ACSC'96

Nineteenth Australasian Computer Science Conference
Proceedings

Melbourne, Australia
31 January – 2 February, 1996

Kotagiri Ramamohanarao (Editor)

ISSN 0157–3055

Australian Computer Science Communications
Volume 18, Number 1

RMIT

THE UNIVERSITY OF MELBOURNE

Effect of Frame Loss on the Holding Time of a Virtual Circuit in ATM LAN Interconnection -

Mohammed Atiquzzaman
Dept. of Elec. & Comp. Sys. Eng.
Monash University, Clayton
Clayton, Victoria 3168, Australia
atiq@eng.monash.edu.au

Mahbub Hassan
School of Comp. & Info. Tech.
Monash University, Gippsland
Churchill, Victoria 3842, Australia
mahbubh@gscit-1.fcit.monash.edu.au

Abstract

We analyse the cost performance of LAN interconnection over ATM using switched VC. Frame loss due to buffer overflow during connection setup has been taken into consideration in the proposed cost function. We investigate the effect of the length of the holding time of an idle VC on the total cost per frame transmitted. Our simulation results suggest that there is no optimum holding time for minimising the total cost. Increasing the holding time either increases or decreases the total cost depending on the cost associated with a connection setup and frame loss. We also show that for a given cost associated with a frame loss the practical range of buffer sizes is bounded by a minimum and a maximum value. Within the practical range, larger buffer reduces total cost. Increasing buffer beyond the maximum size does not have any impact on the total cost. For buffers below the minimum size, a connection should never be closed; such buffer sizes have, therefore, no benefit at all.

Keywords
LAN interconnection, ATM network, virtual circuit.

1 Introduction

ATM 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, 3]. 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]. Several methods for implementing connectionless service over a connection-oriented network have been described in [4]. The methods are broadly divided into a *direct* approach and an *indirect* approach. The direct approach [5] 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 [4, 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 frame loss in the IWU.

3. *Switched VC connection establishment (SVCCE)* where a VC is opened when an IWU has a frame to send. However,

Proceedings of the 19th Australasian Computer Science Conference, Melbourne, Australia, January 31–February 2 1996.

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. The major issue with SVCCE is how long to hold an idle VC open. Too short holding time will cause more loss and opening cost while leaving it open for too long reduces bandwidth utilisation. The objective of this study is to investigate the effect of frame loss due to buffer overflow on adjusting the holding time of an idle VC. Extensive simulations have been carried out to study the dependency among the buffer size, frame loss rate, connection opening rate and the connection set-up delay.

Saran and Keshav [6] have analysed a few holding policies based on two different pricing schemes. The first pricing scheme depends on the total number of connection provided to a site, a per frame charge, and a call setup charge. The second pricing scheme, which is more realistic for guaranteed bandwidth ATM connection, is based on a call setup charge, a holding time charge and the cost associated with loss of user utility due to a call setup delay. The cost due to loss of user utility is provided by the system manager.

Lund, Phillips and Reingold [8] have extended the above work by proposing an adaptive VC holding policy that adapts to the characteristics of the IP traffic. They have used the second pricing scheme mentioned above.

The authors of the above papers have not considered frame loss in their analysis of holding policies. In [6], there is no frame loss at the IWU implying a sufficiently large buffer at the IWU. In contrast, Lund, et.al. [8] assume no queueing of frames at the ATM interface, implying a direct host interface to the ATM network.

Murata, et.al. [1] have considered performance and cost comparison for LAN interconnection over switched ISDN service which was suggested as applicable to the SVCCE scenario as well. However their model explicitly assumed an infinite buffer at the IWU and hence no frame loss.

Manthorpe [9, 10] has stressed frame loss due to a finite buffer at the IWU. However, the study was limited to a direct approach as mentioned above where the assumption was to leave the connection open all the time. Hence, their study is not directly related to our study.

IWUs have a finite amount of buffer space and can result in frame loss in the case of the buffer becoming full. Frame losses in a TCP/IP based network result in retransmission of the lost frame and a consequent drop in the goodput. It is therefore important to take into account the effect of frame loss in such studies, and also to incorporate it in the cost function of a VC.

In this study, we consider an IWU which has a finite amount of buffer and therefore, could result in frame loss in the IWU. Our model is closer to reality than previous studies which do not account for frame losses. We also propose a new cost function to take into account the frame loss due to buffer overflow. We investigate the behaviour of the cost function by changing the length of idle time. Based on our simulation results, we recommend a practical range of buffer size for a given cost associated with a lost frame.

Results of this study will help in determining an optimal holding time policy for a given buffer size and arrival characteristic. It will also help the system manager to decide the required amount of buffer size in an IWU.

The contributions of this paper can be summarized as follows:

- Our study assumes a finite amount of buffer space (which is more realistic than the assumption of an infinite buffer) at the IWU.
- We have developed a cost function which takes into account frame loss at the IWU.
- We recommend a practical range of buffer size for minimising total cost; this information is very valuable for buffer dimensioning in the IWU.

An overview of this paper is as follows. In Section 2, we describe the pricing scheme and the cost function proposed in this paper. Our simulation model is illustrated in Section 3 and the results are discussed in Section 4. Finally, we conclude the paper in Section 5.

2 Cost Function

We associate cost with three parameters — VC holding time, connection opening and frame loss. If too many frames are lost for each opening, a holding policy should minimise the number of opening to reduce frame loss and the total cost. We take the

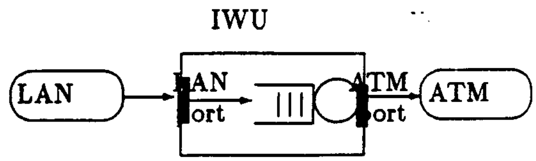

Figure 1: Model of an IWU

cost of one millisecond of holding time as the unit cost. The total cost (C) is defined as:

$$C = N \times C_s + l \times C_f + C_h$$

where, N is number of opening, C_s is cost associated with a opening, l is number of frames lost, C_f is cost associated with a lost frame and C_h is the holding cost.

3 Simulation Model

3.1 Network and Traffic Model

We consider a generic LAN-to-ATM interconnection. 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 a bitrate negotiated during connection setup, 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 ATM is taken as on-off. During the ON period, frames arrive back-to-back without any idle period between the frames. The number of frames that may arrive back-to-back during a ON period is exponentially distributed. The mean number of frames that arrive during a ON period is taken as 25 for this simulation study.

The OFF (silence) period is also exponentially distributed with the mean idle time adjusted to produce the desired load on the VC/VP at the IWU. Load factor is defined as the fraction of the LAN bitrate that arrive at the ATM port for transmission over the ATM network. For example, a load factor of 0.1 means on average 1 Mbps of data will arrive at the ATM port for a 10 Mbps LAN.

3.2 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 the speed of LAN and the ATM port operates at the peak rate that is negotiated with the ATM network at the time of connection opening. Here we assume peak rate allocation because a VC is opened only when there is data enqueued for transmission.

At the ATM port there can be multiple VC/VP setup to interconnect to multiple sites. There is finite buffers 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 buffer occupancy. If there is enough space the frame is placed in the buffer, otherwise it is dropped.

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 the buffer gets empty, the VC becomes idle. If no frame arrives after certain idle period, the VC is closed. Therefore, when a frame arrives it may find the VC closed. If there is no connection when a frame arrives at the head of the queue, a connection set up delay is experienced before the frame can be transmitted. Therefore, the service time for some frames is the transmission time plus the connection set up time.

The segmentation process is done at the ATM port instead of at the LAN-port in order to avoid partial frame transmission. Partial frame transmission reduces the effective bandwidth utilisation and the throughput of the higher layers [2]

3.3 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 maximum Ethernet frame plus some synchronisation bytes. VC bandwidth is chosen as 11.2 Mbps so that the time to transmit all the ATM cells, generated for a 1500-octet (32 cells) frame, on the ATM port is about the same as the time to receive a LAN frame on the LAN port. Queueing does not happen once the VC has been established; it only builds up during connection set up. The variable parameters include buffer size, traffic load to the VC/VP, length of the idle time and connection set up time.

4 Results

With SVCCE, a VC is closed after a certain idle period. From now on, we will refer to this

Parameter	Value
LAN-frame size	1526 octet
LAN-bitrate	10 Mbps
VC bandwidth	11.2 Mbps
Mean no of frames per burst	25
Traffic load factor	0.1

Table 1: Values of simulation parameters

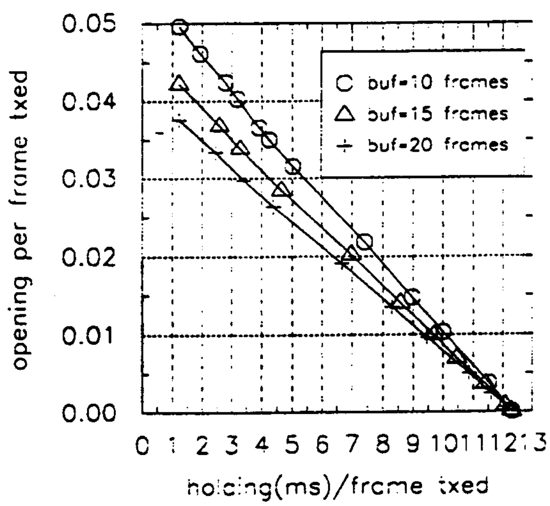

Figure 2: Number of connection opening for increasing holding time

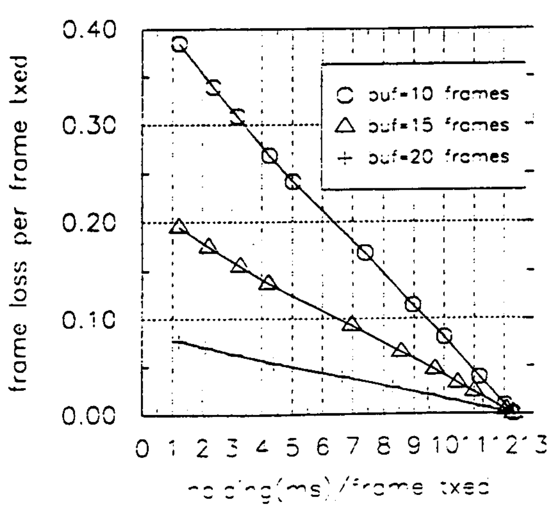

Figure 3: Frame loss for increasing holding time

idle period as t_i. The goal is to investigate a value for t_i which will minimise the total cost C.

Long t_i means fewer number of openings and lower frame loss at the price of increased holding time. For very long t_i, it becomes a permanent connection as it may not be necessary to close the connection at all. For short t_i the holding time is reduced at the expense of more openings and higher frame loss. We refer to costs for frame loss and connection set up as Type-A cost and costs for holding the VC as Type-B cost. From the above explanation, increasing t_i will increase the Type-B cost but reduce Type-A cost.

Since the goal of a holding policy is to minimise the total cost (Type-A cost plus Type-B cost) we investigate the effect of varying t_i on the total cost for different C_s and C_f. Intuitively, we thought there would be an optimum value for t_i which will minimise the total cost. However, extensive simulations showed that there is *no such optimum value for* t_i.

We ran simulations for reasonably long time; the simulation was terminated automatically either when the 95% confidence level for 0.05 relative precision was reached for the frame loss rate at the IWU or 5000 seconds of simulation time elapsed. The numbers of connection set up and frame loss against holding time per frame transmitted are shown in Figures 2 and 3. The values in these figures are obtained by varying t_i.

The straight lines suggest that Type-A cost is linearly dependent on Type-B when t_i is varied. This linear dependence implies that we have to minimize either Type-A or Type-B cost, depending on the slope of the Type-A vs Type-B line (m), in order to minimize the total cost. This slope will depend on C_s, C_f and C_h for a given buffer size and traffic pattern. The significance of m is explained below.

For $m > 1$, a unit increase in Type-B cost reduces more than a unit of Type-A cost. Therefore, the minimum total cost is achieved for zero Type-A cost meaning no frame loss and no connection set up. This is achieved by leaving the VC open all the time — no buffer is necessary. The total cost incurred in this situation is the maximum cost. Maximum cost per frame transmitted is the average VC holding time per frame transmitted and can be calculated as:

$$C_{max} = \frac{t_a}{\rho}$$

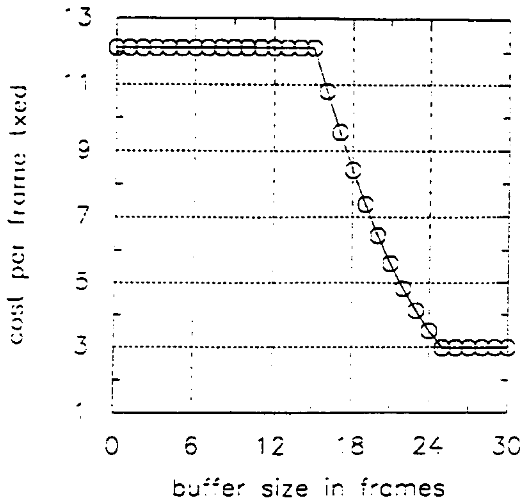

Figure 4: Total cost for different buffer sizes; Cs=50, Cf=50, setup:30ms

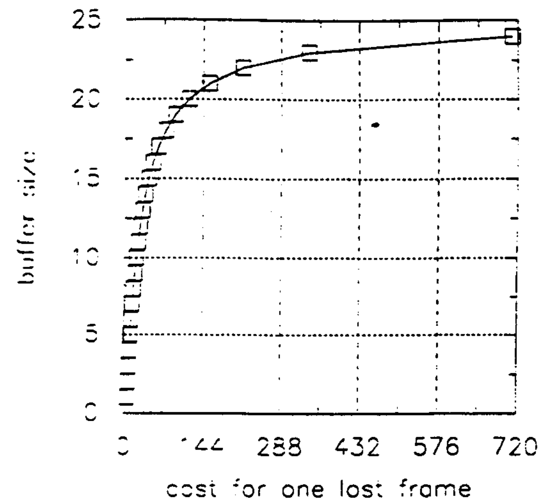

Figure 5: Practical range of buffer sizes; C_s=50,t_d=30ms

where, t_a is the time it takes to transmit a frame on the ATM port and ρ $(0 < \rho < 1)$is the load factor.

On the other hand, for $m < 1$, the minimum total cost means minimum Type-B cost which is achieved for $t_i = 0$. This implies closing the VC as soon as buffer gets empty. In this case total cost will be less than C_{max}.

For $m = 1$, it doesn't matter if the VC is left open permanently or closed as soon as buffer gets empty; the total cost remains the same. From the above three situations, $m > 1$, $m < 1$ and $m = 1$, we see that there is no optimum value for t_i to minimise the total cost.

We now investigate the effect of buffer size on the total cost. For a given buffer size, m depends on C_s and C_f. For large C_s and C_f, $m > 1$ and closing a VC can not be justified. The VC should be left open all the time; otherwise, the cost incurred will be greater than C_{max}.

For a given C_s and C_f, if $m > 1$ then buffer size should be increased. Increasing the buffer size decreases frame loss and connection set up and, therefore, reduces m. The minimum buffer size (b_{min}) for achieving total cost less than C_{max} is the minimum size for which m is less than 1.

There is also a maximum buffer size — b_{max}. No cost saving is achieved beyond this size. This maximum size is the minimum buffer size which will hold all the frames that may arrive during a connection set up and can be calculated as:

$$b_{max} = \lceil \frac{t_d}{t_l} \rceil$$

where, t_d is the connection set up delay and t_l is the time it takes for a frame to be received on the LAN port.

Figure 4 shows the behaviour of the total cost curve against buffer size for a load factor of 0.1 (1Mbps). For buffer sizes less than 16 $m > 1$ and the total cost is minimised for $C_{max} = \frac{1.2}{0.1}$. For buffer sizes 16 or more $m < 1$ and total cost is reduced by setting $t_i = 0$. The curve becomes flat for buffer sizes greater than 25 ($b_{max} = \lceil \frac{30}{1.2} \rceil = 25$). The practical range of buffer sizes in Figure 4 is 16-25.

The practical range of buffer sizes is, therefore, bounded by b_{min} and b_{max}. From the simulation results for different buffer sizes we calculate C_f for which $m = 1$ and plot the buffer size against this C_f. Minimum buffer sizes for different C_f are shown in Figure 5; the maximum buffer size is 25 frames.

5 Conclusion

We analysed the performance of the SVCCE method for LAN interconnection over ATM. Frame loss due to buffer overflow has been taken into consideration in the proposed cost function. Using simulation, we have investigated the effect of idle VC holding time on the total cost for a on-off traffic model. Our results suggest that *no optimum idle holding time exists* for minimising the total cost. In order to minimise the cost, either the VC should be closed as soon as the buffer gets empty or the VC should be left open permanently depending on the cost associated with a VC opening, frame loss or a unit of holding time.

We also show that for a given cost associated with a frame loss the practical range of buffer sizes are bounded by a minimum and a maximum value. Within the practical range, increasing the buffer size reduces the total cost. Increasing buffer beyond the maximum size does not reduce cost. For buffers below the minimum size, a connection should never be closed; such buffer sizes have, therefore, no benefit at all.

We show the practical range of buffers for different costs associated with a frame loss for a given connection set up delay and traffic load.

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] 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.

[3] F. Ball and D. Hutchison, "Traffic control in an ATM LAN," *IFIP Second Workshop on Performance Modelling and Evaluation of ATM Networks*, Bradford, UK, pp. 27/1–12, 4-7 July 1994.

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

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

[6] H. Saran and S. Keshav, "An empirical evaluation of virtual circuit holding times in IP-over-ATM networks," *IEEE INFOCOM '94: Conference on Computer Communications*, Toronto, Canada, pp. 1132–1140, June 12-16, 1994.

[7] V. Jacobson, "Congestion avoidance and control," *ACM SigComm*, California, pp. 314–329, August 16-19, 1988.

[8] C. Lund, S. Phillips, and N. Reingold, "Adaptive holding policies for IP over ATM networks," *IEEE INFOCOM '95: Conference on Computer Communications*, Massachusetts, pp. 80–87, April 2-6, 1995.

[9] S. Manthorpe and X. Garcia, "Shaper dimensioning for MAN to ATM based WAN," *Interop 94 Engineers' Conference*, Las Vegas, May 1994.

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