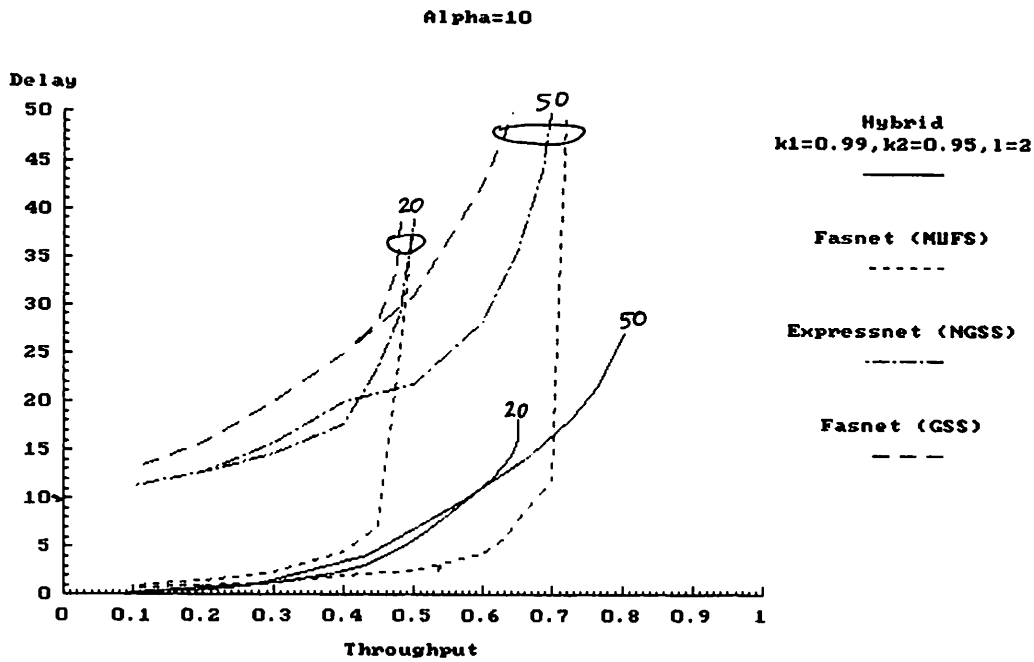

Figure 13: Comparison with Fasnet and Expressnet

5.3 Comparison of Throughput Performance

In this subsection we examine the maximum throughput that can be achieved by various protocols under the assumption of a finite number of stations. The case in which only one station is active is also examined in detail.

In order to determine the maximum throughput achievable by the proposed hybrid protocol, we assume that the M stations are distributed around the ring at equal distances from each other. We consider the case in which these M stations are saturated, i.e., each station always has a message to send. For $M \leq \alpha$, the operation is always in MTM and the maximum throughput that can be obtained is M/α. For $M > \alpha$, the operation is in STM and since each cycle (defined in section 3.1) consists of M transmissions and an overhead of α, the maximum throughput that can be achieved is $M/(M + \alpha)$. The following table summarizes the maximum achievable throughputs for various protocols [13,18] and indicates the superiority of the hybrid scheme:

Buzznet	$\frac{M}{M+6\alpha}$
Expressnet	$\frac{M}{M+2\alpha}$
Fasnet	$\frac{M}{M+2\alpha+1}$
Proposed protocol	$\frac{M}{M+\alpha}$

Finally, let us consider the case in which there is but one active station in the network; this approximates the situation in which traffic from a single station dominates the network. In the case that each station has a *single* buffer which is cleared only after the transmitted packet is received back at the station, the maximum throughput that can be achieved is $1/\alpha$. However, if there are *multiple* buffers at each station, and multiple outstanding messages are permitted (i.e., there may be multiple messages which have been transmitted by a single station but have not yet propagated around the ring back to the transmitting station), a maximum throughput of 1 can be achieved. Among the other protocols previously discussed, only Buzznet and the Contention Ring can attain this limiting throughput value.

6 Conclusions

In this paper, we have described a new hybrid protocol for fiber-optic-based ring LANs. The protocol is asynchronous, completely distributed, and fair, and can also be used for real-time applications in which bounded access delays are critical.

Both simulation and analytic performance models were developed to study the performance of the protocol. A comparison with other existing schemes for Fiber Optic LANs indicated that, for network parameters within the range of practical interest, the proposed protocol generally achieves superior performance.

References

[1] M.R. Finley Jr., "Optical Fibers in Local Area Networks," *IEEE Communications Magazine*, Vol. 22, No. 8, Aug. 1984, pp. 22-34.

[2] H.F. Taylor, "Technology and Design Considerations for a Very-High-Speed Fiber-Optic Data Bus," *IEEE Journal on Selected Areas in Commun.*, Vol. SAC-1, No. 3, April 1983, pp. 500-507.

[3] N. Lee Rhodes, "Interaction of Network Design and Fiber Optic Component Design in Local Area Networks," *IEEE Journal on Selected Areas in Commun.*, Vol. SAC-1, No. 3, April 1983,

pp. 489-492.

[4] J.O. Limb, "Performance of Local Area Networks at High Speed," *IEEE Communications Magazine*, Vol. 22, No. 8, Aug. 1984, pp. 41-45.

[5] M. Gerla, P. Rodrigues, and C. Yeh, "U-Net: A unidirectional fiber bus network," *Proc. FOC-LAN Conf.*, Las Vegas, NV, Sept. 1984, pp. 295-299.

[6] C. Tseng and B. Chen, "D-Net: A new scheme for high data rate optical local area networks," *Proc. Globecom '82*, Miami, FL, Nov.-Dec. 1982, pp. 949-955.

[7] N.F. Maxemchuk, "Twelve Random Access Strategies for Fiber-Optic Networks," unpublished manuscript.

[8] F.A. Tobagi, F. Borgonovo, L. Fratta, "Expressnet: A High-Performance Integrated-Services Local Area Network," *IEEE Journal on Selected Areas in Commun.*, Vol. SAC-1, No. 5, Nov.1983, pp. 898-912.

[9] J.O. Limb and C. Flores, "Description of Fasnet – A unidirectional local area communications network," *Bell Syst. Tech. Journal*, vol. 61, no. 7, Sept. 1982, pp. 1413-1440.

[10] M. Gerla, G.S. Wang, "Performance Models of Buzz-Net, A Hybrid Fiber Optics LAN," *Infocom '87*, San Francisco, CA, Mar.-Apr. 1987, pp. 298-306.

[11] Z. Gao, K.S. Vastola, R.S. Gilbert, "The Contention Ring: A Performance Comparison," *Infocom '87*, San Francisco, CA, Mar.-Apr. 1987, pp. 492-500.

[12] H. Okada and S. Ohno, "Analysis of an Integrated Token-passing and CSMA/CD Ring Network in LAN Environment," *International Comm. Conf. '86*, pp. 615-619.

[13] M. Fine, F.A. Tobagi, "Demand Assignment Multiple Access Schemes in Broadcast Bus Local Area Networks," *IEEE Transactions on Computers*, vol. C-33, no. 12, Dec. 1984, pp. 1130-1159.

[14] E.O. Jensen, M. Tokoro, and L. Sha, "Bus allocation scheme for distributed real time systems," Carnegie-Mellon Univ., Pittsburgh, PA, Rep. Dec. 1980.

[15] F. Borgonovo, L. Fratta, F. Tarini, and P. Zini, "L-Express-net: A communication protocol for local area networks," *Infocom '83*, San Diego, CA, Apr. 1983.

[16] F.A. Tobagi and R. Rom, "Efficient round-robin and priority schemes for unidirectional broadcast systems," *Proc. IFIP 6.4 Zurich Workshop on Local Area Networks*, Zurich, Switzerland, Aug. 27-29, 1980.

[17] P. Nain, N.D. Georganas and W.J. Stewart, "Analysis of a Hybrid Multiple Access Protocol with Free Access of New Arrivals during Conflict Resolution," *Infocom '87*, San Francisco, CA, Mar.-Apr. 1987, pp. 399-408.

[18] D. Blass and S. Thanawastien, "A CSMA/CD-Token ring local network architecture with CSMA/CD-DAMA protocol," *Infocom '87*, San Francisco, CA, Mar.-Apr. 1987, pp. 76-84.

[19] K. Sohraby, "Carrier Sense Multiple Access on the Bus Topology," PhD Thesis, Univ. of Toronto, Canada, Oct. 1985.

[20] M.L. Molle, K. Sohraby, A.N. Venetsanopoulos, "Space-time models of asynchronous CSMA protocols for local area networks," *IEEE Journal on Selected Areaas in Communications*, vol. SAC-5, no. 6, July 1987.

[21] M. Schwartz, "Telecommunication Networks: Protocols, Modeling and Analysis," Addison-Wesley, 1986.

[22] A.G. Konheim and B. Meister, "Waiting lines and times in system with polling," *Journal of the ACM*, vol. 21, no. 3, July 1974, pp. 470-490.

[23] J. Pierce, "How far can data loops go ?" *IEEE Transactions on Communications*, vol. COM-20, June 1972, pp. 527-530.

[24] F.E. Ross, "FDDI – a Tutorial," *IEEE Communications Magazine*, vol. 24, no. 5, May 1986, pp. 10-17.

[25] M. Gerla, G. Wang and P. Rodrigues, "Buzznet: A hybrid token/random access LAN," *IEEE Journal on Selected Areas in Communications*, vol. SAC-5, no. 6, July 1987, pp. 977-988.

[26] F.A. Tobagi, M. Fine, "Performance of unidirectional broadcast local area networks: Expressnet and Fasnet," *IEEE Journal on Selected Areas in Communications*, vol. SAC-1, no. 5, Nov. 1983, pp 913-988.

[27] M. Rios and N.D. Georganas, "A hybrid multiple-access protocol for data and voice-packet over local area networks," *IEEE Transactions on Computers*, vol. C-34, no. 1, Jan. 1985, pp. 90-94.