Simulating a Multimedia FDDI Backbone Network

KISENCHAND RANAI†,* KOON SENG LIM,† AND ROBERT HUIJIE DENG‡

*Department of Information Systems and Computer Science, National University of Singapore, Kent Ridge, Singapore 0511, Republic of Singapore

†Institute of Systems Science, National University of Singapore, Heng Mui Keng Terrace, Kent Ridge, Singapore 0511, Republic of Singapore

‡Department of Electrical Engineering, National University of Singapore, Kent Ridge, Singapore 0511, Republic of Singapore

The Fiber Distributed Data Interface (FDDI) is an ANSI standard for high speed fiber optic ring networks. Its media access control (MAC) protocol operates on the basis of a timed-token ring, and allows the transmission of both time critical and normal prioritized data. Because of these characteristics, it is highly suitable for use in a multimedia environment where there are a mixture of synchronous and asynchronous data classes. This paper describes an implementation of the FDDI MAC protocol using the RESQ2 simulation language and studies the behaviour of the protocol when used as a large multimedia backbone network where stations represent gateway connections to peripheral LANs. It also investigates the behaviour of symmetric and asymmetric loading on such large networks and some potential problems that may be encountered.

Received May 1991, revised February 1994

1. INTRODUCTION

The Fiber Distributed Data Interface (FDDI) (ANSI, 1987, 1988) is an ANSI standard for high speed fibre optic ring networks. The standard essentially specifies a 100 Mbps 200 km fibre optic ring network for general purpose connection among computers and peripheral equipment.

Two types of transmission frames are possible in the ring. The first type, the synchronous frames, have a guaranteed bandwidth and are thus suitable for the transmission of time critical or delay sensitive data. The second type, the asynchronous frames, are allocated bandwidth only after synchronous frame transmissions and are therefore in general highly dependent on the offered load in the network. This makes asynchronous frames suitable for lower priority data transmissions where they may be preempted by higher priority transmissions.

Because of these two distinct data classes, the FDDI is highly suited for integrated multimedia environments where time critical data like voice and video compete with lower priority data-like file transfers. Another attractive aspect of the protocol is that it offers up to eight distinct priority levels for its asynchronous service. This feature allows the implementation of different sub-classes of traffic within the asynchronous service itself, each perhaps with a different grade of service. Such an arrangement is extremely useful in multimedia networks where user requirements may vary greatly over the various mediums.

The performance of the FDDI as an integrated network has been investigated by several authors. For instance in Frotini and Watson (1988), a simulation based investigation of mixed voice/data traffic on the FDDI was conducted. The authors considered a 100 km ring with up to 500 connected stations and investigated the effect of the target token rotation time (TTRT) on the token rotation time of the ring. In Schill and Zieher (1987), a sophisticated traffic generator module was used to simulate bursty arrivals and bimodally distributed packet lengths of data frames in a realistic FDDI network. The main aims of the paper were to study the function of the timed token protocol with respect to various ring parameters and to investigate the sensitivity of the priority mechanism to the target token rotation time and the token holding time thresholds of the various priorities.

This paper investigates the behavior of the FDDI used as a large multimedia backbone network (Figure 1) where the stations represent gateway connections to peripheral LANs. It also studies the behaviour of symmetric and asymmetric loading on such a network. The behaviour of the network is studied based on the definition of two benchmark load patterns which serve to reflect typical loading configurations of large backbone networks. A set of performance requirements for both patterns is then defined and a series of simulation runs executed to observe the effects of varying load on the various ring parameters.

The paper is organized as follows. Following this introduction, Section 2 presents a brief description of the FDDI media access control (MAC) timed token protocol. Section 3 gives a brief description of the RESQ2 simulation language that was used to construct the model and some details of the model's operation. In Section 4, two benchmark load patterns are introduced as a form of performance specification for the purpose of investigating

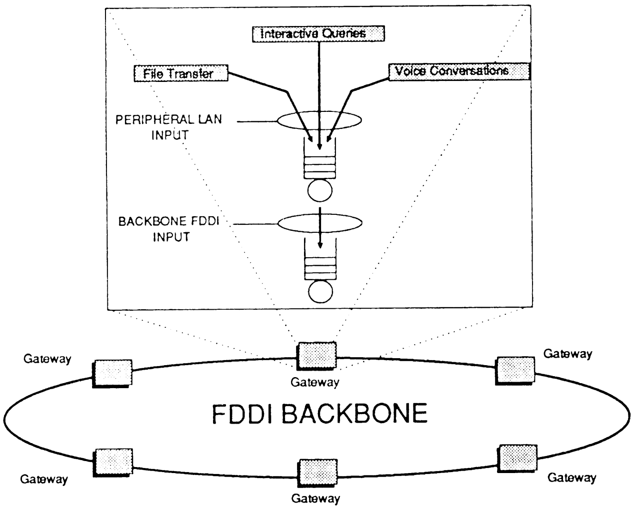

FIGURE 1. Logical model of a multimedia backbone network.

FDDI performance. Section 5 presents the results of the simulations using simulation model developed in Section 3 and the benchmark load patterns defined in Section 4. Finally, Section 6 concludes the paper with a discussion of the results and their significance.

2. THE FDDI TIMED TOKEN PROTOCOL

In this section, a brief overview of the FDDI MAC protocol operating in error-free mode is presented. The discussion will focus mainly on the network access issues of the protocol and other issues such as network initialization procedures, error recovery and other functions unrelated to frame transmission will be ignored. For a thorough overview of the FDDI MAC protocol, the reader is referred to Ross (1989).

The basic operation of the FDDI MAC protocol closely resembles a standard token ring (ANSI, 1985a) in which each station repeats data frames that it has received from its upstream neighbour to its downstream neighbour in a controlled manner determined by the reception of a token. With the inclusion of the priority mechanism, however, the protocol bears more resemblance to the token bus (ASNI, 1985b) MAC standard. Nevertheless, unlike the token bus which has only four classes of traffic (including class 0 which is equivalent to the synchronous class of the FDDI), the asynchronous traffic class of the FDDI has eight levels of priority.

All transmission in the ring is controlled by two timers and a set of preset threshold values: one for each asynchronous priority (denoted by T_Pri(i) for priority i). in addition, there is also a preset value, called the target token rotation time (T_Opr) which defines the ideal length of time between consecutive token arrivals at a station in a ring. The actual length of the token rotation (TRT), however, varies according to load and can be greater or less than T_Opr. The token rotation timer is used to keep track of this quantity and is read and compared with T_Opr with every arrival of the token at a station. If the TRT is longer than the T_Opr, the token is deemed to have arrived late and no asynchronous transmission can take place. Furthermore, the extent of this 'lateness' of the token is carried forward to the next token rotation so that transmission of asynchronous frames can only take place if the token is early by an equivalent amount of time.

Regardless of the value of token rotation timer, however, synchronous frames are always transmitted. The duration of transmission is pre-allocated by higher level protocols. Before transmission, however, the original value of the token rotation timer is first saved into another timer, the token holding timer (THT) and it is then restarted to time the next token rotation. Upon completion of synchronous transmission, the THT is activated and it continues to run upwards as asynchronous frames are being transmitted. A station may initiate a transmission of an asynchronous frame of a particular priority level i if its THT has not expired (run down to zero) and is greater than the threshold value of that priority level (T_Pri (i)). The order of transmission of packets belonging to priorities which satisfy the above criteria is on a FCFS basis. When the THT expires, the station must pass the token to the next station downstream. Because of the manner in which bandwidth is assigned, synchronous frames have a fixed bandwidth regardless of load while asynchronous frames have only the residual bandwidth left after transmission of synchronous frames. Moreover, the bandwidth of asynchronous frames is directly affected by the TRT of the ring, which is in turn dependent on the offered load. Figure 2 gives the flowchart of the FDDI MAC operations.

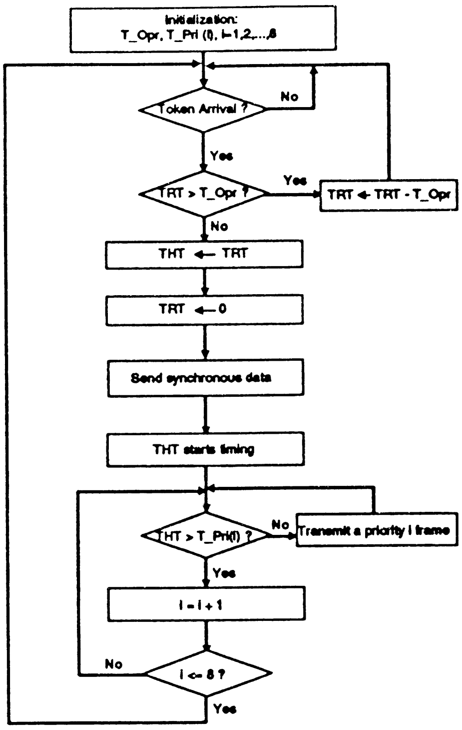

FIGURE 2. Flowchart of FDDI MAC.

As a consequence of the MAC protocol, two fundamental properties of the FDDI were observed:

1. The maximum TRT on the ring cannot exceed twice T_Opr.
2. The average TRT cannot exceed the T_Opr.

These two properties were proven in a study by Sevcik and Johnson (1987). The study also concluded that the protocol's restriction on the transmission of asynchronous traffic can actually be relaxed somewhat while still guaranteeing the above two properties.

Two potential problems were also identified in a recent study by Bux (1989). They are:

1. The relative performance of the priority levels is very sensitive to the number of active stations on the network.
2. The FDDI priority mechanism is unfair in the sense that stations with the same priority level, the same traffic characteristics but at different locations may experience very different frame transmission delay and throughput.

As noted in Bux (1989), the first problem worsens as the number of transmitting stations increases. In environments where the number of active stations on the network varies dynamically, it will be very difficult to determine the timing parameters to reliably meet given performance requirements. The second problem, also noted in Dykeman and Bux (1988) and Johnson (1988), arises when nodes transmitting low-priority frames are located immediately downstream from nodes transmitting higher priority frames. These stations have an unfair advantage over other low-priority nodes. A similar effect can be observed with low and high traffic stations, however, we postpone this discussion until Section 5.2.

3. Modelling of FDDI

This section describes a simulation model of a multimedia backbone FDDI network that was developed to study the FDDI MAC protocol. Because of the complexity of such a configuration, only packets emerging from the peripheral LANs entering the backbone are simulated. Intra LAN communication within peripheral LANs themselves are omitted because they do not contribute to the load on the backbone. A similar approach was taken by Martini *et al.* (1988). Two simulation models were written, one using IBM's Research Queueing 2 (RESQ2) simulation language and another using ANSI Pascal. The RESQ2 model ran on an IBM 3090 mainframe while the Pascal model ran on a VAX8600.

Both models were calibrated against each other and against Dykeman and Bux's (1988) model. Agreement of results among the two models were a primary concern and extensive testing was done to ensure that both models agreed and were accurate. In this paper, however, we only describe the construction of the RESQ2 model.

3.1. The RESQ2 simulation package

RESQ2 (Sauer *et al.*, 1986a,b) is a high level simulation language with powerful elements and functions for supporting the modelling of complex communication systems. Language support for controlling runs and evaluation of results include:

1. Probability distributions and statistical functions for load generation.
2. Active queues with various service disciplines.
3. Passive queues with access control for simulating limited resource allocations.
4. Complex routing functions.
5. Functions for obtaining accurate statistical results (including confidence levels and intervals).
6. Automatic gathering of queue characteristics (e.g. queueing time, queue lengths, etc.).

In RESQ2, a model is constructed by splicing together sub-networks of nodes and queues which have been previously defined. The notion of time in the model is effected by having the packets (or jobs) traverse the networks sequentially and queueing up at the appropriate queues to simulate the passage of time. The model itself can be further subdivided into a number of submodels, making modular design possible. In general, each submodel represents an independent component of the network which is joined to the main model through shared nodes.

The jobs that traverse the networks may be characterized by a job specific parameter array, usually used to hold information relevant to the object being modelled by the job (e.g. if the jobs are used to model packets, then the parameter array could hold values like packet arrival times, packet lengths, queueing time, etc.). In open networks, the jobs are created at a source according to some specified arrival distribution, and are routed through and delayed at the network according to the protocol being modelled. Jobs leave the network through sink nodes before being destroyed. Besides having open networks, closed networks are also possible where jobs circulate around the network indefinitely.

3.2. Model description

Figure 3 gives a simplified view of the RESQ2 model we have constructed. The model consists of an array of traffic generator submodels, a main RING submodel and a sink. The RING submodel (Figure 4) itself is further made up of an array of MAC submodels (Figure 5) with each submodel representing the MAC protocol of a single station on the simulated network.

The traffic generator submodels are used to generate jobs at the proper rate which represent voice and data packets circulating the network. The ring submodel, on the other hand, simulates token passing and physical token transmission while the individual MAC submodels handle the priority mechanism of the FDDI network internal to each station. In the model, there are two types of jobs. The first type is used to represent voice

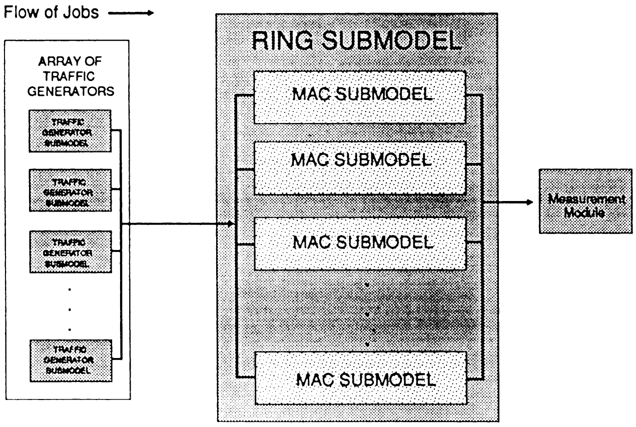

FIGURE 3. Block diagram of the RESQ2 simulation model.

and data packets, while the second type is used to represent the FDDI token.

Voice and data packets are created by the traffic generator submodel and enter the network after passing through nodes SET_BUFF, GET_MEAN and RING_IN.SET_BUFF is dummy node used to simulate the finite message queues within each station and will route packets out into the sink directly above if the queue capacity of the intended station is exceeded. GET_MEAN is used to count the number of packets that have entered the network and RING_IN merely serves as a common node between the RING submodel and the rest of the network.

The FDDI token, on the other hand, is created within the RING submodel and circulates indefinitely within. Once inside the RING submodel, jobs (be they packets or the FDDI token) will be routed to the appropriate MAC submodel by DECIDE1. Once they leave the MAC submodels representing a successful transmission, they will be either:

1. routed out of the RING submodel through DECIDE2 if they are packets; or
2. re-routed back through the network again if it is the FDDI token.

Packets that leave the RING submodel are eventually destroyed when the enter they sink. A FDDI token on the other hand is delayed at the TOK_TI queue, for an amount of time equivalent to passing the token to the next adjacent station. After leaving TOK_TI, it passes through nodes NEXT_STAT, which increments an internal counter to represent the arrival of the token at the next station, and SET_T which calculates the TRT.DECIDE3 is used to route the token either into a MAC submodel, if there are outstanding packets awaiting transmission within the stations, or back to TOK_TI if there are no outstanding packets, or if the token is late.

Within each MAC submodel (see Figure 5), packets and token are separated by M_IN depending on whether they represent:

1. synchronous data frames; or
2. asynchronous data frames; or
3. the FDDI token.

Jobs representing asynchronous frames are further routed by DEC_5 depending on their priorities into

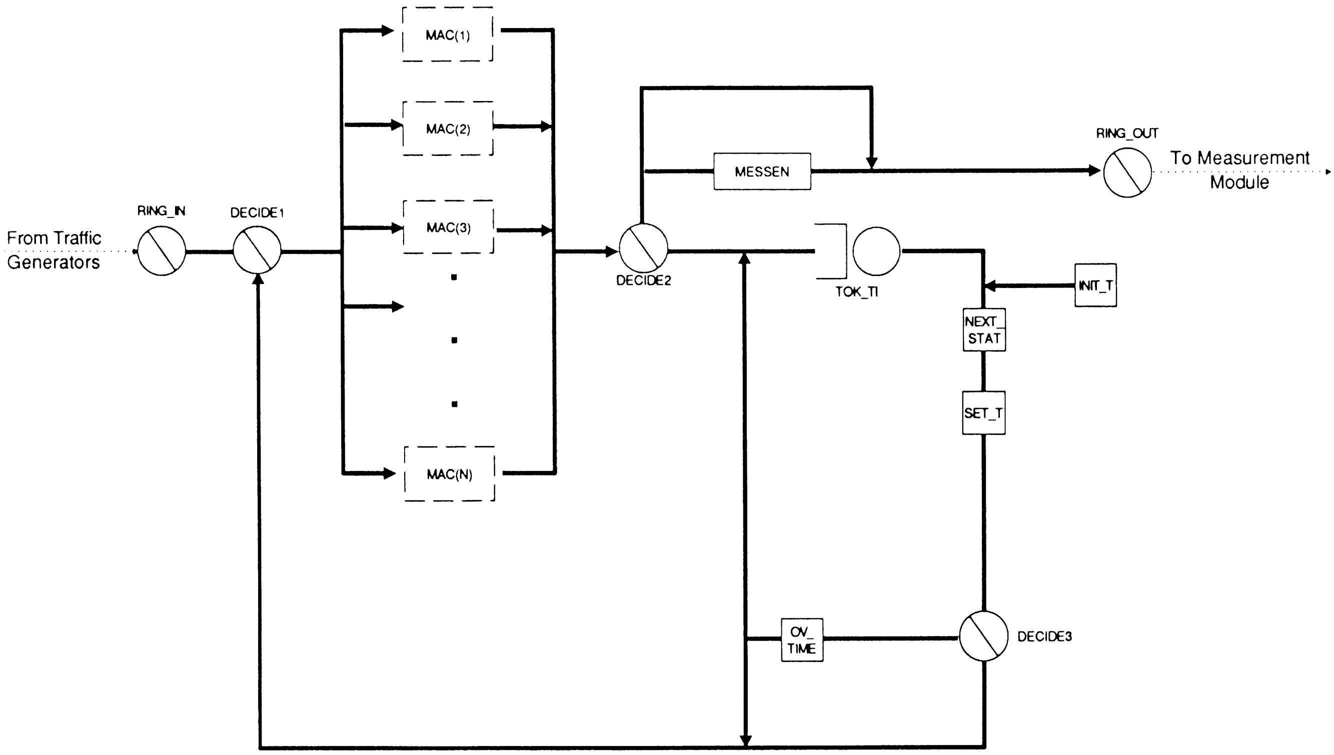

FIGURE 4. The ring submodel.

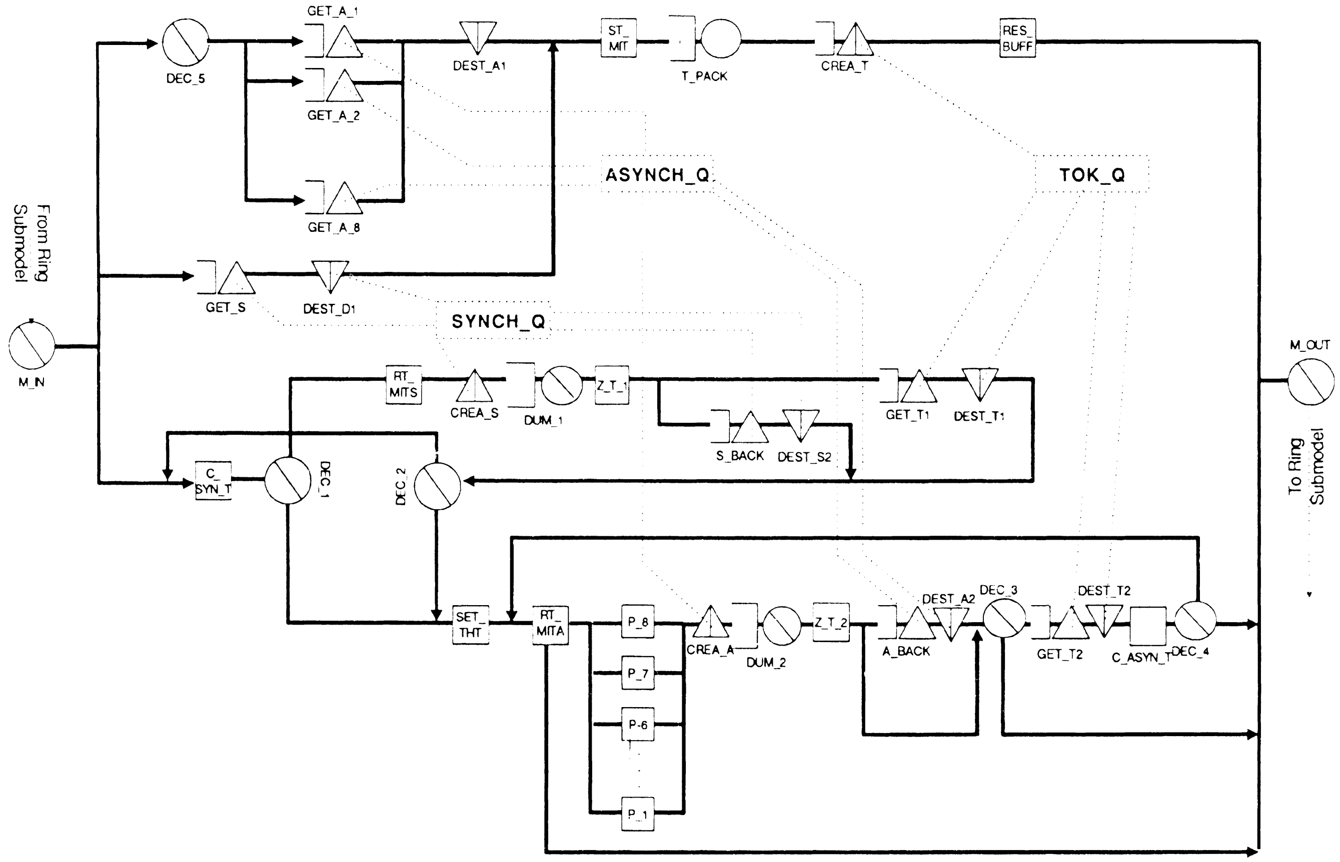

FIGURE 5. The MAC submodel.

eight individual passive queues, GET_A_1 to GET_A_8. These packets await at the queues for resource allocation from a 'pool of passive queue tokens' ASYNCH_Q. Jobs representing synchronous frames, await at GET_S for resource allocation from another 'pool of passive queue tokens' SYNCH_Q. The FDDI token, on the other hand, pass through C_SYN_T which determines whether or not there is sufficient time left for synchronous transmission and is routed at DEC_1 to either:

1. RT_MITS, if there is still sufficient time; or to
2. SET_THT, if there are not enough time left and asynchronous transmission must take place next.

At RT_MITS, an internal flag is cleared and the FDDI token allocates a 'passive queue token' from SYNCH_Q to GET_S at CREA_S, so that transmission of a synchronous frame may take place. Once this happens, the packet previously held up at GET_S is freed and it moves to ST_MIT where an internal flag is set to indicate a transmission. The packet then moves to T_PACK where it is delayed an amount of time equivalent to its transmission time. The FDDI token, in the meantime, gets routed to either:

1. GET_T1 if there was transmission of synchronous frames; or back to
2. DEC_2 through S_BACK, if there were no synchronous frames for transmission.

At GET_T1, the FDDI token is forced to wait for resource allocation from the 'pool of passive queue tokens' TOK_Q which will only happen when a corresponding synchronous packet reaches CREAT_T. This is used to synchronize the transmission of packets with the FDDI token.

The synchronous packet, after releasing a 'passive queue token' from TOK_Q moves out of the submodel through M_OUT. The FDDI token, meanwhile gets routed at DEC_2 back to either:

1. C_SYN_T, if there was a successful synchronous transmission; or
2. SET_THT, if there was no transmission.

This cycle repeats until there is either no more synchronous packets queued for transmission or the time allocated for synchronous transmission expires.

A similar sequence of events happens for asynchronous transmission except that the amount of 'passive queue tokens' allocated from the 'pool' ASYNCH_Q at CREA_A depends on the number of priorities that are allowed to transmit. Thus, suppose if the time only permits priority eight transmissions, then the amount of 'passive queue tokens' allocated at CREA_A will only be sufficient to allow one asynchronous packet queued at GET_8 to be freed.

3.3. Traffic generators

A relatively sophisticated traffic generator comprising further of three submodels was used in the simulation model to provide realistic traffic simulation. Specifically, the three component generators cater for the simulation of:

1. digitized voice packets with silence suppression;
2. bulk traffic sources originating from file transfers; and
3. interactive traffic sources originating from higher level protocols and interactive queries.

Because each station on the network actually represents a gateway connection, a station will have a number of instances of each of these component traffic generators.

3.3.1. Voice source generator

Each source generates a train of voice packets with the length of the train drawn from an exponential distribution (Yatsuzka, 1982; Gruber and Le, 1983) and a period of silence whose length is also given by an exponential distribution. This period of activity and subsequent silence corresponds to a talkspurt and silence period respectively. Figure 6 shows the arrival process of a single stream of voice generated from the model. In this model, voice packets (denoted by vertical markings) arrive at discrete time intervals during a talkspurt. There is a delay between consecutive packet arrivals because time is required to form a packet. This delay is known as the packetization delay and is denoted by T. A silence interval denoted by X, follows a talkspurt and its length is distributed exponentially with a fixed mean. The variable Z denotes the time interval between the arrival of the last packet of the current talkspurt and the first packet of the next talkspurt. For the actual values of X and T, the model used by Sriram and Whitt [21] was adopted which has characteristics given in Table 1.

3.3.2. Data source generator

For the data source generators (both bulk and interactive data), a simplified version of the model used in Martini *et al.* (1988) was adopted. In this model, bulk traffic is represented by a burst transmission of a stream of packets followed by a relatively long period of inactivity. The length of the burst depends on the file lengths, which is in general unknown. However, Martini *et al.* (1988) used an empirical distribution which they had observed in a program development environment (Welzel, 1987) using Ethernet. For this model, a constant file length of 20 packets is assumed.

Furthermore, using the well known 20/80/80/20 rule, it can be shown that the length of packets from the bulk source must be 16 times longer than the corresponding lengths from the interactive source. The characteristics of the bulk source is given in Table 2 with the values for inter-packet gap and inter-file gap calculated based on the mean data rate required and packet lengths.

The interactive source is assumed to have an exponentially distributed arrival rate with a constant packet length. Its characteristics are given in Table 3. The aggregate arrival rate per station on the network is thus

TABLE 1. Voice source characteristics

Vocoding rate	64 kbits/s
Talkspurt length	EXP(352) ms
Silent interval length	EXP(650) ms
Voice packetization delay	16 ms
Voice packet length	1024 bits
Mean no. of packets/talkspurt	22
Mean data rate/source	22.438 kbits/s
Node buffer capacity	500 packets

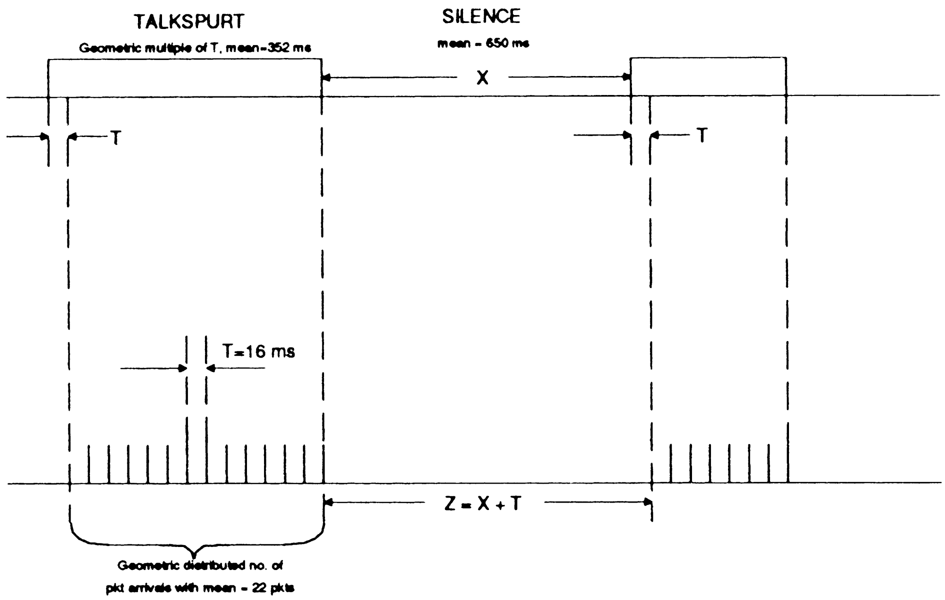

FIGURE 6. A single voice stream generated by the voice source generator.

TABLE 2. Bulk source characteristics

Mean data rate	1.2 Mbits/s
Inter-packet gap	0.6096 ms
Length of burst	20 consecutive packets
Length of packet	6000 bits
Inter-file gap	EXP(87.82) ms
Node buffer capacity	500 packets

TABLE 3. Interactive source characteristics

Mean data rate	0.3 Mbits/s
Packet inter-arrival time	EXP(1.25) ms
Packet length	375 bits
Node buffer capacity	500 packets

given as (1) voice transmission having 1 Mbits/s and (2) combined data transmission having 1.5 Mbits/s. Each station will be fed by 45 voice sources, one bulk and one interactive data source. The aggregate arrival rate per station is therefore 2.5 Mbits/s making the maximum number of stations on the network (without overloading) to be 40. As noted in the beginning of the paper, the urgency in voice packet delivery warrants the assignment of voice frames as synchronous class traffic and data as asynchronous. Furthermore, since interactive traffic requires a faster response time, it is assigned a higher priority of the two.

4. LOADING PATTERNS AND REQUIREMENTS

In this section, an approach similar to that taken by Janetzky and Watson (1987) is followed and two benchmark load patterns, BM1 and BM2 are introduced. Although BM1 and BM2 are artificial patterns, they are nevertheless intended to reflect typical loading configurations of large backbone networks; in particular, campus wide networks in academic institutions.

4.1. Benchmark load patterns

Benchmark load patterns coupled with performance requirements are a useful aid in gaining insight into what performance behavior in terms of waiting times and throughput can be expected from large backbone networks. They are usually used when actual load patterns are not readily available and can be extended as a basis of comparison between multiple systems.

The characteristics of BM1 are given in Table 4. This pattern is used to represent a large backbone FDDI network with variable number of gateways. All gateways (or stations) on the backbone have identical traffic distributions resulting from the superposition of 45 voice channels, a bulk and an interactive data source with respective arrival rates as described in the previous section.

BM2, attempts to simulate a large backbone network, with 40 stations but with asymmetric traffic distribution. This scenario is useful because real life LANSs are

TABLE 4. Benchmark load pattern 1 (BM1)

Voice data rate	1 Mbits/s
Bulk data rate	1.2 Mbits/s
Interactive data rate	0.3 Mbits/s
No. of voice sources/station	45
No. of bulk sources/station	1
No. of interactive sources/station	1
Voice source priority	synchronous
Interactive source priority	asynchronous priority 8
Bulk source priority	asynchronous priority 7
T_Opr	10 ms
T_Pri (8)	10 ms
T_Pri (7)	7.65 ms
Ring latency	0.1617 ms

seldom homogeneous, particularly large campus wide LANs. In such networks, it is not unusual to find a small proportion of the stations generating a large majority of the traffic on the LAN. Such an imbalance is likely to be observed, say when the LANs belonging to the Computer Science department is overloaded by students rushing to complete an assignment. The characteristics of BM2 are given in Table 5.

In this particular model, asymmetric loading is modeled by gradually increasing load on 10% of the stations (or four stations) in the network from an initial symmetrical situation (10% of the stations producing 10% of the load), to a final highly asymmetrical case where 10% of the stations produce more than 80% of the total load. The loading is simulated by increasing the number of bulk and interactive sources that feed stations 1–4 (high load stations) in equal proportions while keeping the number of these sources constant at one in stations 5–40 (low load stations). The physical arrangements of the stations is depicted in Figure 7.

4.2. Performance requirements

The performance requirements for both benchmark load patterns are given in Table 6. For packet voice communications, there are two important factors that must be taken into consideration packet loss rate and round trip delay. It was noted in Gruber and Le (1983) that up to 2% loss of voice packets was acceptable without signi-

TABLE 5. Benchmark load pattern 2 (BM2)

No. of stations	40
T_Opr	10 ms
T_Pri (8)	10 ms
T_Pri (7)	7.65 ms
Ring latency	0.1617 ms
Data rate of traffic sources	identical to BM1
For stations 5–40 (low load)	
no. of voice sources	45
no. of interactive sources	1
no. of bulk sources	1
For stations 1–4 (high load)	
no. of voice sources	45
no. of interactive sources	variable
no. of bulk sources	variable

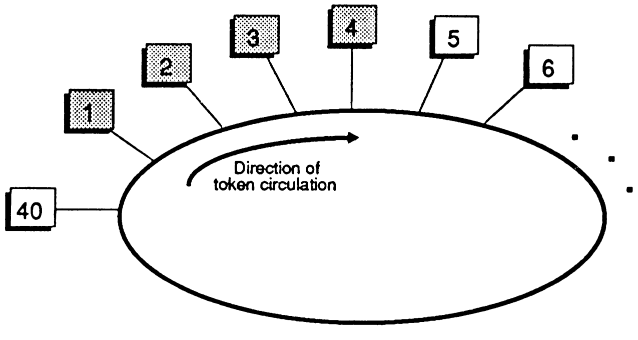

FIGURE 7. Physical arrangements of stations in BM2. □, High load stations. □, Low load stations.

TABLE 6. Performance requirements for BM1 and BM2

Voice traffic	
access delay	$\leqslant 20$ ms
waiting time	> 20 ms with probability $\leqslant 0.005$
Bulk traffic	
throughput	$\geqslant 85\%$ of offered load
Interactive traffic	
access delay	$\leqslant 50$ ms

ficant degradation of conversation. In this model however, the maximum loss rate was fixed conservatively at 0.5%. For the round trip delay, a conversation is liable to be disrupted by echoes of the speech signal if the delay is greater than several tens of milliseconds (Gruber and Le, 1983). Although this can be partially alleviated by the use of echo cancelling equipment, for simplicity, the maximum round trip delay is fixed at 40 ms. This means that any voice packet with a delay over 20 ms will be discarded by the network. For interactive traffic, the maximum delay for a packet is fixed at 50 ms. For bulk traffic, a minimum throughput requirement of 85% (of its offered load) is used.

5. SIMULATION RESULTS

A total of four series of runs were executed; three for BM1 and one for BM2. All simulation runs were for a duration of 10 s of simulated time with run statistics not collected until 3 s of simulated time had elapsed. This was because only the steady state behaviour of the system is of interest to us and visual inspection through the data from numerous runs indicate that 3 s were sufficient for the system to stabalize. The total simulation time of each run is restricted to 10 s in order to keep the actual duration of the run to a reasonable period. For each series of runs, five individual simulations were executed. The solid line drawn in the graphs presented in this section represent the mean of the five individual runs while the high/low values represent the maximum and minimum values observed.

5.1. Results of BM1

Figure 8 is a graph of the average throughput of the network versus the number of stations. During the initial

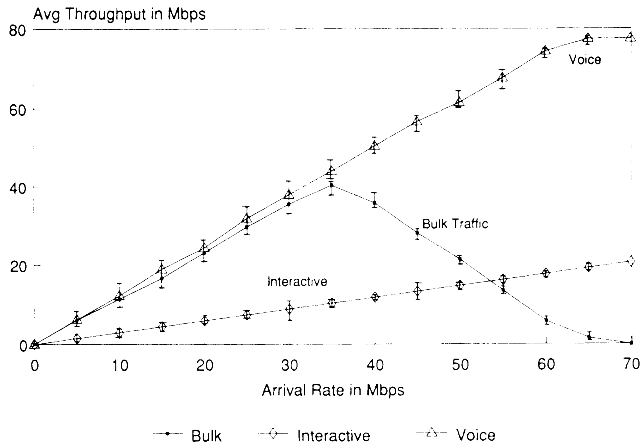

FIGURE 8. Average throughput versus no. of stations for BM1.

stages of loading; from 0 to 30 stations, throughput increases uniformly with load. There is very little waiting time involved, as can be seen from Figure 9, which is a graph of the average delay versus the number of stations, and arriving packets are transmitted well within one token rotation. The delay experienced by each packet is small, the major component being the time taken to wait for the next return of the token. As such, the mean waiting time in each class is therefore approximately equal to one half the average TRT. Also because of the low arrival rate and small transmission time of each station, the TRT value grows only very slowly (see Figure 10).

As the number of stations is increased from 30 to 40, the TRT value jumps from around 0.8 to over 7 ms as shown in Figure 10 which is a graph of the average TRT versus the number of stations. This value is near the token holding time threshold for bulk traffic, thus less bulk traffic is transmitted per token rotation. As a result, delay of bulk traffic frames grow quickly to infinity and throughput starts to decline when there are around 40 stations in the system. Beyond this point, performance of the bulk traffic declines gradually before being totally excluded at around 60 stations when the average TRT

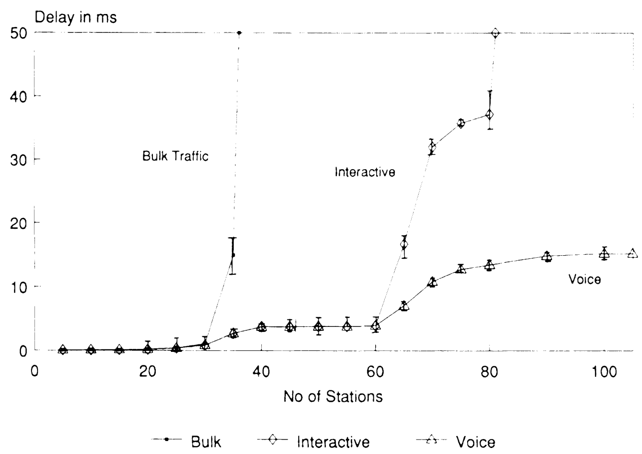

FIGURE 9. Average delay versus no. of stations for BM1.

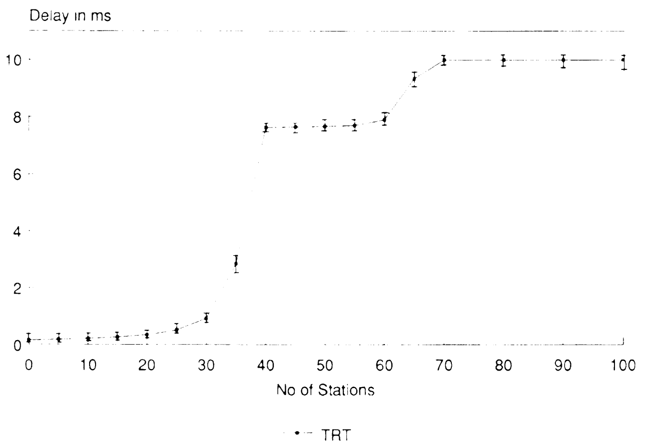

FIGURE 10. Average TRT versus no. of stations for BM1.

exceeds its token holding time threshold (of 7.65 ms as given in Table 4).

Throughput and delay of the two other classes remain stable up to 60 stations after which the delay for interactive traffic grows without bound. At this stage (60 stations), the total load on each station (3.75 Mbits/s) is approximately 150% of the channel capacity. Although the current load is well over 100%, the network is only just beginning to saturate. This can be seen from the TRT which is only around 9 ms. The reason for this is that access class 7 (bulk traffic) is virtually excluded from transmission (because of the high TRT) and hence the effective arrival rate is only equal to the sum of the arrival rates of the synchronous class (voice) and access class 8 (interactive) which totals up to about 78 Mbits/s. The network only becomes saturated when there are around 70 stations. At this point the average TRT is around 10 ms and the delay of synchronous frames begin to build up.

Percentage of voice packets lost are shown in Figure 11. Note that there is very little packet lost until the network saturates at around 70 stations. This is because initially as the number of stations increase, the bandwidth allocated to packet voice grows at the expense of other data classes, causing a proportionate rise in its throughput, thereby checking delay growth. However, as the number of stations reach 60, the delay of interactive packet rises without bound and the network begins to saturate. Between 60 and 70 stations, the average delay for voice packets rises gradually, but does not yet reach the 20 ms bound (which we have defined in our performance specifications) and hence the average packet loss rate is still relatively low. Beyond 70 stations, the network saturates as all possible time left has already been allocated to voice transmission and so average delay rises quickly above the 20 ms threshold. Packet loss rate rises dramatically after this and the network begins to choke itself.

From the configuration given in BM1, the threshold for interactive traffic (T_pri(8)) and bulk traffic (T_pri(7)) is set at 10 and 7.65 ms, respectively. Using the performance requirements specified earlier, we find that the maximum number of stations that the network can support and still satisfy the requirements is around 35. Closer examination of this critical point reveals that the main constraint in preventing more stations from joining the network is the requirement that 85% of the offered load of bulk traffic must be transmitted. Other constraints (such as voice traffic loss rate and interactive traffic delay) only apply when there are over 54 stations. Therefore it would seem that if the threshold value of priority 7 were raised thereby relaxing the throughput constraint, more stations could be supported.

Figure 12 shows the relation between the maximum number of stations that can be supported and the T_Pri(7)/T_Pri(8) ratio, confirming the above hypothesis while Figure 13 plots the ratio of throughputs to the ratio of priority threshold values. In Figure 13, the throughput ratio is obtained by premultiplying the throughput of priority 8 by 4 before being using it to

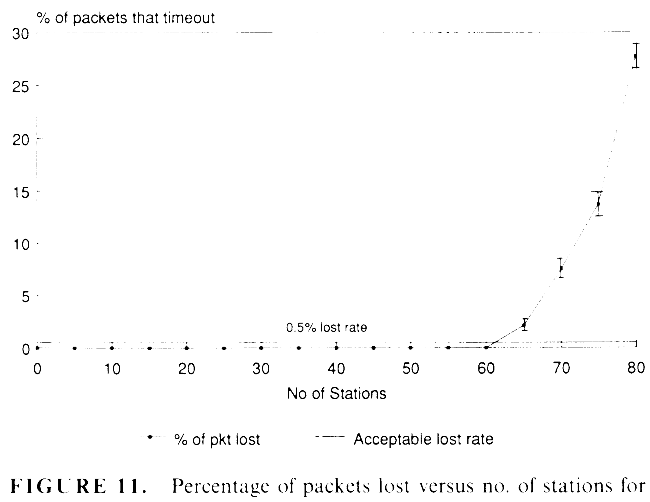

FIGURE 11. Percentage of packets lost versus no. of stations for BM1.

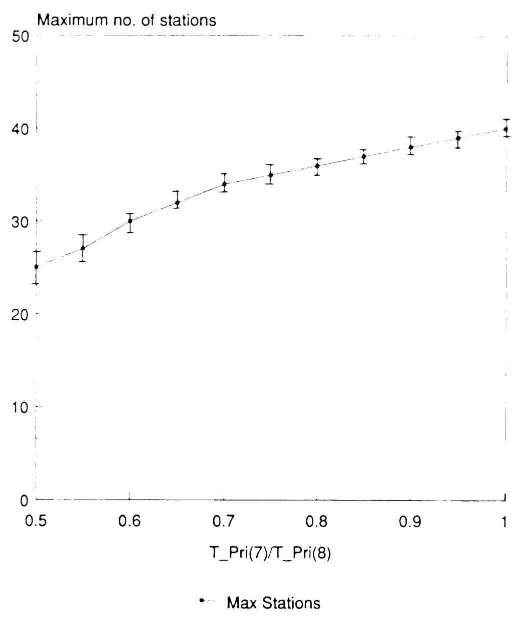

FIGURE 12. Max no. of stations versus T_Pri ratio for BM1.

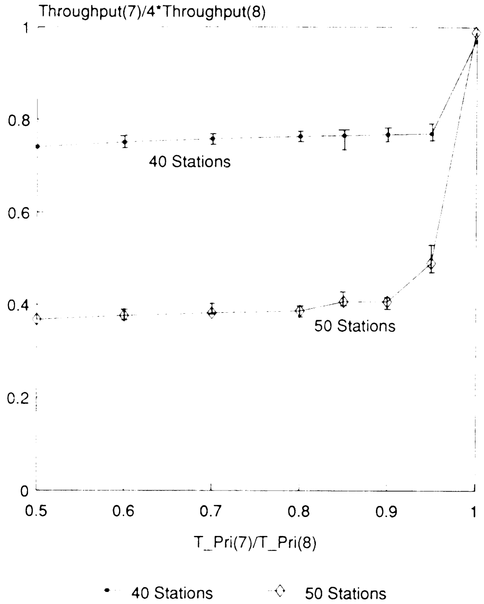

FIGURE 13. Throughput ratio versus T_Pri ratio for BM1.

divide the throughput of priority 7. This is because the offered load of priority 8 is set to be 4 times lower than the corresponding load of priority 7 and the premultiplication is necessary to bring them to an equal basis. Initially, when the ratio of T_Pri(7)/T_Pri(8) is low, the transmission window for priority 7 is small and its corresponding throughput ratio is low. As the priority threshold ratio is raised, the throughput of priority 7 improves, raising the throughput ratio, up to the point when T_Pri(7) equals T_Pri(8) and the two priority levels are indistinguishable.

5.2. Results of BM2

The throughput and delay characteristics of BM2 bulk traffic class (priority 7 traffic) are given in Figures 14 and 15. The horizontal axis of the figures represent the percentage of the total offered load produced by four adjacent high load stations in the network.

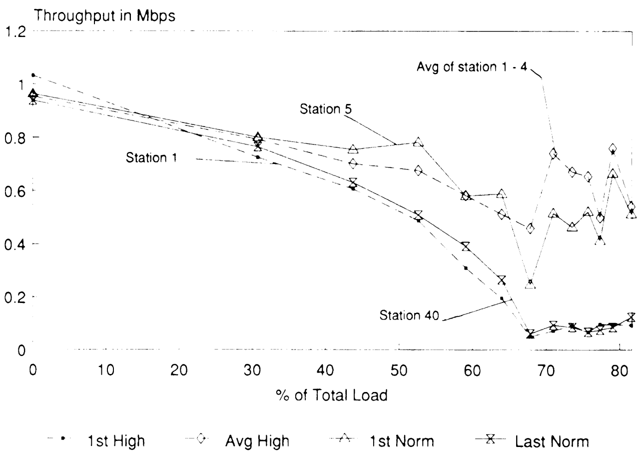

FIGURE 14. Average bulk traffic throughput versus percentage of asymmetric load for BM2.

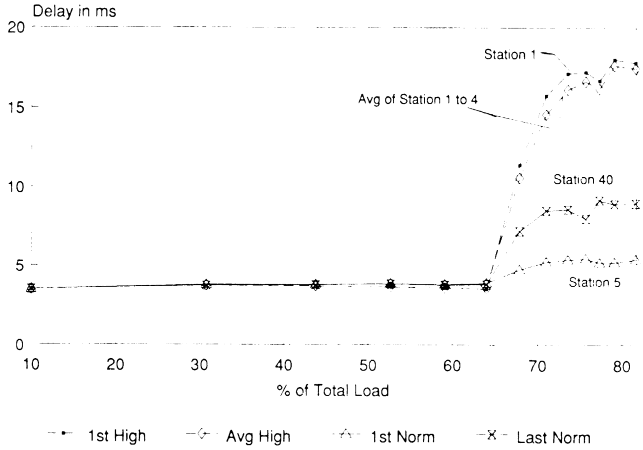

FIGURE 15. Average delay of interactive traffic versus percent asymmetric load for BM2.

As can be seen from Figures 14 and 15, there is much disparity in the values obtained between similar stations as load on the network is increased. Initially when the network is symmetrically loaded, the average total throughput is spread evenly across all stations and each station produces roughly the same level of throughput. As the network becomes increasingly asymmetrical, the disparities in throughput become evident.

For example, when the asymmetric load on the network reaches around 70%, the difference in throughput levels of priority 7 traffic between the station (station 5) that is directly downstream from the high load stations (stations 1–4) and the station that is furthest downstream from the high load stations (station 40) is more than 5 times the lower of the two. A similar effect can be observed for the throughput levels between the first high load station (station 1) and the average high load stations excluding the first (stations 2–4).

An explanation for this unfairness is as follows. When a high load station does not have a full queue of frames to transmit, the unused portion of its bandwidth will be available for use by the lower load stations downstream. This unallocated bandwidth will circulate sequentially among all the stations with each passing of the token, beginning with the first station immediately downstream from the high load stations. When the load is low, high and low load stations are indistinguishable and any unused bandwidth is spread out among the stations more or less uniformly. However, when the asymmetric load becomes high, high load stations will use up more and more of their bandwidth, resulting in less and less unused bandwidth. Because of the decreasing quota of unused bandwidth and the sequential nature of its circulation, stations immediately following high load stations tend to be favoured over other stations further downstream. A similar phenomenon was observed in Dykeman and Bux (1988), except that in their case, asymmetry was achieved by having stations transmit only a single priority with this priority either being high or low.

Priority 7 frames are not the only ones that are subjected to unfairness when asymmetric load is high. Delay characteristics of priority 8 frames also exhibit similar unfairness albeit to a lesser degree because of their higher priority. Figure 14 gives the delay curves of priority 8 traffic. Notice that the same pattern is observed here with the delay of station 5 significantly lower than the delay of station 40 when asymmetric loading is high.

6. CONCLUSION

This paper has presented a simulation model that can be used to study the effects of symmetric and asymmetric loading in a multimedia backbone network. The simulation results based on BM1 show that a far greater number of stations can be supported on such networks than is readily apparent. This is because careful setting of the token holding thresholds of the various priorities allow the selection of an appropriate mix of traffic classes so that any excess bandwidth in one traffic class can be transferred to another. However, these setting would have to take into consideration the various performance requirements of the individual classes so that the quality of service of each class is not compromised.

Another area studied in this paper was the effect of asymmetric loading in large backbone networks (based on BM2). As was pointed out by a number of researchers (Dykeman and Bux, 1987, 1988) the FDDI protocol becomes unfair under increasing asymmetric loads. The inherent unfairness of the FDDI MAC protocol during high asymmetric loads is disturbing for two reasons. The first reason is the fact that many real life networks fall into this category. The second reason is that there is no known 'cut-off' point after which the network becomes significantly unfair. In other words, it is very hard to define the safe operating range of a network so that good and fair performance is guaranteed.

One solution to this problem would be to allow the priority threshold values of each priority in each individual station to be adjusted dynamically to compensate for the bias. Limb (1988) demonstrates the feasibility of such a scheme in controlling the throughput of unidirectional bus and token ring networks and (Deng *et al.*, 1994) applies a similar algorithm to the FDDI. These two works apparently show that a dynamic modification of crucial network parameters through the calculation of a simple modification factor is sufficient to overcome the weakness of the protocol. It is worthwhile to note, however, that the above two algorithms assume a certain amount of independence between the modification of a parameter on one station and its effect on other stations. This assumption is valid so long as there are a large number of stations (typically in the hundreds) and updates occur asynchronously. In backbone networks however, the number of gateways is usually rather small (typically in the tens) and so the algorithms may not be so effective when applied to this context.

REFERENCES

ANSI (1985a) Token ring access method and physical layer specifications. *ANSI 802.5-1985*.

ANSI (1985b) Token-passing bus access method and physical layer specifications. *ANSI 802.4-1985*.

ANSI (1987) FDDI token ring media access control (MAC). *ANSI X3.139-1987*.

ANSI (1988) FDDI token ring physical layer protocol (PHY). *ANSI X3.148-1988*.

Bux, W. (1989) Token-ring local area networks and their performance. *Proc. IEEE*, **77**, 00–00.

Deng, R. H. Lim, K. S. Ranai, K. and Huang, K. T. (1994) Performance control in timed-token passing LANs. *IEEE Trans. Commun.*, submitted.

Dykeman, D. and Bux, W. (1987) An investigation of the FDDI media-access control protocol. Paper presented at *Proc. 5th European Fibre Optics Communications Local Area Networks Conf.*, Basel, Switzerland.

Dykeman, D. and Bux, W. (1988) Analysis and tuning of the FDDI media access control protocol. *IEEE J. Select. Areas Commun*, **SAC-6**, 00–00.

Frontini, M. and Watson, G. (1988) An investigation of packetized voice on the FDDI token ring. In *Proc. 1988 Int. Sem. on Digital Commun.* Plattner, B. and Gunzburger, P. (eds),

Gruber, J. G. and Le, N. H. (1983) Performance requirements for integrated voice/data networks. *IEEE J. Select. Areas Commun.*, **SAC-1**, 00–00.

Janetzky, D. and Watson, K. S. (1987) Performance evaluation of the MAP token bus in real time applications. Kummerle, K, Tobagi, F. A. and Limb, J. O. (eds), *Advances in Local Area Networks*. IEEE Press, New York.

Jayant, N. S. and Christensen, S. W. (1987) Effects of packet losses in waveform coded speech. *IEEE Trans. Commun.*, **COM-20**, 00–00.

Johnson, M. J. (1988) Performance analysis of FDDI. Paper presented at *Proc. 6th European Fibre Optics Communication Local Area Networks Conf.* Amsterdam, Netherlands.

Limb, J. O. (1989) Load-controlled scheduling of traffic on high-speed metropolitan area networks. *IEEE Trans. Common.*, **COM-37**, 00–00.

Martini, P., Spaniol O, and Welzel T. (1988) File transfer in high-speed token ring networks: performance evaluation by approximate analysis and simulation. *IEEE J. Select. Areas Commun.*, **SAC-6**, 00–00.

Ross, F. E. (1989) An overview of FDDI: the fiber distributed data interface. *IEEE J. Select. Areas Commun.*, **SAC-7**, 00–00.

Sauer, C. H. Blum, A. M. MacNair, E. A. and Kurose, J. F. (1986a) *The Research Queueing Package Version 2 CMS Reference Manual*. IBM, New York.

Sauer, C. H. Blum, A. M. MacNair, E. A. and Kurose, J. F. (1986b) *The Research Queueing Package Quick Reference Summary*. IBM, New York.

Schill, A and Zieher, M. (1987) Performance analysis of the FDDI 100 Mbit/s optical token ring. In *High Speed Local Area Networks*. IFIP.

Sevcik, K. C. and Johnson, M. J. (1987) Cycle time properties of the FDDI token ring protocol. *IEEE Trans. Softw. Engr.* **SE-13**, 00–00.

Sriram, K. and Whitt, W. (1986) Characterizing superposition arrival process in packet multiplexers for voice and data. *IEEE J. Select. Areas Commun.*, **SAC-4**, 00–00.

Welzel, T. (1987) Analyse und messung der workload. In *Proc. Commun. Distributed Syst., Informatik Fachberichte 130.*, Springer, Berlin (in German).

Yatsuzka, Y. (1982) Highly sensitive speech detector and high-speed voiceband discriminator in DSI-ADPCM system. *IEEE Trans. Commun.*, **COM-30**, 00–00.