Performance of Alternative Media for LAN Interconnection

H.S. Hassanein†,V.J. Friesen‡, J.W. Wong‡ and J.W. Mark†

‡Department of Computer Science
†Department of Electrical and Computer Engineering
University of Waterloo, Waterloo, Ontario, CANADA N2L 3G1

Abstract

The profusion of LANs and the substantial increase in LAN traffic volume have created a demand for LAN interconnection that facilitates the move towards B-ISDN. This is an important step towards the realization of a "global communications network". In this paper, we study the performance of an interconnection scenario where servers, workstations, and LANs are interconnected by an interconnecting medium. ATM switching technology is an attractive candidate for the interconnecting medium because high-speed switches can be built at relatively low cost. A competing technology is DQDB, which is the IEEE 802.6 standard for Metropolitan Area Networks. A model is developed to study the performance differences of ATM switches and DQDB when used as interconnecting media. Simulation results for different network configurations and traffic scenarios are presented.

1: Introduction

Broadband Integrated Services Digital Network (B-ISDN) [1] is becoming widely accepted as the public network of the future. It uses Asynchronous Transfer Mode (ATM) as its switching and multiplexing technique. Bandwidth is allocated on demand using 53-byte data units called cells [2]. ATM supports a wide range of services, including voice, data and video, and provides a flexible means to multiplex these traffic types onto the same physical network.

Recently, we have seen a profusion of Local Area Networks (LANs) and a substantial increase in the volume of LAN traffic. Such activities create a demand for LAN interconnection that facilitates a move towards B-ISDN. This is an important step towards the realization of a "global communications network". In this paper, we study the performance of an interconnection scenario where servers, workstations, and LANs are interconnected by an interconnecting medium; this medium is also connected to B-ISDN (see Figure 1).

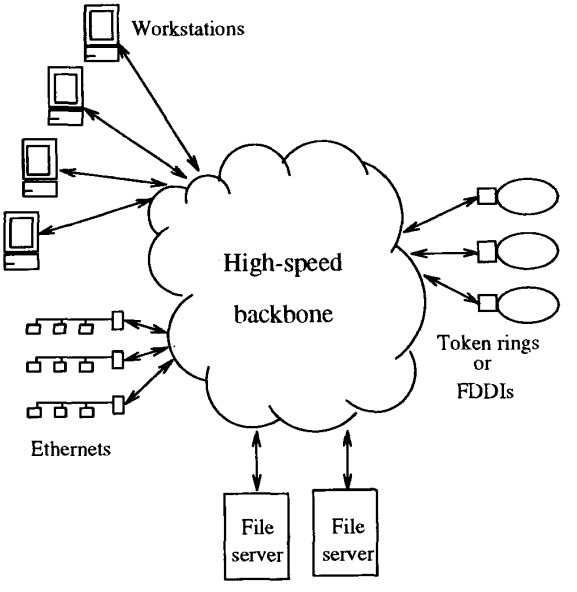

Figure 1
LAN Interconnection

ATM switching technology is an attractive candidate for the interconnecting medium because high-speed switches can be built at relatively low cost. A competing technology is Distributed Queue Dual Bus (DQDB), which is the IEEE 802.6 standard for Metropolitan Area Networks [3]. DQDB is also cell based, with the same cell size as that of ATM. The topology of DQDB consists of two time-slotted unidirectional busses providing

0-8186-4270-X/93 $03.00 © 1993 IEEE

communication in opposite directions. The distributed queue is implemented by requiring each station to maintain a counter which indicates the number of data segments (cells) to be transmitted ahead of this station by downstream stations. When a station has a segment to transmit, it uses of the contents of this counter to determine when to access the channel. DQDB is, therefore, a shared-medium technology. DQDB is compatible with ATM in terms of cell size and payload. It is also capable of providing similar services.

In this paper, we compare the performance of the two interconnecting media mentioned above, namely ATM switching and DQDB. We develop a performance model and obtain simulation results for various network configurations and traffic scenarios. Our results would help in justifying the use of ATM switches or DQDB as interconnecting media for LAN interconnection. Indeed, we show that by configuring the network properly, the delay incurred within the interconnecting media can be made only a small fraction of the total delay.

This paper is organized as follows. Section 2 describes our performance model. Section 3 presents a performance comparison between an ATM switch and DQDB as interconnection media. Finally, Section 4 summarizes our findings and presents some possibilities for future work.

2: Performance model

In this section we describe the performance model used in our simulation.

2.1: Network model

We consider a scenario consisting of LANs and hosts (servers or workstations) interconnected by an interconnecting medium. The LANs (Ethernets) are modeled explicitly while the hosts are modeled as packet sources. In this paper, the data unit within the interconnecting medium is called a cell, and that outside the interconnecting medium is called a packet.

2.1.1: ATM switch: Our model consists of a single ATM switch and a number of multiplexers/demultiplexers. The ATM switch is an internally non-blocking, output-buffered fast packet switch. Multiplexers/demultiplexers may be used to combine/separate streams of traffic from and to a group of LANs and hosts. Interface modules, situated between each LAN or host and the ATM portion of the network, perform segmentation and reassembly of packets, a function of the ATM Adaptation Layer (AAL).

We assume that permanent connections are established between each source-destination pair and maintained throughout the simulation. We further assume that the payload in an ATM cell is 48 bytes and that cell-by-cell forwarding is performed. These assumptions are compatible with the services provided by AAL type 5 [4].

Cells arriving at an output port of the switch, or at a multiplexer, are handled independently, and are buffered in a single queue which is serviced first-in-first-out (FIFO) by the outgoing link. We assume that all links between the interface modules and the switch carry ATM cells at a rate of 45 Mb/s. The time unit used in the simulation is the time required to transmit one 53-byte ATM cell on a 45 Mb/s channel (approx. 9.42 μs). This time unit is referred to as a *slot*. All transmission lines in our model are assumed to be error-free.

Packets arriving at an interface module (from a host or an Ethernet) are segmented into ATM cells before being transmitted on the ATM portion of the network. Each of these cells is given a sequence number and one of the following *type* indicators: SCM (single-cell message), BOM (beginning of message), EOM (end of message), or COM (continuation of message). Sequence numbers and cell types are irrelevant within the ATM portion of the network, where cell-by-cell forwarding is performed. However, they are used by an interface module to reassemble cells from the ATM portion and to detect incomplete packets. Incomplete packets are simply dropped. We assume that the cells comprising a given packet are received in sequential order by the interface module unless they are dropped at the switch. This follows because all cells belonging to a given source-destination pair will follow the same route through the network.

2.1.2: DQDB: The DQDB network in our simulation is modeled explicitly at the medium access control (MAC) layer in accordance with the specification in the IEEE 802.6 standard document [3]. The DQDB network consists of K equidistantly-spaced interface stations (each corresponding to an interface module for an Ethernet or a host), and two Headend station (HOA and HOB for busses A and B, respectively). Headend stations generate empty slots, 53 bytes each, at a rate of 45 Mb/s.

As in the ATM switch model, packets arriving at an interface station (from a host or an Ethernet) are segmented into cells. These cells are then buffered for transmission on a DQDB bus. Cells from station i destined to station j are forwarded on bus A if $i<j$ and on bus B if $i>j$. Within the DQDB portion of the network, each of these cells independently contends for channel access. Cells are 53 bytes in length with a 48-byte payload; this is compatible with AAL type 5.

Destination DQDB stations are responsible for reassembly of packets. Reassembled packets are then forwarded to Ethernets or hosts. It should be noted that segmentation and reassembly are part of the DQDB protocol and that a DQDB station (or an interface module) may be thought of as resembling a remote MAC bridge. This would then comply with the transparent bridge specification for DQDB [5].

Our model also allows for the erasure node segment-clearing function. We implement the erasure node algorithm in [6] which is being considered as the basis for the IEEE 802.6e erasure node standard [7].

2.1.3: Ethernet: The Ethernets are modeled explicitly at the MAC layer. Each of them has N stations spaced equidistantly along the channel. Stations 1 to $N-1$ are regular Ethernet stations, while station N is the interface module. A packet requesting transmission at one of the $N-1$ Ethernet stations is placed in a station buffer and serviced on the channel according to the CSMA/CD protocol. Some percentage of the packets generated within a given Ethernet have destinations external to that Ethernet. These packets are received and forwarded by the interface module. Packets arriving from outside the Ethernet are collected by the interface module, which then contends for the channel with the other stations, to transmit the arriving packets to their respective destinations.

2.2: Traffic models

We assume that for each Ethernet, packet arrivals at a given station follow a Poisson distribution with rate λ_E. A variety of packet length distributions are considered. All inter-arrival times, packet lengths and destination assignments are assumed to be independent.

The host traffic is assumed to consist mainly of bursty interactive traffic and file transfers. We therefore model the arrival of packets from a host to its interface module as an Interrupted Poisson Process (IPP) (see [8] and [9]). The IPP source model alternates between an active state, during which arrivals are Poisson with rate λ_H, and an idle state, during which no arrivals occur. This model is commonly used to account for bursty traffic sources. As with the Ethernet packets, various packet-length distributions are considered, and no dependencies are assumed.

3: Performance comparison

In this section we present experiments performed using our simulation model. We compare the performance of a number of possible ATM switch and DQDB configurations for connecting eight Ethernets and two hosts.

Two ATM switch configurations are considered. The first configuration consists of two multiplexers, each combining the traffic from four Ethernets and one host (see Figure 2). In the second configuration, we again use two multiplexers, but each is connected to only the four Ethernets, while the hosts have direct links to the ATM switch (see Figure 3). Interface modules in these two diagrams are labeled "I". As well, three DQDB configurations are considered. In the first, hosts are placed on extreme ends of the DQDB bus close to Headend nodes (see Figure 4). In the second, hosts are placed in the middle of the bus (see Figure 5). The third configuration makes use of erasure nodes to achieve spatial reuse on the bus (see Figure 6). For ease of reference, these configurations will be referred to as configuration ATM-1, ATM-2, DQDB-1, DQDB-2 and DQDB-3, respectively. To simplify the presentation of the DQDB configurations, Figures 4-6 do not explicitly show either Ethernets or hosts. Instead the figures only show interface modules marked "E" (denoting an Ethernet interface) and "H" (denoting a host interface).

Note that configurations DQDB-1 and DQDB-2 are similar to ATM-1 which uses 2 output ports of the switch. That is, all three configurations have 90 Mb/s of total bandwidth available to them (two 45 Mb/s channels). Configuration DQDB-3, in which the host traffic is isolated by erasure nodes, is an attempt to construct a DQDB configuration which is similar to ATM-2 where hosts have separate links.

To ensure a fair comparison between ATM switches and DQDB as interconnecting media, we set the propagation delays such that all configurations span an equivalent geographical area. Assume that the ATM switch is at the center of an imaginary circle, with the interface modules on the perimeter, in either configuration ATM-1 or ATM-2. We set the distance between DQDB stations such that the one-way end-to-end propagation delay is equal to the circumference of the circle. For the experiments in this section, the radius of the ATM configurations is 4 slots in length, resulting in a one-way end-to-end propagation delay of 25 slots for the DQDB configurations.

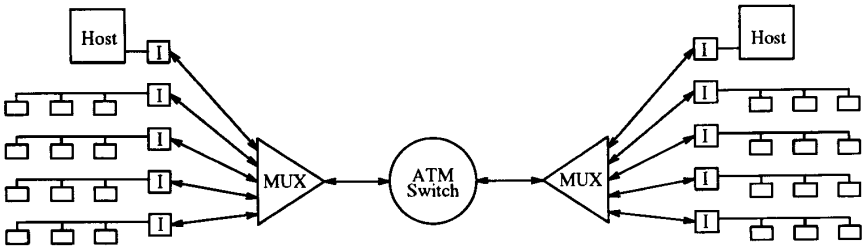

Figure 2
Configuration ATM-1

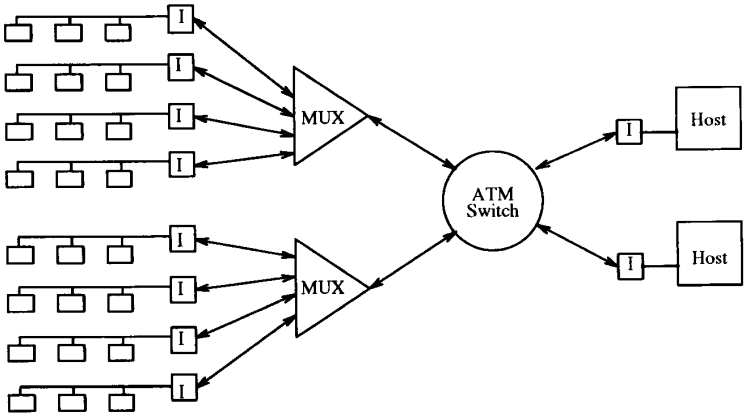

Figure 3
Configuration ATM-2

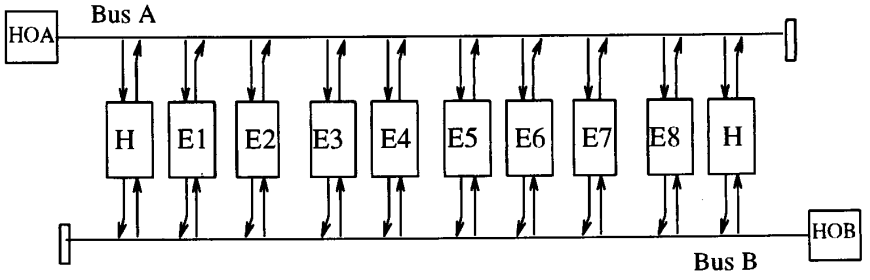

Figure 4
Configuration DQDB-1

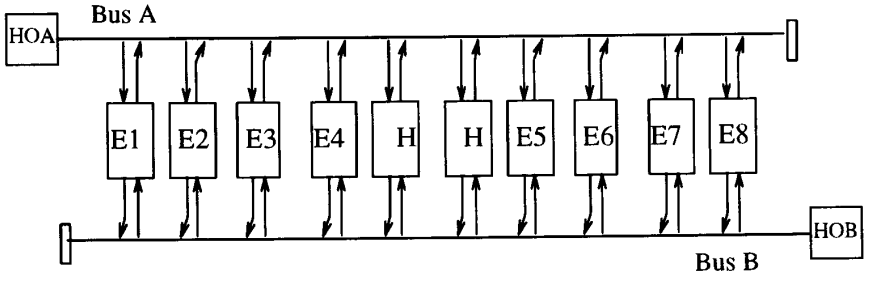

Figure 5
Configuration DQDB-2

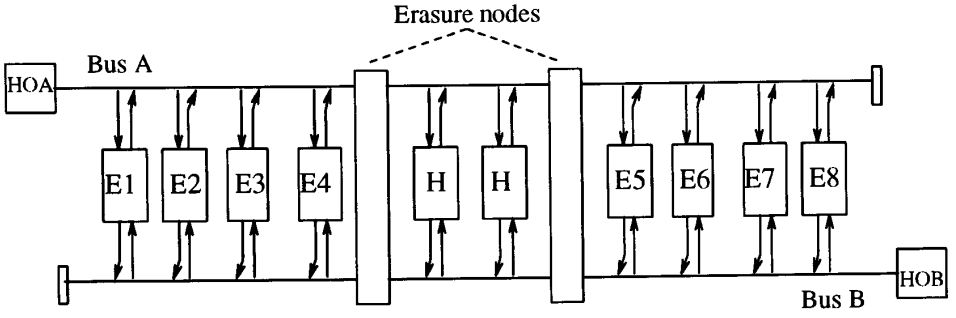

Figure 6
Configuration DQDB-3

Each Ethernet in the configurations considered is composed of $N = 4$ stations, one of which is the interface module. The propagation delay between consecutive stations is 5 bytes in length. Of the packets generated at the regular Ethernet stations, 50% are 64 bytes in length (control packets), while the other 50% are 1500 bytes in length (file transfers). This bimodal model for packet-length distribution is used to approximate traffic measurements performed on Ethernets having client-server traffic (see [10] and [11]), where a large proportion of the packets were found to be either the minimum or the maximum Ethernet packet length. The offered load from the regular stations is 30% (i.e., 3 Mb/s), and an additional load is added as a result of the incoming traffic from the interface module. Half of the packets generated by the regular stations leave the Ethernet. 20% of these are addressed to hosts (10% to each) and the other 80% are equally dispersed amongst the remaining seven Ethernets. Each Ethernet station can buffer up to 4096 bytes of data. The interface module is also equipped with a reassembly buffer.

The host traffic sources have a peak bit rate of 45 Mb/s and a mean burst length of 100 slots (close to 1 ms). The mean bit rate is varied to observe different loads on the network. Host packet lengths vary as follows. 50% of the packets are short (64 bytes), corresponding to control packets. The other 50% of the packets are long, corresponding to file transfers. Long packets destined for an Ethernet are 1500 bytes in length, while those transferred between hosts are 2400 bytes (50 cells) in length. 80% of host packets are addressed to the other host, and the other 20% of the host packets are uniformly dispersed amongst the eight Ethernets.

In the following experiments we assume that all buffers in the ATM configurations, and the interface station buffers in the DQDB configurations, are infinite. We, therefore, do not report packet/cell loss statistics. Even though such a performance measure is of importance in B-ISDN, setting ATM switch and DQDB configurations and parameters to ensure a fair comparison with finite buffers is not a simple task. This issue is currently under further study.

Table 1 shows the mean delay (in slots) and delay variance experienced by Ethernet packets when the mean bit rate of each host is 27 Mb/s. (The 95% confidence intervals for the mean values are within ± 5% of the presented values.) The first set of results (station) represent the elapsed time from packet generation to delivery for packets transmitted from a station on one Ethernet to a station on another Ethernet. The second set of results (interface) represents the delay from when an Ethernet packet is segmented at the source interface module until it is reassembled at the destination interface module. These results, then, represent delay incurred within the interconnecting medium. Since all DQDB configurations are symmetric, and as the DQDB busses are identical, we report station and interface delays for Bus A only.

Delay Type	Configuration	Mean Delay (slots)	Delay Variance
Station	ATM-1	411	1.4×10^5
	ATM-2	333	1.4×10^5
	DQDB-1	421	2.1×10^5
	DQDB-2	353	1.4×10^5
	DQDB-3	343	1.4×10^5
Interface	ATM-1	116	1.1×10^4
	ATM-2	35	4.4×10^2
	DQDB-1	123	4.3×10^4
	DQDB-2	52	1.9×10^3
	DQDB-3	42	1.3×10^3

Table 1
Delay statistics for Ethernet packets

We first note that, for the scenario in Table 1, the delays within the Ethernet portion of the network dominate the end-to-end delay performance. (This Ethernet delay is almost constant regardless of the configuration or the interconnecting media — about 300 slots for the scenario in Table 1.) For both interconnecting media, both the station and interface delays drop quite significantly as we isolate the traffic-intensive hosts (see results for ATM-2, DQDB-2 and DQDB-3). In fact, the interface delay becomes a small fraction of the end-to-end delay. This is further demonstrated by studying the effect of host traffic on Ethernet packet delays. Figure 7 plots interface-to-interface delay for various loads provided by the host traffic. By isolating the traffic-intensive hosts, the performance experienced by Ethernet packets becomes much less sensitive to the load.

Comparing the Ethernet interface-to-interface delay for the different configurations, we note the following:

- DQDB-2 clearly outperforms DQDB-1, a configuration with equivalent bandwidth. This is because the delay performance of DQDB stations is location dependent [12]. As the hosts are traffic-intensive devices, the placement of these hosts in the middle of the

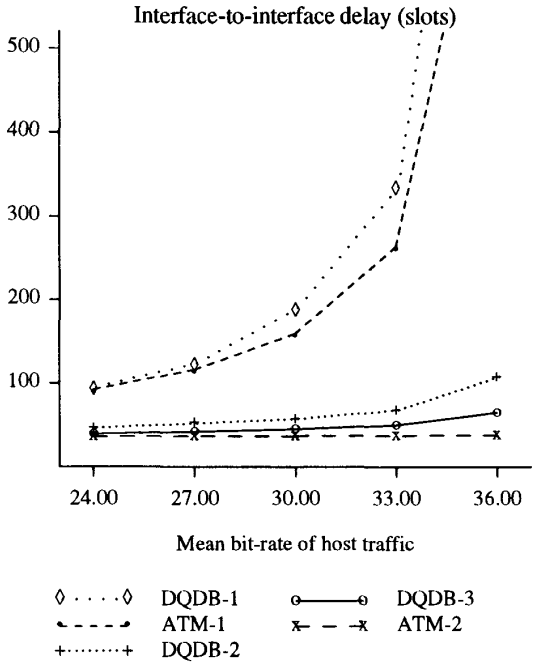

Figure 7
Ethernet packet delay for various loads

bus (DQDB-2) is more advantageous to Ethernet interface modules. This, however, comes at the expense of longer delays for host interface modules. (At host mean bit-rate of 27 Mb/s, the average host packet interface delay was 139 slots for DQDB-2 as opposed to 101 slots for DQDB-1.)

- Not only does configuration DQDB-2 result in an improvement in Ethernet interface delay, that delay is also less sensitive to the host load.
- Little improvement in Ethernet delay performance is achieved by isolating the host traffic using erasure nodes (DQDB-3). This is clearly demonstrated in Figure 7. The use of erasure nodes, however, results in a much improved *host* delay performance, as more bandwidth is now available to the host interface modules.
- Isolating the host traffic also improves the delay variance within the interconnecting medium. We note that the delay variance for ATM-2 is much lower than for the other configurations (see Table 1). As the load increases, delay variances for DQDB configurations increase more rapidly than for ATM-2.
- While DQDB-2 was shown to outperform ATM-1 (a configuration with equivalent bandwidth), it does not provide uniform access delays for the different Ethernet interfaces. Figure 8 plots the interface delays of the Ethernet interface modules for the different DQDB configurations. Note that except for DQDB-3, the interface delay generally increases with the distance from HOA. The use of erasure nodes (DQDB-3) results in a more uniform interface delays. This effect has been demonstrated earlier in [13].

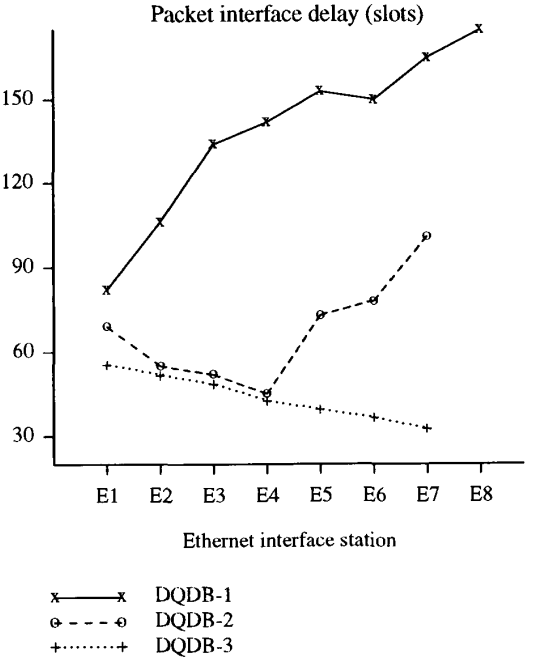

Figure 8
Ethernet packet interface delay for different interface stations

We have compared the performance of ATM switches and DQDB as interconnecting media under similar network configurations and traffic scenarios. Our observations show that ATM-2 outperforms all other configurations. This is because ATM-2 is the only configuration that totally isolates inter-host traffic from the rest of the network. As the load and the number of devices in the network increase, the expandability of ATM switches, i.e., by using more switch ports, makes them a more attractive interconnecting medium.

4: Conclusions and future work

We have developed a simulation model to study the performance differences of ATM switches and DQDB as interconnecting media for LAN/host interconnection. Our results revealed that the placement of traffic-intensive hosts greatly impacts the performance of inter-LAN traffic. When using ATM switches, such hosts should not be multiplexed with LAN traffic. For inter-LAN traffic on DQDB, it is more advantageous to place the hosts in the middle of the bus.

While our experiments show that ATM switches and DQDB (with the proper configuration) have similar performance characteristics, ATM switching has the following advantages. First, it provides more uniform delays among interface stations as opposed to the location-dependent delays in DQDB. Second, an ATM interconnecting medium yields a lower delay variation. This is significant because the public interconnecting medium will need to provide acceptable delay variation to some traffic types. Third, compared to shared medium DQDB, ATM has a higher aggregate throughput because it allows multiple simultaneous data transfers through the switch. Moreover, switches may be interconnected to further increase the aggregate capacity.

The interconnection of dissimilar LANs and devices over a high-speed backbone has a number of inherent problems. First, a high-speed device such as a file server or a traffic-intensive host can be sending a traffic stream to a low-speed LAN, e.g., Ethernet. Second, more than one device or LAN may attempt to simultaneously communicate with the same device or LAN. Since the speed of the backbone is much higher than that of the individual LANs, packets may be admitted to the backbone network, but could be dropped later at the destination LAN interface module. Without an effective congestion control mechanism, this can lead to performance degradation. Such a mechanism requires further investigation.

Acknowledgement

This research was supported by a grant from the Canadian Institute for Telecommunications Research under the NCE program of the Government of Canada.

References

[1] CCITT Rec. I.121 "Broadband Aspects of ISDN," Dec. 1990.

[2] CCITT Rec. I.327 "Broadband ISDN Functional Architecture," June 1992.

[3] IEEE Standard 802.6, "Distributed Queue Dual Bus (DQDB) Metropolitan Area Network (MAN)," Dec. 1990.

[4] CCITT Rec. I.363 "ATM Adaptation Layer (AAL) Specification," June 1992.

[5] IEEE Working Document, "Remote LAN Bridging Over DQDB Metropolitan Area Networks," IEEE 802.6i/D5 Unapproved Draft, Mar. 1993.

[6] H. S. Hassanein, J. W. Wong and J.W. Mark, "An Effective Erasure Node Algorithm for Enhancing Slot Reuse in DQDB," Contribution 802.6-1992/43, IEEE 802.6 working group, Nov. 1992.

[7] IEEE Working Document, "Erasure Node Algorithm for Slot Reuse in Distributed Queue Dual Bus (DQDB) Subnetwork of A Metropolitan Area Network (MAN)", IEEE 802.6e Unapproved Draft, Mar. 1993, Rapporteurs: H. S. Hassanein, J. W. Wong and J.W. Mark.

[8] A. Kuczura, "The Interrupted Poisson Process as an Overflow Process," *Bell Syst. Tech. J.*, vol. 52, no. 3, pp. 437-448, Mar. 1973.

[9] J.Y. Hui, "Resource Allocation for Broadband Networks," *IEEE Journal on Selected Areas in Comm.*, Vol. 6, No. 9, Dec. 1988, pp. 1598-1608.

[10] R. Gusella, "A Measurement Study of Diskless Workstation Traffic on an Ethernet," *IEEE Trans. on Commun.*, Vol. 38, No. 9, Sept. 1990, pp. 1557-1568.

[11] L.J. Bottomley and A.A. Nilsson, "Traffic Measurements on a Working Wide Area Network," *Teletraffic and Datatraffic in a Period of Change*, ITC-13, 1991, pp. 767-772.

[12] H. Kaur and G. Campbell, "DQDB — An Access Delay Analysis," *Proceedings INFOCOM '90*, pp. 630-635.

[13] H. S. Hassanein, J. W. Wong and J.W. Mark, "An Effective Erasure Node Algorithm for Slot Reuse in DQDB," *Submitted for Publication, IEE Proceedings-I, Jan. 1993.*