Effectiveness of Timer-Based Connection Management Mechanisms for IP/ATM Networks

Mahbub Hassan, Richard Egudo and Kieran Power

ATM Research Lab
School of Computing & Information Technology
Monash University, Gippsland Campus
Churchill, Victoria 3842, Australia.
Tel: +61-3-9902 6568, Fax: +61-3-9902 6842
Email: mahbub.hassan@fcit.monash.edu.au

Mohammed Atiquzzaman

Department of Electrical and Computer Engineering
The University of Dayton
Dayton, Ohio 45469-0226, USA.
Tel: +1-513-229 3611, Fax: +1-513-229 2471
Email: atiq@engr.udayton.edu

Abstract. A timer-based connection management mechanism has been recently standardised to manage ATM connections dynamically for transporting IP traffic over ATM networks. However, implementing timers will increase the complexity and cost of ATM network interface cards. In this paper, the effectiveness of timer-based connection management techniques for minimising transmission costs for IP traffic over ATM networks for a typical intranet-ISP connection is analysed. Our study shows that to minimise the transmission cost, the SVC should be closed immediately or left open depending on whether the mean packet interarrival time is greater than or less than the ratio of set-up to holding cost respectively. Further study is required to manage the timer when the interarrival time is close to the ratio of set-up to holding cost. The analytical results are validated by simulating an IP/ATM router with the well known fractal-like, self-similar traffic collected from real networks.

1 Introduction

Intranet is a rapidly growing internetworking technology for connecting local computers and resources within an organisation using the standard Internet Protocol (IP). Intranet has enormous benefits over proprietary networking as it provides ready access to the global Internet which connects millions of computers all over the world. Organisations world-wide are rapidly deploying intranets as their communication structure within the organisation.

The intranet of an organisation is connected to the global Internet via an Internet Service Provider (ISP). The connection between the intranet and the ISP is usually established over a leased line from the public telecommunication companies. The leased line connects the IP router of the intranet to the IP router of the ISP. For bursty data traffic, the bandwidth of the leased line is not utilised efficiently. The organisation has to pay for the bandwidth for the entire lease period whether there is traffic or not. The cost of maintaining leased lines is, therefore, very high.

With Asynchronous Transfer Mode (ATM) [Prycker, 1995], a cell-switching network, set to become the world-wide standard [ITU-TS, 1991] for the next generation of telecommunication networks, it is expected that such leased lines will be replaced by ATM virtual circuits (VCs). Interconnection of an intranet to an ISP via such an ATM VC is shown in Figure 1. All traffic from the intranet will be multiplexed onto a single high speed ATM VC.

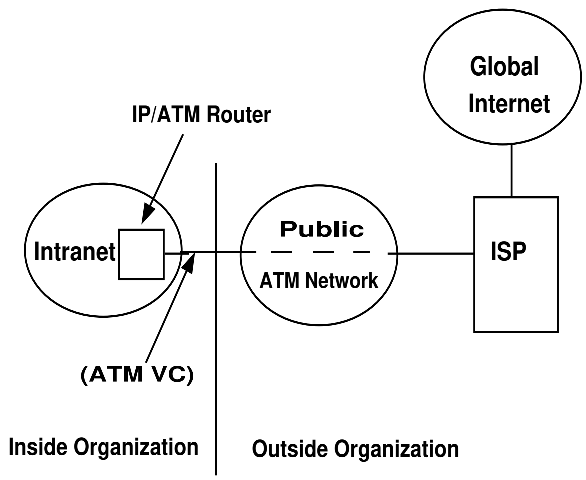

Fig. 1. The interconnection of an IP-based intranet to an ISP through the public ATM network.

Since ATM is a connection-oriented network, the intranet router needs to set up an ATM VC with the ISP router first before it can transmit the IP packet. The connection may be released after the transmission of the IP packet.

ATM connections which are set up and released dynamically as needed are called *switched virtual circuits* (SVC). One problem with SVCs is the large number of connection set-ups required. Since setting up of an SVC creates signalling traffic in the ATM network and reserves resources in the network, it is likely that there will be a separate charge for each set-up (called *set-up cost*) in addition to a time-based charge for holding an SVC open (called *holding cost*). The data transmission cost through ATM networks may increase significantly if the routers generate too many SVC set-up requests.

A timer-based connection management technique has been standardised [Perez et al., 1995] to reduce the number of SVC set-ups for IP/ATM networks. Instead of releasing an SVC immediately after transmitting an IP packet, a router starts a timer and leaves the SVC open. The SVC is released if it remains idle for a specified timeout period. However, implementing such a timer will increase the complexity and cost of the ATM network interface card (NIC).

No study was reported in [Perez et al., 1995] to evaluate the effectiveness of the timer-based connection management mechanism in minimising the transmission cost (holding cost *plus* set-up cost) for IP traffic over ATM networks. Later, Keshav et al. [Keshav et al., 1995] reported that timer-based connection management mechanisms can effectively minimise the transmission cost if the timeout value is set to *half of the ratio of the set-up cost to holding cost*. This result was based on a simulation of an IP/ATM router with traces of IP traffic collected from actual networks.

There are a number of limitations in the study reported in [Keshav et al., 1995] which prevent it from applying to the intranet-ISP connections. Firstly, the authors of [Keshav et al., 1995] partitioned the IP traffic into separate groups based on source-destination pairs in the IP headers and considered *separate* ATM SVCs for each group. Since for an intranet-ISP connection, all IP packets are transmitted through the *same* ATM SVC, the packet interarrival distribution for the ATM SVC for an intranet-ISP connection will be different than the distribution observed in [Keshav et al., 1995]. Secondly, the sample sizes in the traces used for simulation in [Keshav et al., 1995] were very small; they range from 2000 packets to 20,000 packets only. For very high speed ATM SVCs, this represents traffic for only a few seconds.

In this paper, we analyse the effectiveness of timer-based connection management mechanisms for an IP/ATM router connecting an intranet to an ISP through the public ATM networks without the above limitations. We consider a single ATM SVC to transmit all traffic from the intranet and conduct simulations with trace size of one million packets. *In contrast to the results reported in [Keshav et al., 1995], our results suggest that the transmission cost cannot be minimised by setting the timeout value to half of the ratio of the set-up cost to holding cost.* To minimise transmission cost, the SVC should be closed immediately if the mean packet interarrival time is greater than the ratio of set-up to holding cost. If the mean packet interarrival time is smaller than the ratio of set-up to holding cost, the SVC should be left open. Further study is required to manage the timer when the interarrival time is close to the ratio of set-up to holding cost.

Our results are first derived from an analysis based on exponential packet interarrival times. Later, we show that our analytical results compare favourably with simulation results obtained by simulating an IP/ATM router with the well known fractal traffic [Leland et al., 1994] of one million IP packets collected from real networks. Although the analysis is based on a traffic model with exponential interarrival times, simulation results confirm that the analysis is equally valid for most case of fractal-like, self-similar traffic measured in real networks. This further strengthens the conclusion that timers may not always be effective in minimising transmission costs for IP/ATM networks.

The rest of the paper is organised as follows. The analysis for evaluating the effectiveness of timers is presented in Section 2. The simulation model is presented in Section 3 followed by the results in Section 4. Finally, concluding remarks are provided in Section 5.

2 Analysis

We consider a pricing scheme for ATM networks similar to that of existing telecommunication networks. There is a holding cost for every time unit the SVC remains open. No cost is incurred if the SVC is closed. In addition to the holding cost, there is a set-up cost for every SVC set-up.

If the next IP packet arrives before the timer expires, there is only a holding cost for the packet. However, if the IP packet arrives after the timer expires, there is a holding cost and a SVC set-up cost. Therefore, the average cost for a packet for a given timeout value depends on the probability distribution of the packet interarrival times.

The traffic from the intranet is modeled by the Poisson arrival process with exponential packet interarrival times. For exponentially distributed interarrivals, the average cost (normalised to the holding cost of one time unit) for a packet for a given timeout value t can be obtained by adding the average cost for interarrivals shorter than the timeout value to the average cost for interarrivals longer than the timeout value as

$$\begin{aligned} C(t) &= \int_0^t x\lambda e^{-\lambda x} dx + (t+\mathcal{S}) \int_t^\infty \lambda e^{-\lambda x} dx \\ &= \frac{1}{\lambda} - \frac{e^{-\lambda t}}{\lambda} + \mathcal{S} e^{-\lambda t}. \qquad (1) \end{aligned}$$

where λ is the mean packet arrival rate ($\frac{1}{\lambda}$ is the mean interarrival time) and $\mathcal{S}$ is the ratio of SVC set-up cost to holding cost of one time unit.

For the above cost model, the optimum timeout value for minimum cost is obtained from the first derivative of $C(t)$ given by

$$C'(t) = e^{-\lambda t}(1 - \mathcal{S}\lambda) \qquad (2)$$

From Equation (2) we consider the following three cases to study the effectiveness of a timer in minimising transmission cost:

Case 1. $\frac{1}{\lambda} > \mathcal{S}$. For this case, $C'(t)$ is positive for any t meaning $C(t)$ is an increasing function. The minimum cost is achieved for minimum t which is $C|_{t=0} = \mathcal{S}$. Therefore, the SVC should be released immediately upon becoming idle. *No timer is necessary.*

Case 2. $\frac{1}{\lambda} < \mathcal{S}$. $C'(t)$ is negative for any t and $C(t)$ is a decreasing function. The minimum cost is achieved for maximum t which is $C|_{t=\infty} = \frac{1}{\lambda}$. For this case, the SVC should be left open. *Again, no timer is required.*

Case 3. $\frac{1}{\lambda} = \mathcal{S}$. The first derivative $C'(t)$ equals to zero for any t suggesting that $C(t)$ is constant; the cost is $C|_{\text{any } t} = \frac{1}{\lambda} = \mathcal{S}$. *A timer has, therefore, no effect on the cost.*

The above observations suggest that for traffic with exponential packet interarrival times, the transmission cost for IP/ATM networks *can not be minimised by implementing a timer*. The validity of this result is investigated in the following section by simulating an IP/ATM router with the well known fractal traffic [Leland et al., 1994] collected from real networks.

3 Simulation Model

In order to validate the results obtained in the previous section, we have developed a simulation model using OPNET, a high performance commercial tool for simulating communication networks. Figure 2 shows the node level diagram of our simulation model. It can be seen that the node diagram contains three nodes: *Measured IP Traffic, Queue (fifo)* and *ATM_network.*

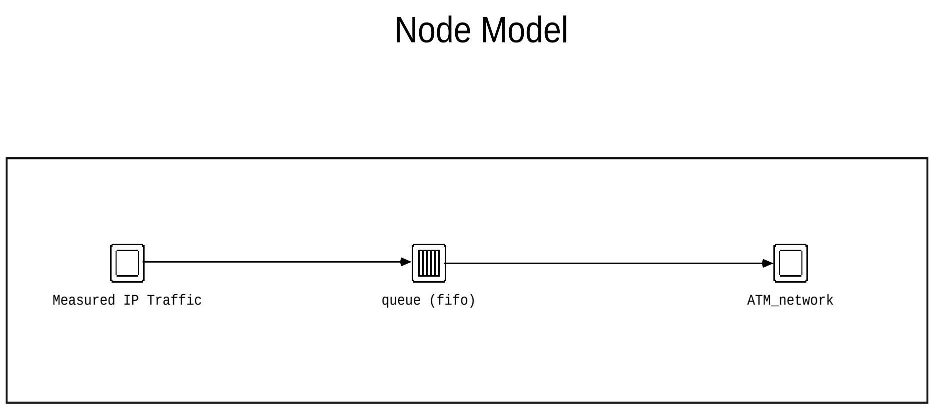

Fig. 2. OPNET model for an IP/ATM router.

The *Measured IP Traffic* node generates packets according to the arrival information stored in an external data file, which contains a line of text specifying the size and interarrival time for every packet to be generated. In the case of our simulation, we used the packet arrival information collected by Leland et al. [Leland et al., 1994]. This arrival information was collected from networks at the Bellcore Morristown Research and Engineering facility, which carry a major portion of the Lab's traffic, including all traffic to and from the internet. The collected data are available in four files, *pAug.TL, pOct.TL, OctExt.TL* and

OctExt4.TL, available via anonymous FTP at ftp.bellcore.com/pub/lan_traffic. Each of these files contains packet size and packet interarrival times for approximately one million packets for different times of the year. For our simulations, we used the data contained in the file *pOct.TL* . A summary of the packet arrival data in *pOct.TL* is provided in Table 1.

Total number of packets	1,000,000
Mean packet size	628 bytes
Mean packet interarrival time	1.76 msec

Table 1. Summary of packet arrival data in file *pOct.TL*.

The *Queue (fifo)* node simulates a buffer in the IP/ATM router in the intranet and is based on the standard fifo queue provided with the OPNET simulator. However, a couple of alterations have been made to simulate a timeout mechanism for the server (ATM SVC in this case). If a packet arrives while the SVC is busy transmitting a packet, the packet is buffered in this queue. This can effectively simulate an SVC with an arbitrary bandwidth. In our simulation, we have set the SVC bandwidth to a very high value to avoid any queuing in the IP/ATM router.

The *ATM_network* node is based on the sink process provided with OPNET. In the case of our simulations, the *ATM_network* node simply discards the packet once it has been received.

The cost of the next arriving packet is calculated as the holding cost of the SVC before the packet arrives plus a set-up cost if the SVC must be re-opened for this packet. At the end of the simulation, an average cost is calculated for all the packets simulated.

4 Results

In this section, we compare the average costs for the transmission of a packet obtained from the simulation with the ones obtained from the analysis for each of the three cases discussed in Section 2. The analytical cost values are obtained by replacing $\frac{1}{\lambda}$ in Eqn. (1) by the mean packet interarrival time of the one million packets, which is 1.76 msec.

Figures 3 through to 5 illustrate the comparison of simulation and analytical results. It can be clearly seen, that for $\frac{1}{\lambda} > \mathcal{S}$ (Figure 3), the average costs obtained from both simulation and analysis increase as a function of the timeout value and asymptotically approaches $\frac{1}{\lambda}$, the mean packet interarrival time. Similarly, Figure 4 shows that for $\frac{1}{\lambda} < \mathcal{S}$, the cost is a decreasing function of the timeout value. Figure 5 shows that the cost remains stable for $\frac{1}{\lambda} = \mathcal{S}$ for Poisson packet arrival. Therefore, the simulation of fractal traffic obtained from

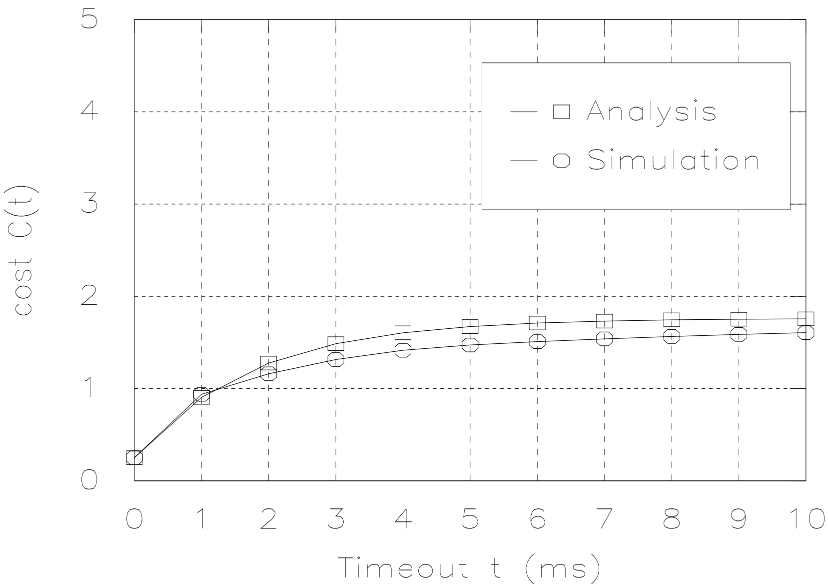

Fig. 3. Average costs for $\mathcal{S} = 0.25$ ($\frac{1}{\lambda} > \mathcal{S}$).

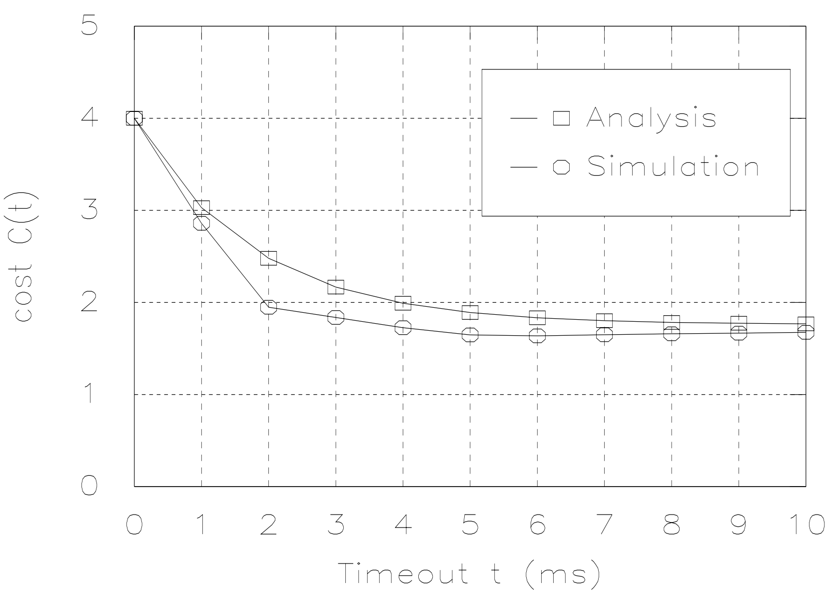

Fig. 4. Average costs for $\mathcal{S} = 5.0$ ($\frac{1}{\lambda} < \mathcal{S}$).

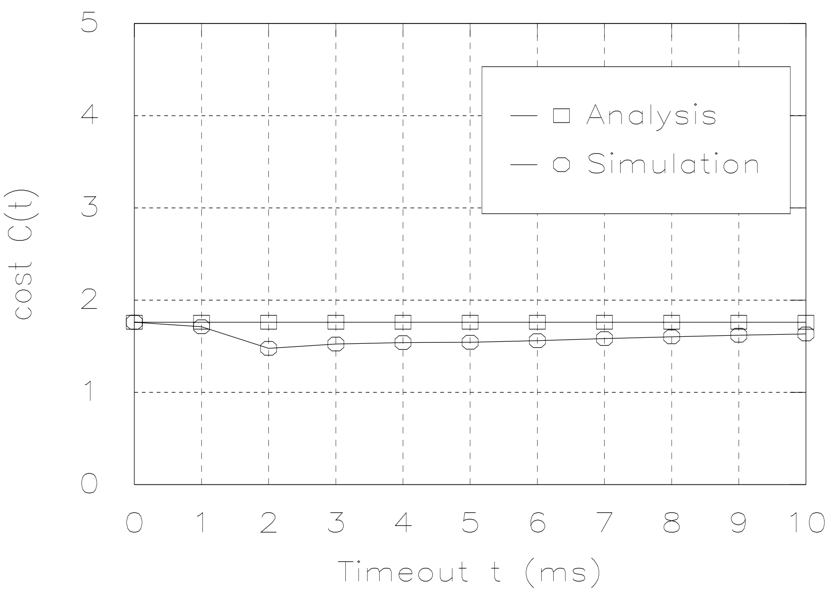

Fig. 5. Average costs for $\mathcal{S} = 1.76$ ($\frac{1}{\lambda} = \mathcal{S}$).

real networks compares favourably with the results obtained from the analysis based on Poisson packet arrival model. These simulation results agree with the conclusion that, in most cases, timer-based connection management mechanisms may not be effective in minimising transmission costs for IP/ATM networks. In this paper we have used only one file from the Bellcore collection. Further study involving more empirical data is required to study the effectiveness of the timer for the case when $\frac{1}{\lambda} \simeq \mathcal{S}$.

5 Conclusion

We have analysed the effectiveness of timer-based connection management techniques for minimising transmission costs for IP traffic over ATM networks for a typical intranet-ISP connection. In contrast to previously reported results, we have shown that timers are not effective in minimising the transmission cost when all IP traffic from an intranet is multiplexed onto a single ATM SVC. The decision regarding whether to release an SVC or to leave it open for minimising the transmission cost depends only on the relationship between the mean packet arrival time and the ratio of set-up cost to holding cost for the SVC. Eliminating the timer will reduce the complexity and cost of an ATM NIC.

In deriving the above result, we first obtain the average transmission cost of a packet as a function of the timeout value using a simple Poisson packet arrival process for the ATM SVC. Later, we validate the analytical results by simulating an IP/ATM router with the well known fractal traffic collected from real networks by other researchers. The simulation results suggest that, in most cases, the conclusions obtained from the analysis with Poisson arrival process are similar to those obtained for fractal-like, self-similar traffic present in real networks. This is a significant observation since in most cases it eliminates the need for analytical modeling with fractal packet arrival process which is much more complex and difficult than modeling with Poisson arrival process.

Acknowledgment

This work was partly funded by a Monash Research Grant and the authors greatly acknowledge this support.

References

ITU-TS (1991). *Recommendation I.121: Broadband Aspects of ISDN.* ITU, Geneva.

Keshav, S., Lund, C., Phillips, S., Reingold, N., and Saran, H. (1995). An Empirical Evaluation of Virtual Circuit Holding Policies in IP-Over-ATM Networks. *IEEE JSAC*, 13(8):1371–1382.

Leland, W.E., Taqqu, M.S., Willinger, W., and Wilson, D.V. (1994). On the Self-Similar Nature of Ethernet Traffic (Extended Version). *IEEE Transactions on Networking*, 2(1):1–15.

Perez, M., Liaw, F., Mankin, A., Hoffman, E., Grossman, D., and Malis, A. (1995). *RFC 1755: ATM Signalling Support for IP over ATM.*

Prycker, M.D. (1995). *Asynchronous Transfer Mode.* Prentice Hall, third edition.