On the Performance of Bandwidth Allocation Strategies for Interconnecting ATM and Connectionless Networks

Edward Chan, Victor C. S. Lee
Department of Computer Science, City University of Hong Kong
e-mail: csedchan@cityu.edu.hk

Jim M. Ng
Information Communications Institute of Singapore
e-mail: jimmee@icis.gov.sg

Abstract

Using ATM networks as the switching fabric for interconnecting LANs and MANs means that a strategy for bandwidth allocation must be developed to map connectionless traffic in the LAN/MAN to the ATM network which is essentially connection-oriented. This paper presents a bandwidth allocation algorithm based on bandwidth advertising and burst drop when overflow occurs. The performance is evaluated by simulation and shown to reduce Burst Loss Rate significantly.

Keywords High Speed Networks, ATM, Internetworks

1. Introduction

Asynchronous transfer mode (ATM) networks have been widely touted as the switching fabric that would interconnect LANs and MANs. This would allow LAN and MAN services to be extended geographically. However there is a bandwidth allocation problem at the LAN/MAN and ATM interconnection point since most LAN/MAN traffic is connectionless and ATM network is basically connection-oriented. Some means must be derived to map the bursty connectionless traffic on to the connections provided by the ATM switching fabric.

2. Previous Work

Various approaches have been proposed to deal with the bandwidth allocation problem in the interconnection of LAN/MAN and ATM. One scheme is to use semi-permanent Virtual Paths (VP) between all InterWorking Unit (IWU) pairs to form a full mesh. However, this scheme suffers from a number of problems. As the size of the network increases, the number of VPs grows rapidly. Since each ATM switch needs to maintain a VP table, the overhead from large VP tables may be unacceptable in a large ATM network. Bandwidth allocation of the VPs is also a problem in this approach. If fixed maximum bandwidth is allocated to each VP, bandwidth waste will be unacceptably high. To alleviate this problem, a number of approaches have been suggested. Two common ones are dynamic bandwidth renegotiation and bandwidth advertising.

In the dynamic bandwidth renegotiation scheme, a small bandwidth is allocated to each VP initially. When traffic exceeds the allocated bandwidth, the reserved bandwidth will increase through bandwidth negotiation. Similarly, the bandwidth of the VP will decrease as the traffic volume decreases. The main purpose of this scheme is to fully utilize the available bandwidth.

Mongiovi [1] proposed assigning a 'peak' bandwidth to each VP and a number of buffers in the IWU. When the burst traffic exceeds the reserved bandwidth, the data will overflow into the buffer. As the queue in the buffer exceeds a certain threshold level, the bandwidth renegotiation mechanism will be triggered to reserve more bandwidth. On the other hand, the allocated bandwidth will be decreased as the queue in the buffer drops below

the lower threshold level. The reserved bandwidth in the idle state is generally set to a small value so that renegotiation occurs frequently. This can lead to substantial renegotiation overhead. There are also a number of problems with this approach. A very large buffer is usually needed to accommodate bursty traffic. Bandwidth negotiation takes time, and consume resources at the ATM switch. Finally, determining the threshold level for bandwidth increase and decrease is very critical and highly dependent on the traffic pattern, rendering accurate estimation quite difficult. An alternative is to use a fast reservation scheme. While this scheme has its advantages, the IWU must know the bandwidth requirement of the burst traffic in advance in order to make reservation, which again is a difficult task [3].

It is possible to allow traffic to enter into the ATM network instead of staying in the buffer if there is enough residual bandwidth in the path. Variations of this approach has been proposed by Gerla [2,3] and Crocetti [4,5]. The advantage of this approach is that overall delay is lowered since bandwidth reservation and renegotiation time is reduced. However, a Bandwidth Advertising mechanism must be set up to provide up-to-date residual bandwidth information to each IWU. The effectiveness of this scheme depends heavily on the bandwidth advertising strategy. If updates are not handled frequently enough, it is possible for the network to be flooded with excess traffic.

3. Proposed Approach

In this section, a modification to Gerla's Bandwidth Advertising (BA) scheme will be proposed. This paper will use his network model to facilitate comparison. Detailed models of the components in the network will be examined before the presentation of simulation results in the next section.

3.1 Overall Network Model for BA Scheme

The overall network topology is shown in Fig. 1. For the purpose of this study, we will only focus on a single 150 Mbps ATM trunk line, and the system is simplified into that depicted in Fig. 2. The whole system can be divided in 5 components: the traffic source, gateway for foreground traffic, background traffic multiplexer, ATM switch and data sink.

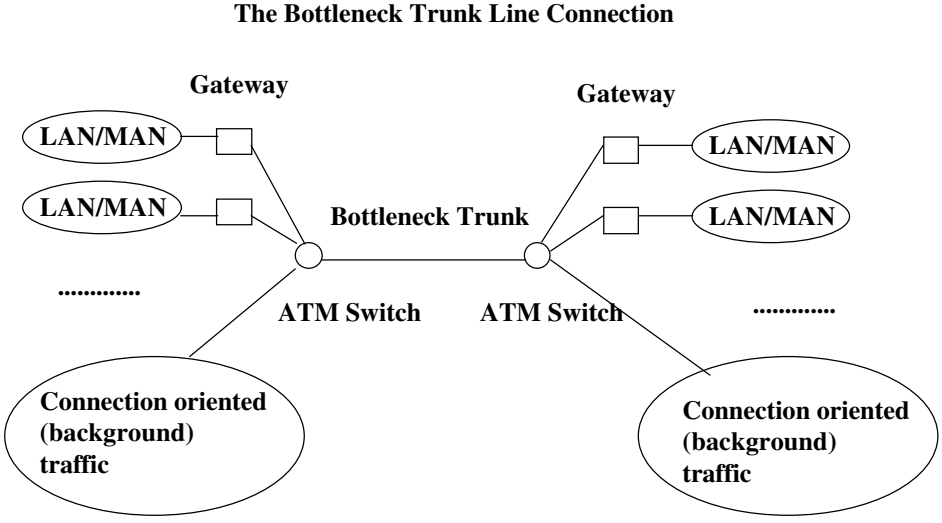

Fig. 1 Overall Network Configuration

Both connectionless and connection-oriented traffic are included in the model. Since data in LANs and MANs are predominately connectionless in nature, such traffic is considered as foreground traffic and connection-oriented traffic modeled as background traffic. The number of foreground VPs is varied to represent different

traffic loading levels, and 0.1 Mbps is allocated for each foreground VP. There are 3 background traffic VPs for the ATM switch. Ten bursty traffic sources are connected to each foreground and background VP.

Fig. 2 also shows the ATM switch with a Bandwidth Advertising Center. The bandwidth advertising interval is set to a certain number of cell transmission times, where one cell transmission time is the time to transmit an ATM cell on a 150 Mbps trunk line. Abusive cells i.e. cells that exceeded the allocated 0.1 Mbps for each VP are marked, and only sent if there is excess advertised bandwidth.

The ATM switch uses threshold buffer management. The cell buffer size is 50 in the switch, with threshold level set to 40. When the threshold level is exceeded, incoming marked cells will be dropped.

Bandwidth Advertising with Cell DropDetection (BA-Drop) Network Architecture

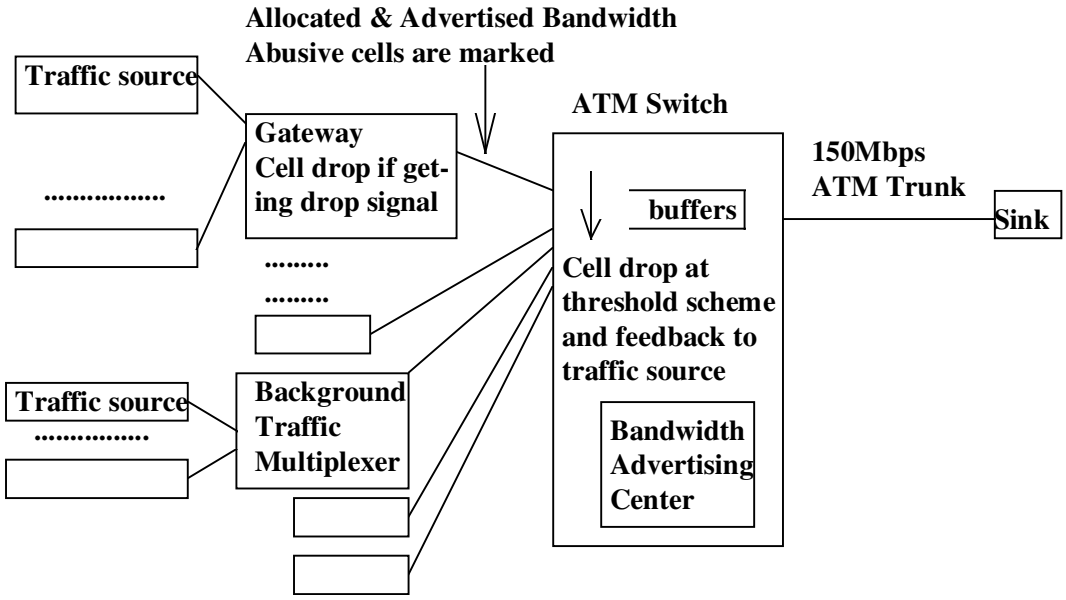

Fig. 2 Simplified Network Model Showing Bandwidth Advertising with Cell Drop Detection

3.2 The BA-Drop Scheme

While the BA scheme has been shown to be effective, a number of enhancements are possible. It has been pointed out by researchers that policing functions should consider higher levels of abstraction in discarding excess traffic. [6] In many applications, discarding a certain number of cells renders the entire data burst useless to the end-user. This is because a data burst may represent an integral unit such as a graphics object. In this case, it may be preferable to discard the entire burst.

The idea that Burst Loss Rate (BLR) is more significant than Cell Loss Rate (CLR) is the key idea in our proposed scheme. We propose to make a very simple modification to the BA scheme. In the ATM switch, a feedback signal is sent back towards the traffic source whenever cells are being discarded. The gateways will then drop all the following cells in the 'violating burst'. Since it is not possible to identify a particular burst, the current burst in the gateway will be considered as belonging to the same burst. Cells in the same burst that precedes cell drop detection will of course pass through to the ATM trunk line.

The network model of Bandwidth Advertising with Cell Drop Detection (BA-Drop) is shown in Fig. 2. Note the gateway is used for foreground traffic only. The gateway concentrates bursty traffic to the ATM switch. It consists of two portions, the gateway multiplexer and gateway processor. The multiplexer concentrates incoming traffic, whereas the gateway processor implements the required policing function. The structure of the gateway is shown in Fig. 3.

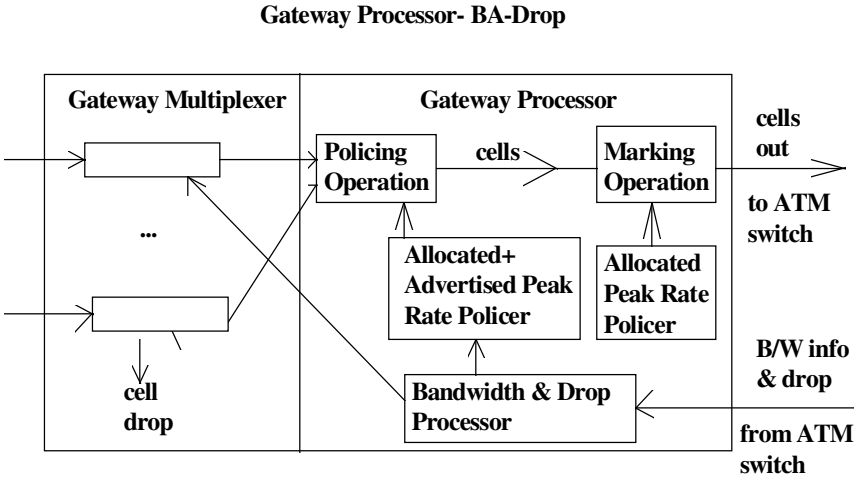

Fig. 3 Gateway Structure: BA with Cell-Drop

The structure of the ATM switch is depicted in Fig. 4. The switch consists of four portions in our model: the ATM switch multiplexer, buffer, bandwidth advertising center and drop signal handler. Note that the ATM switch multiplexer contains an access control function for discarding marked cells from foreground VPs. Background VPs, on the other hand, are not policed. The bandwidth advertising center is assumed to use dedicated network channels to convey control signal in real-time. The bandwidth advertising method described and evaluated in Crocetti's original paper [5] is used in our simulation.

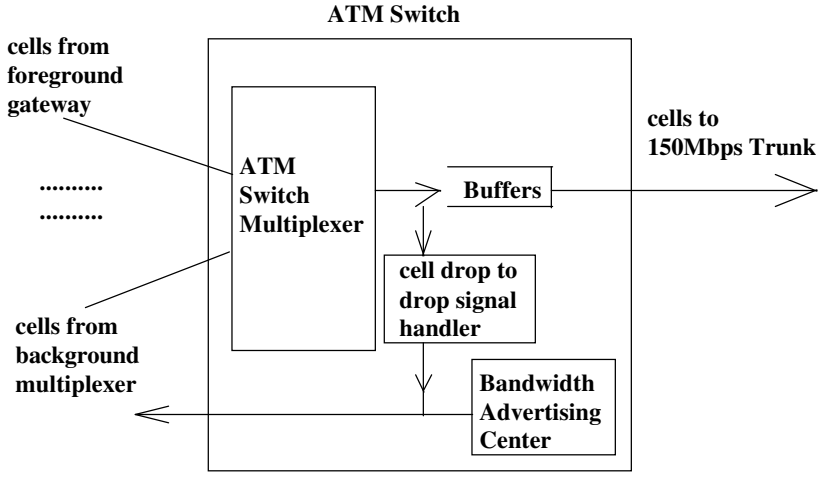

Fig. 4 ATM Switch Structure

4. Simulation Model

The performance of the BA-Drop scheme is studied and compared with the original BA scheme using the network model and configuration described in the previous section. The simulation was done using OPNET, a proprietary graphical simulation package.

Two types of traffic (connectionless and connection-oriented) are modeled in this study. Both types of traffic sources are represented by an on/off model. Active period, in which sources transmit at peak rate, is interleaved with idle period, in which sources remain silent. Active and silent periods are exponentially distributed. The on-off type bursty source is shown in the diagram below (where the 'pulses' are cells) and characterized by the following parameters:

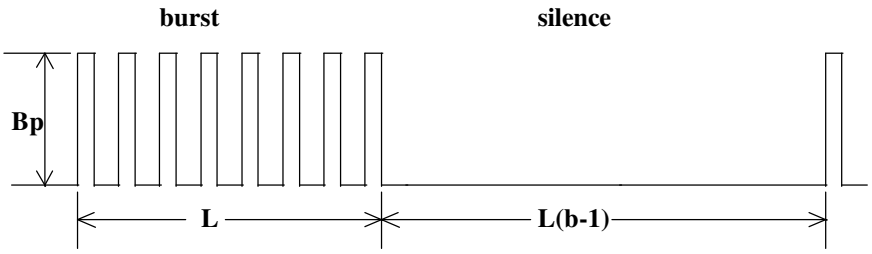

Bp : peak bit rate (number of bits / time duration of burst)
b : burstiness - defined as the ratio of peak bit rate Bp to mean bit rate
L : mean burst duration - in term of ATM cells. (53-byte long)

In our simulation, it is assumed that an appropriate mechanism is available to determine the boundaries of a burst. Various mechanisms have been proposed by researchers that map upper layer entities such as video frames or TCP packets to cell level 'bursts'. One researcher, for instance, uses the ATM-layer user-to-user (AUU) parameter in the ATM cell header as an end-of-packet delimiter in an attempt to map TCP packets to cells. Such schemes would allow the switch to drop cells of a violating burst (i.e. TCP packet) until the AUU parameter is set in a cell [8]. However, which mechanism is used for implementing burst delimitation is not the focus of this paper: our main objective is to investigate, given that such mechanisms are available, the effectiveness of a simple mechanism that would improve the BLR for the well known BA scheme.

Each gateway pair maintains a 'thin' VP, which supports 10 bursty traffic sources. In the foreground traffic, the parameters are assumed as Bp = 10 Mbps, b = 50 and L = 200. Each thin VP is rated at 0.1 Mbps.

There are three conventional VPs, each of which accommodates ten background traffic sources. The background connection-oriented traffic is carried by Virtual Connections, which are statistically multiplexed with each VP. The bandwidth required for each VP to guarantee Grade of Service (GOS), in terms of cell loss rate, for a given set of input traffic parameters is calculated using the Uniform Arrival and Service (UAS) model [7] for the finite buffer case. The parameters are assumed to be Bp = 5 Mbps, b = 10, and L = 100.

As for background sources, each background VP is allocated a fixed bandwidth. Based on the parameters specified and the UAS model, 26.236 Mbps are needed for each background VP to guarantee a GOS (Grade Of Service) of 10^{-5} with an ATM buffer of 50 cells. These values are basically the same as that in Crocetti's original study [5] to facilitate comparison.

5. Simulation Results

5.1 BA-Drop Scheme

The performance of the BA-Drop scheme and the original BA scheme are shown in Fig. 5(a) and 5(b) with the bandwidth advertising interval set to 1000 cell transmission times. It is found that while overall CLR is higher for the BA-Drop scheme, BLR is considerably reduced, especially at heavy load. To look into how the BA-Drop scheme achieves these results, the CLRs in both the gateway and the switch are collected and are shown in Fig. 5(c) and 5(d) respectively. In Fig. 5(c), since there is no cell drop mechanism in the original BA scheme, the CLR in the gateway is zero. On the other hand, when BA-Drop scheme is used, the CLR in the gateway increases as the number of VPs increases. However, this is not really a problem since nearly all of the dropped cells belong to bursts whose cells are already being dropped in the switch.

The advantage of dropping the cells in the gateway using the BA-Drop scheme is shown clearly in Fig. 5(d). Much fewer cells are dropped in the switch in the BA-Drop scheme than the original BA scheme since excess cells are already dropped in the gateway. Dropping cells in the gateway is much preferred to dropping them in the switch, particularly at the burst level since the cells dropped in a switch typically belong to different bursts due to multiplexing. As a result, though the total CLR for the BA-Drop scheme is higher than that for the

original BA scheme, the total BLR is lower for the BA-Drop scheme. Thus the objective in devising the scheme, namely the reduction of burst loss, has been achieved. Another salutary effect of the scheme is shown in Fig. 5(e), which shows that the BA-Drop scheme can help to reduce cell delay when the load is heavy.

5.2 Bandwidth Advertising Interval

To investigate the performance of these schemes in depth, we varied some parameters in the schemes. One of the parameters is the bandwidth advertising interval (BAI). In Fig. 6(a)-(d), the effect of varying the BAI is shown. In Fig. 6(a) and 6(b), when the load is light (no. of VPs=30), both the total CLR and BLR increase when the BAI increases. This phenomenon is easy to understand. When the BAI increases, the information about the congestion level of the switch in the gateways is not updated frequently enough. This means that when the switch becomes congested, the gateways will not reduce their token generation rate and slow down the traffic volume passing through them until they receive the next bandwidth advertisements which may be too late to remedy the situation. The argument can be verified by the results shown in Fig. 6(d). As the BAI increases, the CLR in the switch increases sharply. The congestion is in turn propagated back to the gateways through the cell-drop mechanism; hence the gateway shows an increase (although less pronounced than at the switch) in CLR. This effect is shown in Fig. 6(c).

A similar trend is observed when the traffic loading is heavy (no. of VPs = 35 or 40). However, when the number of VPs is large, an interesting effect is noted when BAI is small. One would expect that the CLR and BLR would decrease as BAI becomes quite small (about 1000 to 2000 cell transmission times), since frequent advertisements help to keep the gateways up-to-date. In reality, as Fig. 6(a)-(d) shows, both CLR and BLR actually increase for small values of BAI. The reason is that although the gateways can respond to the congestion level of the switch promptly for small BAI, queue oscillations will occur in the switch. All the traffic flow from the gateways will shrink simultaneously when a short burst arrives at the switch causing it to mistakenly advertise a drastically-reduced bandwidth. Similarly, traffic will surge into the switch when there is a short silence. The basic problem is that at such a small BAI, the advertised bandwidth is simply not accurate enough. Moreover, with every broadcast of bandwidth information there is a period of instability, and with a small BAI the system does not stabilize. As a result, the bandwidth efficiency will be reduced and the loss ratio will increase. This observation is further demonstrated by the CLR in the switch, shown in Fig. 6(d).

Hence, we find that optimum BAI values exist for different values of VP. When the traffic loading is light, frequent bandwidth advertising causes no harm. When the traffic loading increases, the optimum BAI shifts to a larger value. From Fig. 6(a) and 6(b), we can see that as the number of VPs increases from 30 to 50, the optimum BAI value which leads to minimum loss ratios shifts to about 4,000 cell transmission times. This value will be used in the experiments in the next few sections when the effects of other parameters on the BA-Drop scheme is investigated.

The performance of the original BA scheme is also shown in these figures for comparison. The same observations and explanations for the BA-Drop scheme can be made for the BA scheme. Consistent with the results shown in Fig. 5(a)-(d), as the load increases, the difference between the two schemes becomes large. It should be noted that while both schemes experience higher cell loss in the gateway (see Fig. 6(c)) when the BAI is large, they are caused by different mechanisms. Thus the loss in BA-Drop scheme is due to the cell drop signal, while that for the BA scheme is simply due to the buffer overflow in the gateway when a outdated bandwidth information has been used in the policing function for a certain period of time.

5.3 Burstiness

The performance of the BA-Drop scheme may be affected by factors such as the burstiness of the traffic sources and the propagation delay of the cell drop signal sent from the ATM switch back to the source gateway. Fig. 7(a) shows the total BLR when the burstiness of the traffic sources is halved, which means that the traffic volume is doubled. The results show that the performance of the BA-Drop scheme is still superior to that of the original BA scheme. In fact, the performance of the original BA scheme degraded much more sharply.

In this study, the current burst in the gateway will be assumed to be the violating burst. However, if the burstiness of the traffic source is reduced, the inter-burst duration decreases as well. Thus the chance that the current burst in the gateway is not the violating burst will increase. As a result, some cells not belonging to the violating burst may be falsely dropped. However, our results show that this is rare. It can be seen in Fig. 7(b) that the false CLR remains very low (< 0.2%) even though the burstiness was reduced by 75%.

5.4 Propagation Delay of the Backward Cell-Drop Signal

Another factor which may degrade the performance of the BA-Drop scheme is the propagation delay of the backward cell-drop signal. Two problems may arise if the cell-drop signal arrives late at the source gateway. Firstly, a late signal means that most of the cells of the violating burst have already left the gateway to the switch. As a result, the objective of reducing the congestion caused by the cells of the violating burst cannot be achieved. Even worse, as mentioned before, if the cell drop signal arrives when the next burst is in the gateway then part or all of this burst will be falsely dropped. Fig. 8(a) shows the total BLR as the propagation delay is increased from a negligible value to 0.01 sec, which corresponds to a distance between the switch and the gateway of up to about 2000 km. It can be seen that the impact of this factor becomes prominent when the number of virtual path exceeds 30, and that the effect of loading seems to be more significant than the propagation delay itself.

6. Conclusion

In this paper, a simple modifications to the BA scheme, which we call the BA-Drop scheme, is shown to have a substantial improvement on BLR of connectionless traffic injected into an ATM network. Performance of the BA-Drop scheme with different values of the bandwidth advertising interval is studied. It is found that optimal values exist for different traffic loading. The effect of burstiness and propagation delay of the feedback signal do not seem to have a major impact on the effectiveness of the BA-Drop scheme, allowing it to be used in a variety of network environments. We are currently studying some additional enhancement such as better management of the gateway and switch buffers as well as the effect of the scheme on specific type of traffic such as real-time traffic as well as self similar traffic [9]. We hope to report on the results of these investigations in a future paper.

References

[1] Mongiovi, L., Farrell, M. & Trecordi, V., "A Proposal for Interconnecting FDDI Networks through B-ISDN," *Proc. IEEE INFOCOM '91*, Bal Harbour, FL, April 1991.

[2] Gerla, M., Tai, T., Gallassi, G., "LAN/MAN Interconnection to ATM: A Simulation Study," *Proc. IEEE INFOCOM '92*, 1992.

[3] Gerla, M., Tai, T., Gallassi, G., "Internetting LAN's and MAN's to B-ISDN's for Connectionless Traffic Support," *IEEE JSAC*, Vol. 11 No. 8, Oct. 1993, pp. 1145-1159.

[4] Crocetti, P., et. al., "Multicast in SMDS over an ATM Network," *Proc. Globecom '92*, Milan, Italy, 1992, pp. 1062-1066.

[5] Crocetti, P., et. al., "Interconnection through LAN/MANs through SMDS on Top of an ATM Network," *Computer Communications*, Vol. 16, No. 2, Feb. 1993, pp. 83-92.

[6] Mateescu, M., "On Allocation Schemes for the Interconnection of LANs and Multimedia Sources over Broadband Networks," *Local Area Network Interconnection*, Onvural, R.O., Nilsson, A. (ed.), Plenum Press, 1993.

[7] Monterio, J.A.S. et. al., "Statistical Multiplexing in ATM Networks," *Proc. 4th Int. Conf. Data Commun. Systems and Their Performance*, Barcelona, Spain, 1990.

[8] Romanow, A., Floyd, S., "Dynamics of TCP Traffic over ATM Networks," *IEEE JSAC,* Vol. 13, No. 4, May 1995, pp. 633-641.

[9] Willinger, W. et. al., "Self-Similarity Through High-Variability: Statistical Analysis of Ethernet LAN Traffic at the Source Level," *Proc. ACM SigComm '95*, Cambridge, U.S.A, 1995.

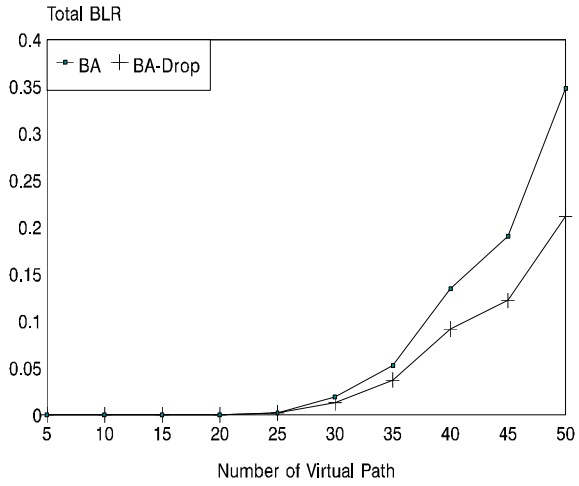

Figure 5(a): Total BLR

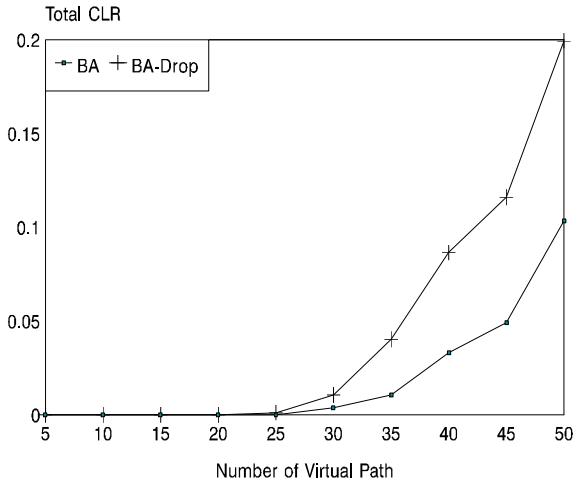

Figure 5(b): Total CLR

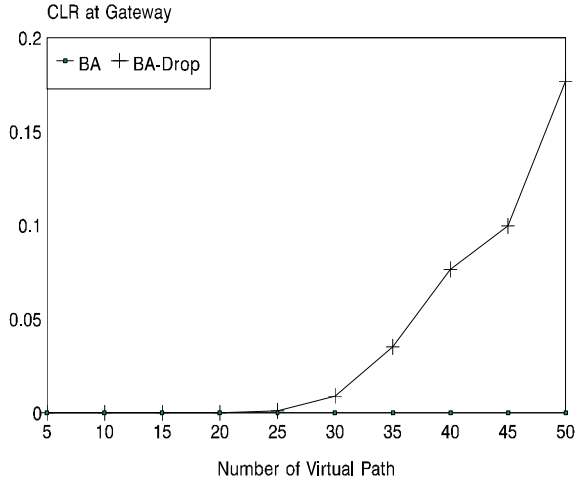

Figure 5(c): CLR in the Gateway

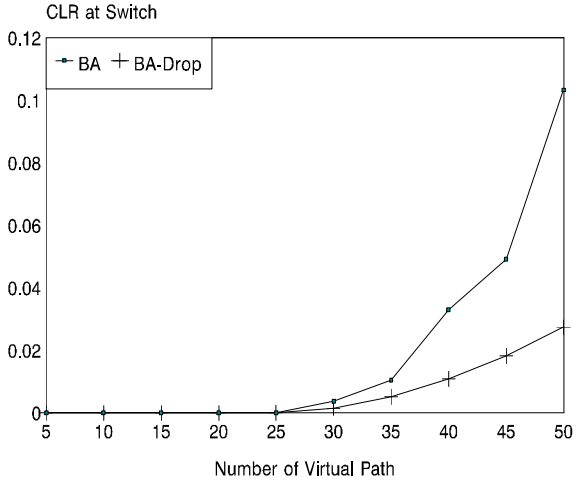

Figure 5(d): BLR in the ATM Switch

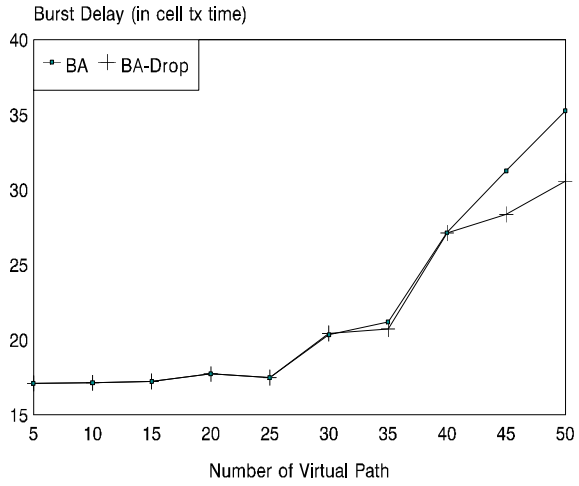

Figure 5(e): Burst Delay

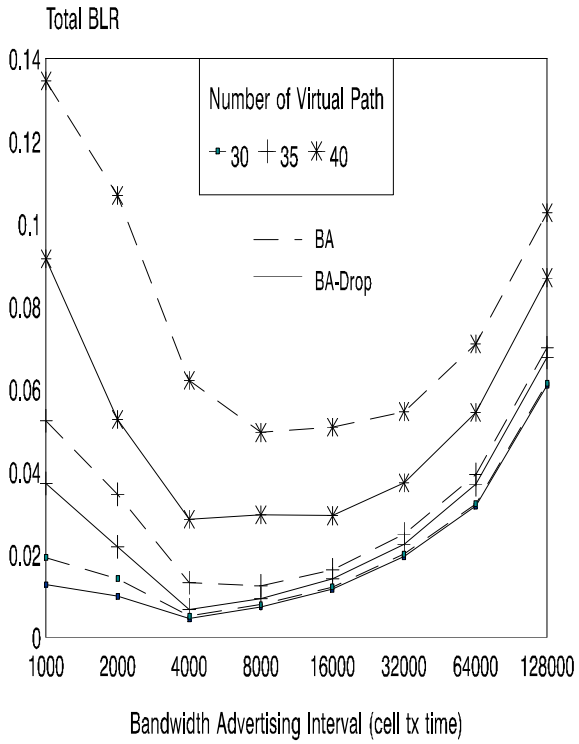

Figure 6(a): Total BLR

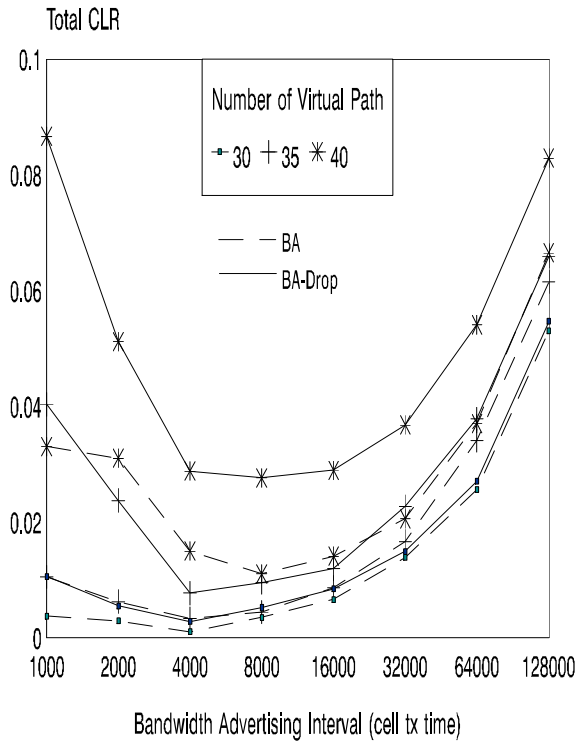

Figure 6(b): Total CLR

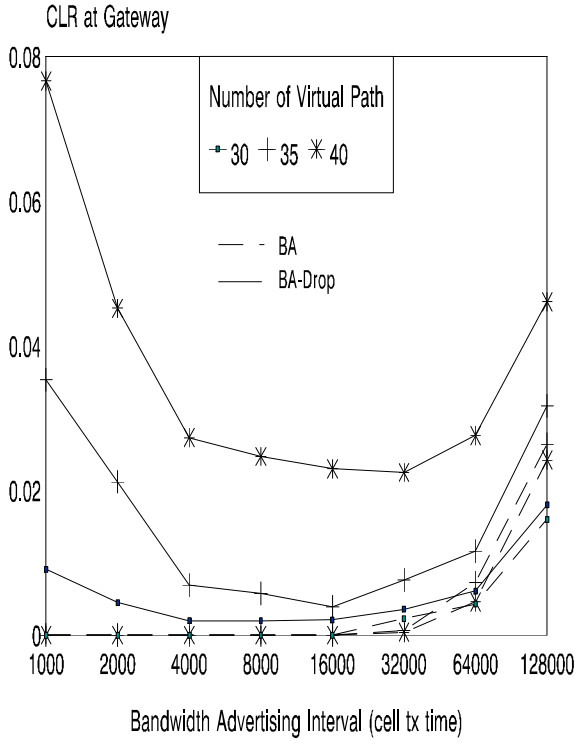

Figure 6(c): CLR in the Gateway

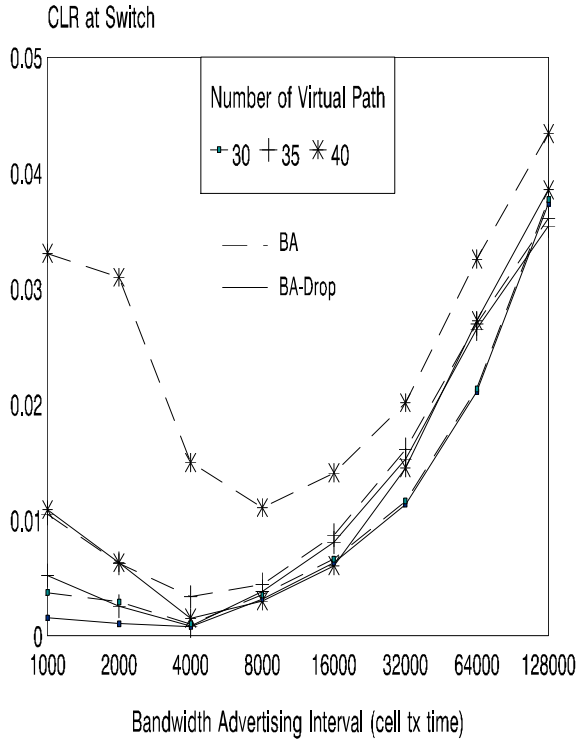

Figure 6(d): CLR in the ATM Switch

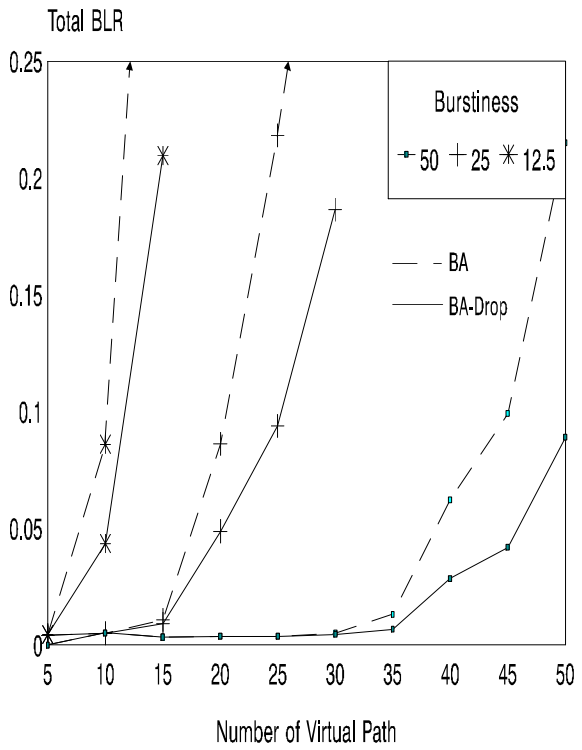

Figure 7(a): Total BLR with Different Burstiness

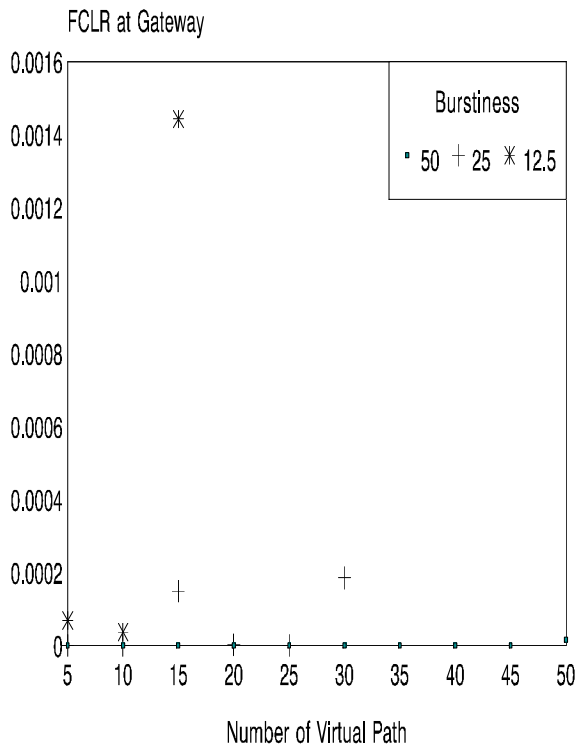

Figure 7(b): False CLR in the Gateway

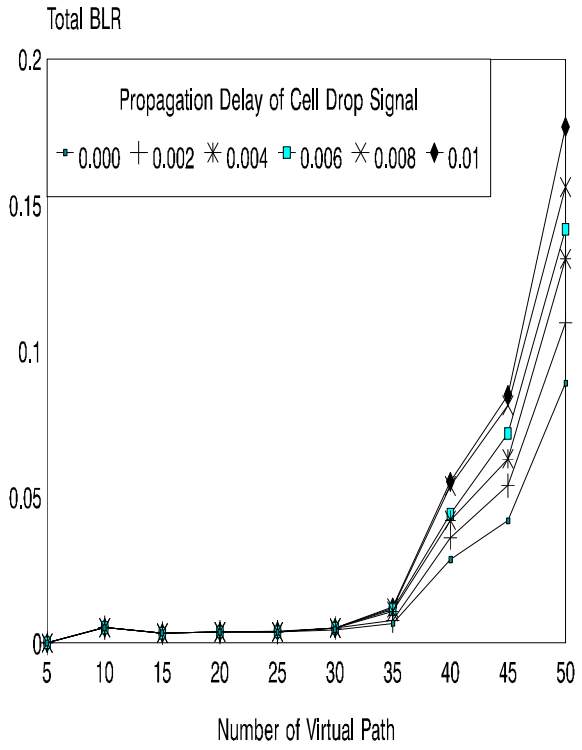

Figure 8(a): Total BLR with Different Propagation Delay of the Cell Drop Signal

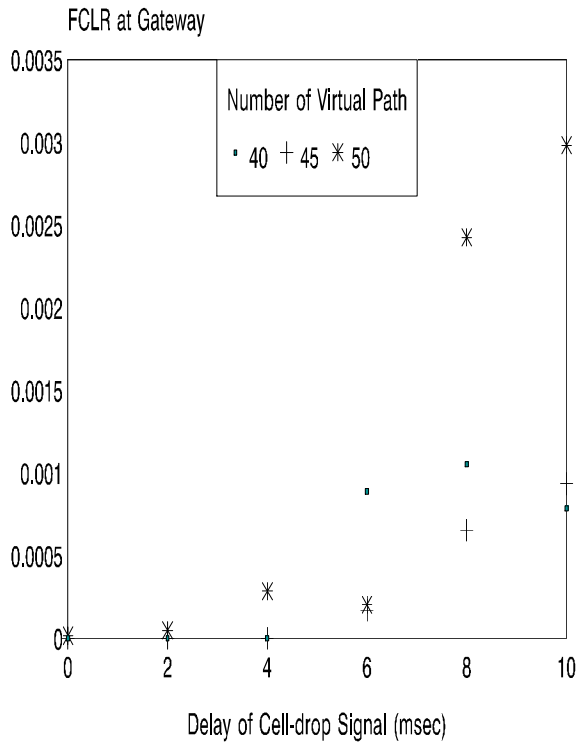

Figure 8(b): False CLR in the Gateway