Loss and Delay Study of LAN Interconnection over ATM

Mahbub Hassan
School of Computing
and Information Technology
Monash University, Gippsland Campus
Churchill 3842, Australia
E-mail:Mahbub.Hassan@fcit.monash.edu.au

Mohammed Atiquzzaman
Department of Electrical
and Computer Systems Engineering
Monash University, Clayton Campus
Clayton 3168, Australia
E-mail:atiq@eng.monash.edu.au

Abstract

LAN interconnection is expected to be the first service offered by ATM based B-ISDN. Since traffic on the current LANs is connectionless and bursty, the main challenge will be to provide an efficient connectionless service over connection oriented ATM network. One of the methods proposed for transporting bursty LAN traffic over ATM is to set up an ATM connection upon arrival of a burst and close the connection while there is no traffic. We investigate the performance of such method of LAN interconnection. We present results, obtained from extensive simulations, to show the effects of buffer size, VC bandwidth and connection set-up time on the frame loss rate, average frame delay and VC bandwidth utilisation at the LAN-to-ATM interworking unit.

1 Introduction

ATM (Asynchronous Transfer Mode) has been accepted as the transport mechanism for future broadband ISDN (B-ISDN) networks. It allows the integration of voice, video and data on a single transport protocol. ATM is connection-oriented and is based on statistical multiplexing using small cells of fixed size. Service classes with different loss and delay priorities can be integrated onto a single network.

An ATM network is capable of being used as a backbone network to interconnect a number of LANs [1, 4, 6]. LAN interconnection is expected to be the first service offered by B-ISDN. Since data transfer through LANs are connectionless, the main difficulty in LAN interconnection is providing a connectionless service over a connection-oriented network [1]. Methods for implementing connectionless service over a connection-oriented network have been described in [2, 3]. The methods are broadly divided into a *direct* approach and an *indirect* approach. The direct approach is based on a connectionless virtual overlay network consisting of Connectionless Service Function (CLSF) nodes in the network. The indirect approach can use one of the following three methods [1]:

1. *Semi permanent end-to-end path establishment* where a semi permanent Virtual Path (VP) (or Virtual Circuit (VC)) is established between every pair of Interworking Units (IWU) of every pair of LANs. This results in a mesh of VPs (or VCs) between LANs, resulting in a scalability problem for large number of LANs. Since most data applications are bursty in nature and can not adequately describe their traffic behavior, this method results in an inefficient use of bandwidth.

2. *VC connection establishment on a frame basis* where a VC connection is established for every frame to be transmitted from a LAN to another LAN. Although the bandwidth utilization is high in this method, the processing overhead for establishing and tearing down a VC on a frame basis may result in unacceptable delays in transferring the data. This delay might result in buffer overflow and packet loss in the IWU.

3. *Switched VC connection establishment (SVCCE)* where a VC is opened when an IWU has a frame to send. However, the VC instead of being closed at the end of every frame, transmits the next frame if one is available; otherwise it is left open for some idle period of time before being closed down. The method reduces the overhead required for connection establishment and also has a reasonable utilization of the bandwidth.

The SVCCE method offers a compromise between bandwidth utilization and frame transfer delay. In this paper, we study the performance of the SVCCE method for transferring data between LANs interconnected through an ATM backbone.

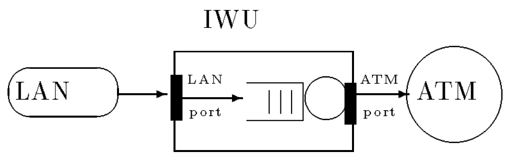

Figure 1: Model of an IWU

The performance criteria considered are the bandwidth utilisation (U) of the ATM connection, frame loss rate (P) at the IWU due to buffer flow and the average delay (D) experienced by a frame at the IWU.

The objective of this study is to investigate the effect of buffer size (K), duration of idle period (T), connection set-up time (C) and the amount of VC/VP bandwidth (B) on the performance measures described above. Extensive simulations have been carried out to study the above effects.

Results of this study will provide an insight into the required duration of the idle period, size of the buffer required at the IWUs for a given quality of service (loss and delay requirements), and the effect of connection setup time on the overall performance.

An overview of this paper is given below. In Section 2, we describe our simulation model. The results of our study are provided in Section 3. We conclude the paper in Section 4.

2 Simulation Model

2.1 Network and Traffic Model

We consider a generic LAN-to-ATM interconnection (Figure 1). The only physical characteristic of the LAN that is important for our simulation is the speed (bitrate) of the shared medium. The ATM network is used as a pipe, with *constant bitrate* negotiated during connection set up, to carry LAN frames.

We simulate one-way traffic - from LAN to ATM. Constant frame sizes are considered. The arrival process of the LAN traffic to the IWU is taken as ON-OFF. The number of frames that may arrive back-to-back during a ON period is exponentially distributed. The OFF (silence) period is also exponentially distributed with the mean adjusted to produce the desired load on the IWU. Load factor is defined as the mean bit rate of the frames arriving at the IWU to the LAN bit rate. For example, a load factor of 0.1 means on average 1 Mbps of data will arrive at the LAN port for a 10 Mbps LAN.

2.2 IWU Model

We simulate a two-port IWU - one port connected to the LAN and the other to the ATM network as shown in Figure 1. The LAN-port operates at LAN bit rate and the ATM port operates ate the rate negotiated with the ATM network at the time of connection set-up. There is a finite amount of buffer associated with each VC/VP. A single VC/VP is modeled to simulate a single LAN interconnection; the results will be valid for any number of VC/VP.

When a LAN frame arrives to the LAN port of the IWU, a check is made on the availability of buffer space. If there is enough space the frame is placed in the buffer, otherwise it is dropped. Since partial frame transmission reduces the effective bandwidth utilisation and the throughput of the higher layers [7], partial frames are not buffered in our simulation.

The ATM-port works as a server to this buffer. The service includes segmentation of a frame into cells using AAL5 and transmit the cells to the ATM network via the established VC/VP connection at the negotiated rate. The segmentation time is considered negligible and therefore, the service time of a frame is the transmission time of the frame at the negotiated rate.

If there is no connection when a frame arrives at the ATM port, a connection set-up delay (C) is experienced before the frame can be transmitted. Therefore, the service time for frames which arrive at a time when there is no previously established connection, is the transmission time plus the connection set-up time.

2.3 Simulation Parameters

The values of the parameters used in the simulation model are shown in Table 1. A LAN frame size of 1526 octet represents the maximum Ethernet frame plus some synchronisation bytes. The variable parameters include buffer size (K), length of the idle period (T), connection set-up time (C) and VC/VP bandwidth (B).

3 Results

In this Section we present the results obtained from our simulation study. All simulations stop automatically when 95% confidence interval is reached with a 0.05 relative precision for the frame loss rate at the IWU. Simulation time varied between 200 to

Simulation Parameter	Value
LAN-frame size	1526 octet
LAN-bitrate	10 Mbps
Load Factor	0.1
Mean no of frames per burst	25

Table 1: Values of simulation parameters

2000 seconds depending on the values set for various parameters.

In the following three sections we present the effects of the variable parameters — T, C, B and K — on the three performance measures — P, D, and U. When we study the effect of a particular variable, we keep the other three variables fixed; this makes the effect of the variable under study more tractable.

3.1 Frame Loss Rate (P)

We measure the frame loss rate as the ratio of frames dropped to the total frames arrived at the IWU from the LAN port.

We find that P decreases with increase in T, B and K and increases with increase in C.

P decreases linearly with increasing T. This means that longer we leave the connection open after the queue gets empty the lower the frame loss. This is also intuitively correct. However, the decrease in P is not significant for the range of T shown in Figure 2. P decreases faster for smaller buffer.

Figure 3 shows the behaviour of P against C. P increases as C increases. However, the rate of increase is not uniform. P increases rapidly up to a certain threshold value of C. After the threshold, the rate of increase slows down. This behaviour is consistent with different buffer sizes.

For a given buffer size, frame loss strongly depends on the bandwidth of the ATM connection. Figures 4 and 5 show that P decreases linearly as B and K increases.

3.2 Average Frame Delay (D)

Frame delay in the IWU is defined as the time a frame spends in the IWU from entering in the queue to leaving the IWU through the ATM port. Therefore, the delay includes the queueing delay in the buffer and the time it takes to transmit the frame on the ATM port. It does not include the time it takes to receive the frame at the LAN port. D is calculated as the average delay of all the frames buffered at the IWU and transmitted through the ATM port. A dropped frame has no delay and is not considered in the calculation of D.

Values of D against T are shown in Figure 6. D decreases linearly with increase in T. This is because less connection set up is required for larger T and each connection set-up introduces some delay.

From Figure 7 it is clear that D directly depends on C. D increases linearly as C increases.

It can be seen in Figure 8 that D decreases exponentially with an increase in B. This is because of shorter service time due to increased bandwidth of the ATM connection.

For large ATM bandwidth of the connection, not many frames are queued at the IWU and the delay experienced by a frame is not affected by the buffer size. However, D increases linearly with buffer size for low ATM bandwidth. This is illustrated in Figure 9.

3.3 VC Bandwidth Utilisation (U)

U is defined as the ratio of the time a VC spends transmitting cells to the total time the VC is open. 100% utilisation can be achieved by closing a connection as soon as buffer gets empty.

From Figure 10 we see that U decreases exponentially with increase in T. This is because there is more idle bandwidth for larger T.

Figure 11 shows the behaviour of U as a function of C for T=10ms. For large buffer U does not depend much on C. However, for smaller buffer U decreases as C increases. This is because more arriving frames are dropped for smaller buffer while the first frame is waiting for connection set-up. When the connection is established there are not many frames in the buffer to transmit, the queue gets empty quickly and the VC enters an idle waiting period. For T=0, however, U will not depend on C for any buffer size; U will be 100%.

Utilisation falls for increasing ATM bandwidth. This has a serious connotation — faster (and more expensive) ATM connection results in a waste of bandwidth. Figure 12 shows this phenomenon. Utilisation can be improved by decreasing T as suggested by Figure 10.

The behaviour of U against K is shown in Figure 13. Utilisation of VC bandwidth increases as K increases but reaches a saturation point for a sufficiently large buffer. Higher the bandwidth the faster it approaches to the saturation point.

4 Conclusion

We have studied the performance of switched VC based LAN interconnection over ATM. In particular, we have presented results to quatitatively show the effects of buffer size, idle time, connection set-up delay and VC bandwidth on the frame loss rate, average frame delay and the VC bandwidth utilisation at the interworking unit between LAN and ATM.

Our results suggest that frame loss rate is not very sensitive to the length of idle period but decreases linearly with increasing buffer size and VC bandwidth. Frame loss rate increases for longer connection set-up time.

Delay experienced by a LAN frame at the IWU decreases exponentially as VC bandwidth is increased and increases linearly for longer connection set up time. Delay is insensitive to buffer size for large VC bandwidth but increases linearly for small bandwidth. Increasing the length of idle period decreases delay because of fewer number of connection set-ups are required.

ATM bandwidth utilisation can be increased quite significantly by reducing the length of idle-period without much affecting the frame loss rate and the frame delay. This demonstrates the benefit of SVC over PVC for carrying bursty LAN interconnection traffic. Bandwidth utilisation can also be increased by providing larger buffer at the IWU. However, U approaches a saturation point for a sufficiently large buffer; increasing buffer after this point does not have significant effect on U.

References

[1] M. Murata and H. Miyahara, "LAN interworking through broadband ISDN," *IEICE Transactions Communications*, vol. E77-B, no. 3, pp. 294–305, March 1994.

[2] ITU-T Draft Recommendation I.211, *B-ISDN service aspects*. ITU-T, July 1992.

[3] ITU-T Draft Recommendation I.364, *Support of broadband connectionless data service on B-ISDN*. ITU-T, March 1993.

[4] S. Manthorpe and X. Garcia, "TCP performance over ATM based LAN interconnection services," *Interop 95 Engineers' Conference*, March 29-30, 1995.

[5] J.Y. Le Boudec, A. Meier, R. Oechsle, and H.L. Truong, "Connectionless data service in an ATM-based customer premises network," *Computers Networks and ISDN Systems*, vol. 26, no. 11, pp. 1409–1424, August 1994.

[6] Y. Ohba, M. Murata, and H. Miyahara "Analysis of Dynamic Bandwidth Control for LAN Interconnection through ATM Networks" *IEICE Transactions Communications*, vol. E76-B, no. 3, pp. 367–377, March 1994.

[7] Allyn Romanow and Sally Floyd, "Dynamics of TCP Traffic over ATM Networks" *IEEE Journal on Selected Areas of Communications*, vol. 13, no. 4, pp. 633–641, May 1995.

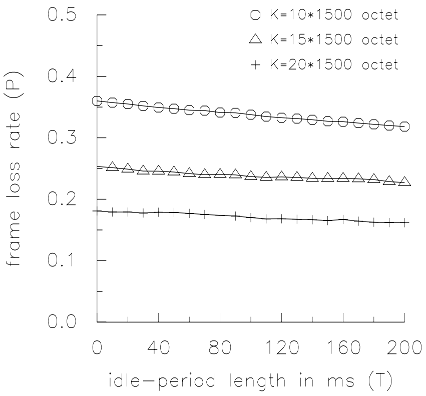

Figure 2: Frame loss rate against T; C=10ms, B=5Mbps.

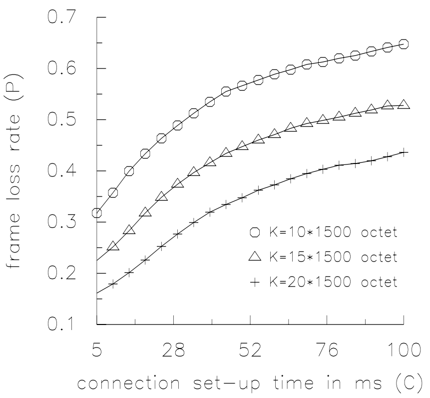

Figure 3: Frame loss rate against C; T=10ms, B=5Mbps.

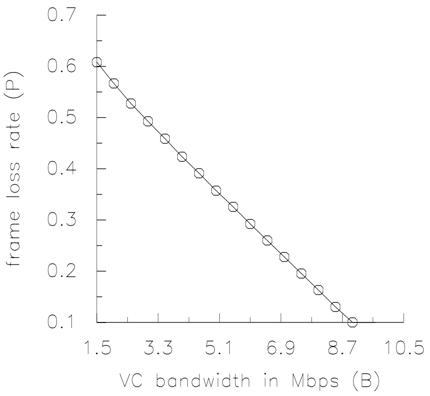

Figure 4: Frame loss rate against B; K=10*1500oct, C=10ms, T=10ms.

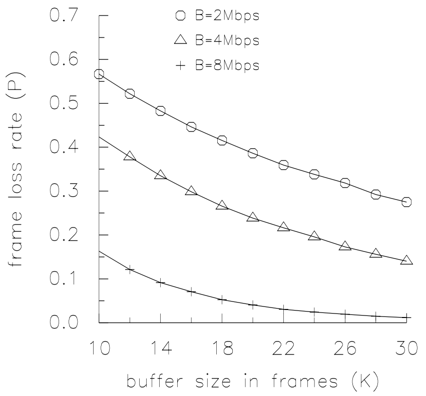

Figure 5: Frame loss rate against K; C=10ms, T=10ms.

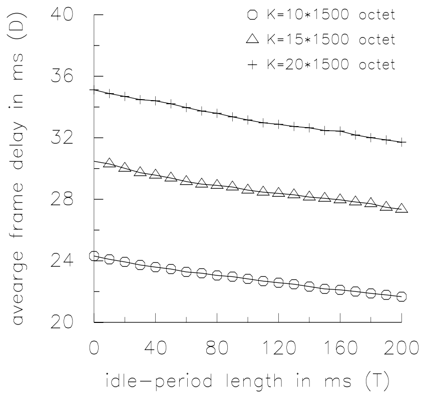

Figure 6: Frame delay against T; C=10ms, B=5Mbps.

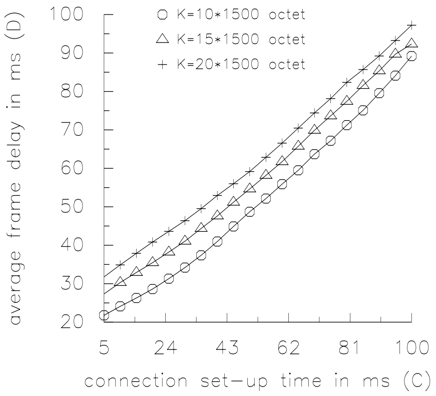

Figure 7: Frame delay against C; T=10ms, B=5Mbps.

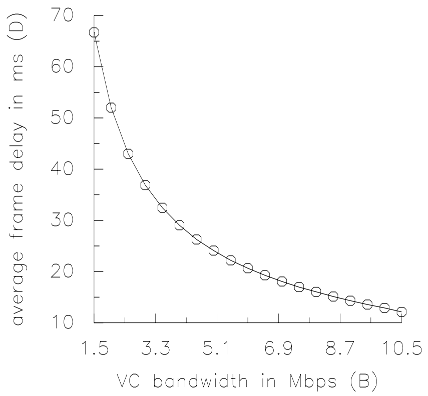

Figure 8: Frame delay against B; K=10*1500oct, C=10ms, T=10ms.

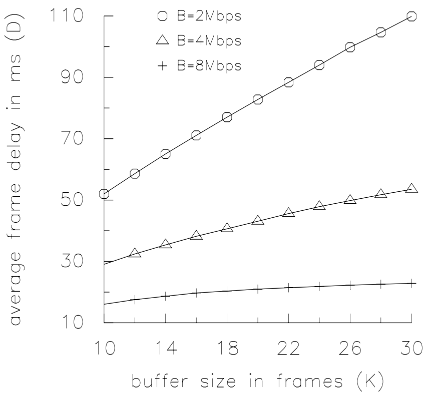

Figure 9: Frame delay against K; C=10ms, T=10ms.

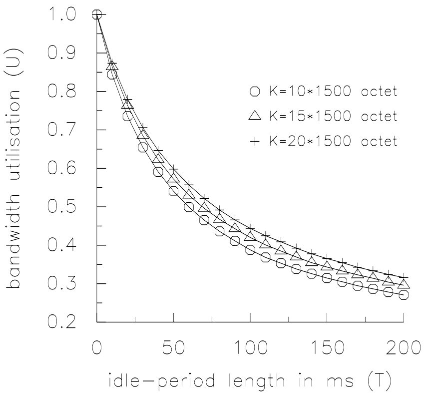

Figure 10: Bandwidth utilisation against T; C=10ms, B=5Mbps.

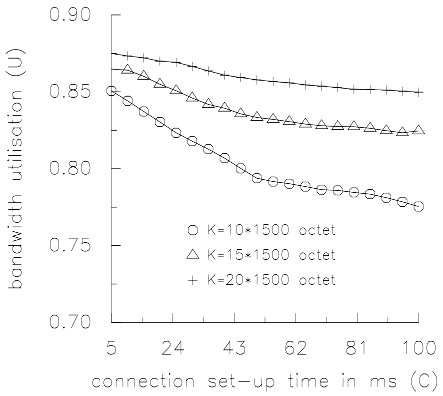

Figure 11: Bandwidth utilisation against C; T=10ms, B=5Mbps.

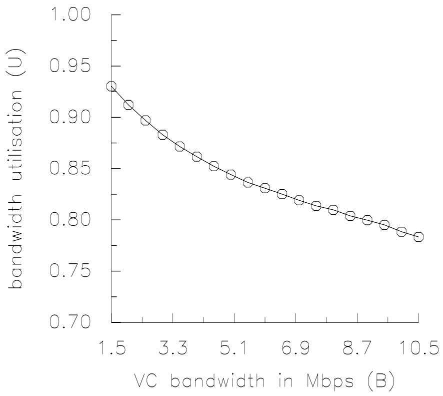

Figure 12: Bandwidth utilisation against B; K=10*1500oct, C=10ms, T=10ms.

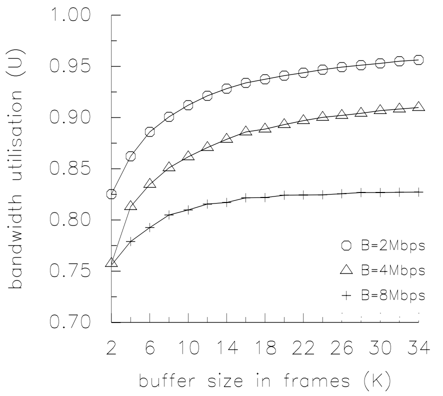

Figure 13: Bandwidth utilisation against K; C=10ms, T=10ms.