

30

Indicators for the assessment of congestion in TCP over ATM-UBR¹

Mohammad A. Rahin and Mourad Kara

School of Computer Studies

The University of Leeds, LS2 UK

E-mail: {rahin, mourad}@scs.leeds.ac.uk

Telephone: +44 113 233 6590

Fax: +44 113 233 5468

Abstract

In this paper we analyse congestion in TCP over ATM-UBR, in particular the association between the congestion measures at the ATM physical and TCP transport layers. Two new indicators, protocol level congestion association factor (PLCAF) and CLR-PRR correlation coefficient are used. ATM cell loss ratio (CLR) and TCP packet retransmission ratio (PRR) are good measures of congestion at their respective layers, but do not address the relation between cell loss and its effect on TCP congestion control schemes. Both PLCAF and CLR-PRR correlation coefficient are derived to better capture the association between congestion measures at ATM and TCP levels. It is shown through simulation of multiple TCP connections over a ATM-UBR service under selected number of different operating conditions that the two new indicators can be successfully used for the analysis of congestion in TCP over ATM-UBR.

Keywords

TCP, ATM, congestion indicator, congestion association

¹This work is partially funded by a UK-ADF grant (University of Leeds).

1 Introduction

Performance evaluation of TCP over ATM has attracted the attention of a number of researchers [5][6][9][10]. The main causes for the performance degradation in TCP over ATM may be broadly classified into three different categories [7], *i*) due to the dynamics of TCP, *ii*) due to the behaviour of ATM and *iii*) due to the interaction between TCP and ATM layer congestion control mechanisms. A thorough knowledge about the interaction between congestions at these two distinct levels would lead to the understanding of the dynamics of TCP over ATM and would lead to a systematic appraisal of the TCP over ATM performance. A number of analytical and numerical studies e.g., [12][2] sought to estimate the burstiness of cell loss, i.e., the short overload period when cell losses occur followed by long loss-free intervals from the conditioned cell loss and string loss (loss of contiguous cells) probabilities. It is however, noted that conditional and string cell loss probabilities are not always sufficient to predict higher layer performance.

In this paper the relationship between congestion measures at TCP and underlying cell level ATM layers is investigated through simulation. Strength of association between these two measures are made using correlation coefficient of the congestion measures as well from their ratios. Of the two ATM service classes suitable for carrying TCP traffic, the UBR class is chosen in particular because, unlike ABR, UBR does not employ any congestion control mechanism at the ATM level and hence would allow the investigation of the relationship between the unhindered cell level loss to higher layer TCP level loss. The investigation is carried out with and without the presence of any intelligent packet drop policies e.g., EPD, PPD [10]. In section 2 the congestion control mechanisms in TCP and ATM-UBR are reviewed. Two congestion assessment indicators, CLR-PRR correlation coefficient and PLCAF are introduced in section 3. Section 4 describes the simulation experiments for the performance evaluation and congestion assessment of TCP over ATM-UBR under both LAN and WAN scenarios. Section 5 analyses the simulation results focusing on the congestion phenomenon. The paper is concluded in section 6 with a few remarks and direction for future works.

2 Congestion control

The congestion control mechanism in TCP is window-based providing end-to-end flow control to limit the number of packets in the network. There are two windows, one at the receiving end (receiver's window, RCVWND) as a measure of its buffer capacity and the other at the sending end (congestion window, CWND) as a measure of the network capacity. The sender can only place one window full of data at a time, but is not allowed to inject more than the minimum of RCVWND and CWND into the network.

The TCP congestion control mechanism has two distinct phases, *Slow Start* and *Congestion Avoidance* (Figure 1). During the slow start phase, either at the very

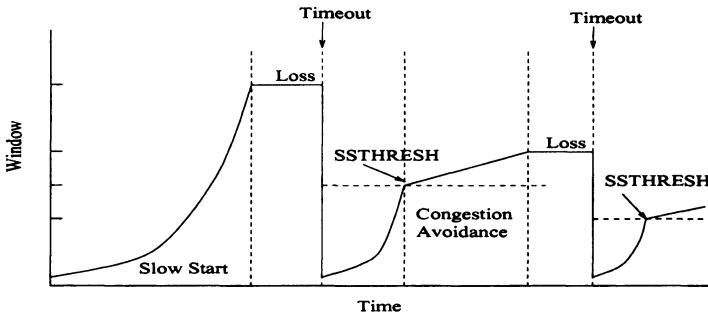


Figure 1: TCP congestion window vs. time.

beginning of the connection or when recovering from a packet loss, CWND is set to 1 TCP segment (usually 512 bytes) and is doubled once every *Round Trip Time* (RTT), i.e., on receipt of an acknowledgement. A delay-estimation algorithm is part of TCP collection of algorithms which provides a good estimate of the round-trip delay. During the congestion avoidance phase CWND is linearly increased and in doing so probes for additional available bandwidth in the network. The congestion avoidance phase starts whenever CWND reaches the value of another parameter STHRESH, usually initially set to 64 kB.

If a TCP packet is lost en route or received damaged, the destination informs this to the sender through duplicate acks for each out-of-order segment received. When duplicate acks are received the source refrain from sending any new segment in order not to increase the network load. At the source side, a timeout counter is maintained for the last unacknowledged segment which is reset on receipt of its ack. The source assumes congestion in the network whenever the timeout counter goes off. So as to remedy the congestion in the network the source sets STHRESH to half of the current value of CWND and CWND is set to one. Consequently, the slow start phase restarts. The source then retransmits the lost segment and increases its CWND by one every time a new segment is acked until CWND reaches STHRESH whereupon the congestion avoidance phase begins. All segments since the lost segment must be retransmitted before the source can send any new segment. This constitutes the go-back-N policy.

The TCP source continuously monitor RTT by measuring the time between the sending of a segment and when its ack is received. The retransmission timeout value is calculated from the estimates of the average and mean deviation of the RTT [4]. The fast-retransmit and fast-recovery mechanism [11] improves the TCP performance by not restarting the slow start phase for congestions indicated by the receipt of duplicate acks. However, even with fast-retransmit method the source must refrain from sending any more segments till either n (typically $n = 3$) duplicate acks are received when the lost segment is retransmitted or when the fast-retransmit timer goes off indicating that the apparently lost segment has at last reached the destination through an alternate route.

3 Assessment of congestion

Congestion in ATM network is generally specified using the metric *Cell Loss Ratio* (CLR). CLR is defined as the ratio of the number of cells that are lost en route to the total number of cells transmitted by the sources. Similarly, the TCP level congestion may be specified by the metric *Packet Retransmission Ratio* (PRR), defined as the ratio of the total number of packets retransmitted to the total number of packets sent by the sources. CLR and PRR are both long-term as opposed to instantaneous indicators of congestion at their respective layers. CLR has been found not to be a good indicator of TCP level performance [6]. In [5] TCP performance is analysed in terms of PRR in TCP over ATM. It was found that PRR is inversely related to switch buffer size, but the actual interaction between congestion measures at ATM and TCP layers under different circumstances was not established.

A single cell dropped at a switch results in the loss of one entire packet. In the worst case scenario, cells dropped are in the form of bursts of one cell long only evenly distributed in the traffic stream. This would cause very high levels of packet retransmissions and throughput loss. Ideally, unavoidable cell losses should be in bursts of a large number of cells such that at most only one or a couple of packets are affected. Both PPD and EPD aim to achieve this goal of concentrating cell losses into single packets. The interaction between cell loss and packet retransmission is difficult to predict as it depends on the nature of traffic flows as well as the cell occupancy levels at the switch. An assessment of the level of interaction, i.e., how cell losses affect TCP throughput by causing packet retransmission would lead to the understanding of effectiveness of cell drop policies, buffer allocation schemes or the source characteristics.

The association between cell loss and the consequent packet retransmission may be analysed using the correlation coefficient between CLR and PRR over a set of observations as follows.

$$\text{Corr}(C, R) = \frac{E[(C_i - \bar{C}_i)(R_i - \bar{R}_i)]}{\sqrt{E[(C_i - \bar{C}_i)^2]} \sqrt{E[(R_i - \bar{R}_i)^2]}}, \quad (1)$$

where C_i and R_i are the i th CLR and corresponding PRR observations. The correlation between CLR and PRR is a measure of the association between them. If C_i and R_i are independent, $\text{Corr}(C, R)$ is zero. On the other hand, a value close to either 1 or -1 would suggest strong association between them. Correlation between cell loss and packet retransmission may be calculated from the observed CLR and PRR of a number of TCP flows passing simultaneously through a switch or from the aggregate CLR and aggregate PRR at a switch for different values of switch buffer sizes, a factor affecting congestion. The CLR-PRR correlation may be used as a tool to compare the effectiveness of different cell drop policies or buffer allocation schemes in terms of strength of association between congestion measures at different levels in protocol stack.

The CLR-PRR correlation provides summarised information on the strength of

association between ATM and TCP layer congestion measures over a set of observations. However, snapshot information on congestion correspondence for a single observation can not be readily obtained from CLR-PRR correlation. Such information may be obtained from the ratio of normalised packet retransmission to normalised cell loss, i.e. the ratio of PRR to CLR for the connection. In TCP over ATM-UBR, packet retransmission may only take place as a result of actual cell loss. The PRR to CLR ratio may thus be used as the true measure of association between cell loss and packet retransmission. We call this ratio *Protocol Level Congestion Association Factor* (PLCAF). The notion of PLCAF however, by definition is non-existent in congestion free i.e., zero cell loss situations. PLCAF may be formally expressed as:

$$PLCAF = \begin{cases} \frac{PRR}{CLR} & ; CLR \neq 0, PRR \neq 0 \\ \text{not defined} & ; \text{otherwise} \end{cases} \quad (2)$$

In non-zero cell loss situations, PLCAF is a positive number within the range $0 < PLCAF \leq (MTU/48)$. MTU is a TCP parameter specifying the largest size of TCP packets and the figure 48 is derived from the payload size of ATM cells. PLCAF may be interpreted as a measure of the number of packets affected as a result of single cell loss. For infinite TCP sources the average number of packets affected per cell loss can be obtained by multiplying PLCAF with $(48/MTU)$. A large PLCAF would indicate that cell losses are spread across the traffic stream in small bursts causing large number of TCP retransmissions. On the other hand, long bursts of cell losses causing relatively small number of retransmissions would result in smaller PLCAF, a desirable property. PLCAF plotted against switch buffer size can be used for congestion analysis in terms of buffer constraints. Also, PLCAF vs. switch buffer size plots can be used as effective tools for the comparison of different cell drop policies or traffic parameters. With an intelligent packet drop policy (e.g. PPD or EPD) active it is expected that PLCAF would decrease with increasing switch buffer size. An increase of PLCAF with increasing switch buffer size would indicate indiscriminate cell dropping and a failure to exploit relaxed buffer constraints.

Both the CLR-PRR correlation coefficient and PLCAF indicators provide information relating to the correspondence between congestion measures at ATM physical and TCP transport layers and are better indicators for the assessment and analysis of congestion in TCP over ATM as both ATM & TCP layers are considered in unison and not in isolation. However, it must be borne in mind that they alone or together do not fully specify the actual level of congestion for which either CLR and PRR must be used in conjunction.

4 Experimental setup

In this section the network configurations and parameters used in the simulations of TCP congestion behaviour over ATM-UBR are described. Fig 2 shows the network configuration used in the simulations. It consists of 2 switches, N TCP connections and $2N$ end-stations. Only one-way TCP connections (discounting the TCP ack

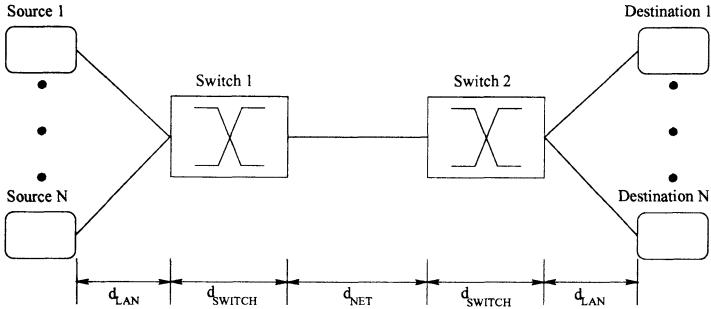


Figure 2: The N -source network configuration.

in the reverse directions) are considered. The link between the two switches is the congestion bottleneck. As study of congestion is the primary objective of the simulations, this particular configuration is selected specifically with a view to concentrate all congestions in a single link/switch. No additional background or cross-traffic sources are considered.

The N -identical TCP sources on the left side of the figure each establish a single connection with the similarly numbered destination across the two switches and the bottleneck link. The TCP sources are greedy sources with infinite supply of data and always have data ready to send as and when permitted by the TCP flow-control. The two end-station groups are considered in close proximity ($d_{LAN} = 0.5\mu s$) with their respectively connected switches. The internal delay in both switches are identical and is set to $d_{SWITCH} = 450\mu s$, the maximum such value sanctioned in ITU-T standard [1]. Two network scenarios, LAN and WAN are considered. The delay in the bottleneck link is set to a typical value $d_{NET} = 10\mu s$ equivalent to 2 km of transmission path and $d_{NET} = 10ms$ equivalent to 2000 km of transmission path for LAN and WAN scenarios respectively. Each connection in the network configuration chosen have only one path and as such the lower bound of RTT is $2 * (2d_{LAN} + d_{SWITCH} + d_{NET})$ and the upper bound adds another term, a function of the buffering at the switches to it. All link bandwidths are 155 Mbps for LAN scenario and 622 Mbps for WAN scenario. The switch speeds are set accordingly, i.e., 155 Mbps and 622 Mbps for LAN and WAN respectively. Finally N , the number of sources is set to 5 and 20 for LAN and WAN respectively similar to those used in a number of published works [3][5][10] and deemed sufficient to generate traffic flows required in our study.

The STCP simulation package [8] is used in the simulation study. The retransmission timers are set to 200 ms and 500 ms for fast & slow timeouts respectively. The MTU is set to 9180 bytes as per RFC1626. The default maximum window size in STCP is 64 kB. This is sufficient to fill the network pipe in LAN, but inadequate in WAN. Despite this the maximum window size in WAN scenario is also set at 64 kB. This particular decision would present a relaxed network pipe capacity to the multiple TCP sources resulting in the facing of a much less bandwidth contention

environment. The sources will be less aggressive against each other as opposed to the situation in LAN scenario.

To study the congestion behaviours at ATM and TCP layers the following three different cases for each of LAN and WAN scenarios in the N -source network configuration are simulated².

- A. All N sources starting simultaneously. Switches do not employ any cell drop policy.
- B. The N sources start at random times. Switches do not employ any cell drop policy.
- C. The N sources start at random times. Both EPD & PPD cell drop policies are in force.

The above three simulation cases are particularly chosen to allow comparison of the congestion behaviour at both ATM and TCP level under different operating environments such as where TCP synchronisation plays an important role and also where intelligent cell drop policies (EPD/PPD) are employed at the switches

The switch architecture used in the simulation is a nonblocking, output-buffered type. Each input port have a private buffer of 128 cells. Each output port have additional buffer and implements a single queue on it. All VCs passing through an output buffer share this buffer. The scheduling policy at the output buffers is FIFO. In order to induce variable congestion in the switches, the size of the shared buffer in each output queue are varied for all simulation cases in both LAN and WAN scenarios. The buffer sizes arbitrarily chosen are 100, 200, 400, 600, 800, 1000, 2000, 4000, 6000, 8000 and 10000 cells. The earlier part of the range offer very tight buffer constraint, while on the other hand, the latter part of the range offer much more relaxed buffer constraint. The switches may be configured with no specific cell drop policy, PPD or EPD. The latter two may be configured either independently or in combination. For EPD, we set the threshold for discarding incoming cells according to the following formula:

$$\text{threshold} = \max \left\{ 0.8, \frac{\text{buffer size} - 3 \cdot \text{segment size}}{\text{buffer size}} \right\}. \quad (3)$$

The threshold is thus set equal to the buffer size minus three times the TCP segment size. The first term 0.8 in the above formula avoids setting the threshold to a very low value when the buffer size is small and segment size is large.

Each simulation was run for a time equivalent to the transfer of 10^{10} bits at the link speed being simulated. This turned out to be 100.01 seconds for LAN and 24.92 seconds for WAN. Each simulation run was divided into 5 batches and average of these 5 batches are taken. Also no additional warm-time times were allowed. This was to capture the TCP behaviour during simultaneous opening of multiple TCP connections. The simulation run times are sufficiently large to capture long-term congestion behaviour.

²The notation used to refer to a particular experiment is scenario.case, e.g., LAN.A.

Table 1: Buffer requirement for zero cell loss

Number of Sources	Network Configuration	Network Capacity (Mbps)	Max. Throughput Achieved) (Mbps)	Maximum Queue (Cells)
5	LAN.A	155	133.54	4000
5	LAN.B	155	130.57	4000
5	LAN.C	155	130.57	6000
20	WAN.A	622	328.54	8000
20	WAN.B	622	207.68	800
20	WAN.C	622	207.68	1000

5 Results & evaluation

In this section the simulation results for the three cases discussed in section 4.1 under both LAN and WAN scenarios are presented³. The focus is mainly on the association between the congestion measures at ATM and TCP levels using the indicators CLR-PRR correlation coefficient and PLCAF as defined earlier in section 3. The switch buffer size was chosen as an artificial means to vary the level of congestion. In this study we only consider the long-term congestion behaviour of the network.

The summary of the results for zero cell loss situations are shown in Table 1. The maximum throughput achieved in all cases is smaller than the available capacity. The loss in throughput is due to various protocol (TCP, IP, AAL5, LLC, SONET) overheads, additionally for cases LAN.B, LAN.C, WAN.B & WAN.C due to some sources waiting for small periods of time before they could begin transmission. Also, in all WAN cases the decision of not using the window scaling option for TCP maximum window size resulted in severe underutilisation of the network. The last column shows the amount of buffer at which zero cell was observed for a particular network configuration. Cases LAN.C & WAN.C required slightly larger buffer compared to LAN.B & WAN.B respectively because of the EPD requirement of setting out a small overflow buffer space beyond the threshold. The maximum window size in WAN cases is kept to a value more suitable for LAN and as a result full exploitation of WAN pipe capacity is not possible by individual TCP sources. However, this allowed multiple TCP sources share the WAN pipe capacity in a less contentious manner.

The association between ATM and TCP layer congestion is first evaluated using the indicator PLCAF. PLCAF measured using the aggregate CLR and aggregate PRR for various switch buffer sizes and for all three cases are shown in Figure 3 for both LAN and WAN scenarios. In LAN.A, PLCAF is almost uniform except for a small blip at around the buffer size of 800 cells (Figure 3a). The relatively small PLCAF

³For a fuller description of the experiments and results please see:

ftp://agora.leeds.ac.uk/scs/doc/reports/1997/97_42.ps.Z.

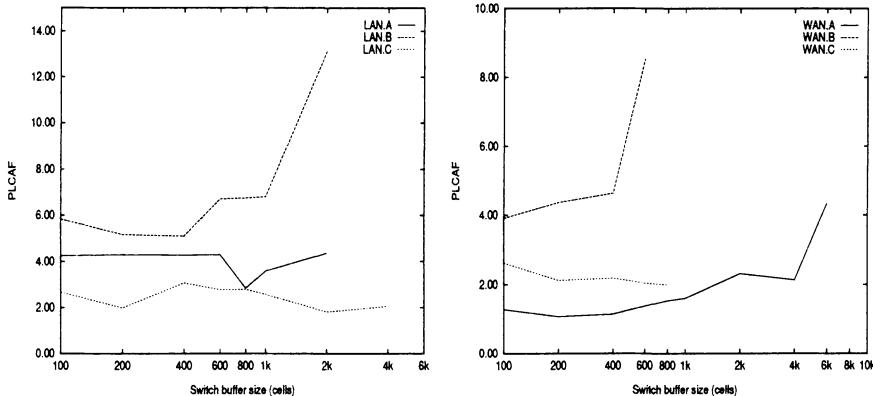


Figure 3: PLCAF as function of switch buffer size in (a) LAN (left) and (b) WAN (right) scenarios.

values in LAN.A suggest that cell losses are taking place in relatively longer bursts. This and the near uniform behaviour of PLCAF against switch buffer size in case LAN.A is due to strong source correlation. Aggregate PLCAF in LAN.B on the other hand increases with switch buffer sizes. A sharp rise in PLCAF is especially observed at buffer size of 2000 cells, the maximum buffer size beyond which no cell loss was noticed. This rise in PLCAF with switch buffer size is explained by the observation that with increasing switch buffers the cell loss phenomenon become smaller in burst size and more distributed along the traffic stream and connections. The ability of EPD to maintain strong association between CLR and PRR is seen from LAN.C data. The effectiveness of EPD is confirmed by the consistently small and fairly uniform PLCAF values in LAN.C. EPD is able to string together dispersed cell loss phenomenon spread across the traffic stream as observed in LAN.B into loss of contiguous cells affecting only a few packets. However, EPD is seen not to exploit the relaxed buffer constraints fully. It was expected that with EPD, PLCAF would show a downward trend with increasing switch buffer sizes. The greedy nature of TCP sources along with the aggressive growth of TCP window sizes are thought to be responsible for this less than ideal behaviour.

The PLCAF curves showing the correspondence between ATM and TCP layer congestion measures for WAN simulation are shown in Figure 3b. WAN.A shows slow growth of PLCAF with switch buffer size unlike the case in LAN.A where PLCAF is found to be uniform. WAN.B mirrors the LAN.B behaviour of PLCAF of abruptly rising to a high value at very relaxed switch buffer constraint. WAN.C on the other hand show decline of PLCAF with increasing switch buffer size as is expected with EPD. Under the relaxed WAN network pipe capacity considered, EPD is found to exploit the relaxed switch buffer constraints. However, unlike in LAN simulation PLCAF values for WAN.C is numerically larger than those for WAN.A. But, the numerical values of PLCAF for all three are smaller in WAN than in LAN

Table 2: CLR-PRR correlation coefficient in LAN scenario

<i>Switch Buffer Size</i>	<i>LAN.A</i>	<i>LAN.B</i>	<i>LAN.C</i>
100	0.944	0.893	0.987
200	0.980	0.974	0.996
400	0.978	0.999	0.965
600	0.998	0.819	0.857
800	0.941	0.999	0.960
1000	0.999	0.953	0.995
2000	0.951	0.957	0.969
4000	-	-	0.988
Aggregate	0.951	0.116	0.883

(Figure 3a & b). This different behaviour may be explained by the fact the sources are unable to expand their window sizes sufficiently large to fill the network pipe in the WAN scenario considered. As a result, the network is underutilised. Sources, at full speed transmission act as an on-off source, period of activity followed by period of inactivity when it is waiting to receive acknowledgement from the receivers. As a result, the sources on the whole tend to be more nicely behaved and less aggressive against each other. Because of the on-off like behaviour of the sources whose transmission start times are interleaved, when cell loss take place they are more likely to be in the form of longer bursts than before resulting in smaller PLCAF.

In the remainder of this section the interaction between ATM and TCP layer congestions using the CLR-PRR correlation coefficient is examined. This analysis for the LAN scenario is shown in Table 2. For each switch buffer size, the correlation coefficient between the set of CLRs of 5 connections and their corresponding PRRs is calculated. For all three cases, very strong CLR-PRR correlation is observed for all different buffer sizes. This is due to the sharing of common switch resources by all 5 connections, thus their behaviours are inherently related. The last row in Table 2 shows the aggregate CLR-PRR correlation coefficients. These are calculated using the set of aggregate CLRs for all buffer sizes (100, 200, 400, 600, 800, 1000 and 2000 cells & also 4000 cells for LAN.C) and the set of corresponding aggregate PRRs. LAN.A shows strong aggregate CLR-PRR correlation while in LAN.B, CLR & PRR appear to be quite independent variables. The strong correlation in LAN.A stems again from the synchronous start times of the sources and matches with the previous PLCAF analysis. While the almost independent CLR-PRR relationship in LAN.B is due to the short bursty cell loss phenomenon in indiscriminate form. Longer bursts of cell losses (maintaining the same or higher CLR) is expected to have stronger CLR-PRR correlation and smaller PLCAF. This is in evidence in LAN.C.

Finally, the association between ATM and TCP layer congestion for WAN scenario is analysed using the CLR-PRR correlation coefficients (Table 3). CLRs and PRRs of competing connections at any particular switch buffer size show strong correlation as before. However, the aggregate CLR and aggregate PRR in WAN.A

Table 3: CLR-PRR correlation coefficient in WAN scenario

<i>Switch Buffer Size</i>	<i>WAN.A</i>	<i>WAN.B</i>	<i>WAN.C</i>
100	0.999	0.960	0.964
200	0.999	0.830	0.942
400	0.999	0.955	0.908
600	0.998	1.000	0.995
800	0.996	-	1.000
1000	0.992	-	-
2000	0.961	-	-
4000	0.956	-	-
6000	0.853	-	-
Aggregate	0.873	0.997	0.983

now show slightly weaker correlation as suggested by PLCAF curve in Figure 3b. Also, very much unlike the LAN scenario, WAN.B shows very strong aggregate CLR-aggregate PRR correlation. Even though the PLCAF vs. switch buffer size plot for WAN.B appear to suggest otherwise, it was observed that the both PRR and CLR exhibit similar behaviour with changing switch buffer size and as a result show strong correlation between aggregate-CLR & aggregate-PRR in WAN.B. As expected WAN.C shows strong aggregate CLR-aggregate PRR correlation.

6 Conclusions & future work

CLR-PRR correlation coefficient and PLCAF are proposed to study the association between user perceived congestion at the TCP layer and the hidden, underlying congestion at the ATM layer. When taken an aggregate view of all connections over a number of different operating environments (changing switch buffer size), CLR and PRR are strongly correlated for sources with synchronous start times in both LAN and WAN scenarios. A similar, characteristic was also observed for sources with asynchronous start times and the ATM switches equipped with EPD & PPD cell drop policies. Using the PLCAF formulation, the association between congestion measures at ATM and TCP layers are analysed. The PLCAF analysis successfully highlighted the effectiveness of intelligent cell drop policies like EPD & PPD in terms of strength of association between CLR & PRR congestion measures.

The simulation study has shown the efficacy of both these indicators in assessing the congestion interaction. It is worthwhile to investigate effect of scaling the number of sources and also the interaction between the rate-based flow control of ABR and the window-based congestion control by TCP in TCP over ATM-ABR environment using the two indicators CLR-PRR correlation and PLCAF.

References

- [1] ITU-T Recommendation I.731, Types and General Characteristics of ATM Equipment. 1996.
- [2] Baumann, M., *Analysis of Cell and Frame Loss Ratios in an ATM Multiplexer*, Proc. 6th Open Workshop on High Speed Networks, pp. 31-38, Stuttgart, Germany, 8-9 Oct., 1997.
- [3] Goyal, R., Jain, R., Kalyanaraman, S. and Fahmy, S., *UBR+: Improving Performance of TCP over ATM-UBR Service*, ICC'97, Montreal, Canada, 8-12 June, 1997.
- [4] Jacobson, V., *Congestion Avoidance and Control*, SIGCOMM'88 Symposium, pp. 314-329, August 1988.
- [5] Kalampoukas, L. and Varma, A., *Performance of TCP over Multi-hop ATM Networks: A Comparative Study of ATM Layer Congestion Control Schemes*, Technical Report - UCSC-CRL-95-13, Computer Engineering & Information Sciences, University of California, Santa Cruz, CA 95064 USA, February 1995.
- [6] Kalyanaraman, S., Jain, R., Fahmy, S., Goyal, R., Lu, F. and Srinidhi, S., *Performance of TCP/IP over ABR Service on ATM Networks*, Globecom'96, London, November 1996.
- [7] Kara, M. and Rahin, M.A., *Towards a Framework for Performance Evaluation of TCP Behaviour over ATM*, ICCC'97 - 13th Intl. Conf. on Computer Communication, pp. 49-60, France, Nov. 18-21, 1997.
- [8] Manthorpe, S., *STCP User Manual*, Tech. Rep. 1.0, LRC-DI-EPFL, <http://lrcwww.epfl.ch/~manthorp/stcp>. 1996.
- [9] Moldeklev, K. and Gunningberg, P., *Deadlock situations in TCP over ATM*, 4th Intl. IFIP workshop on Protocols for high speed networks, pp. 219-235, Vancouver, B.C., Canada, 10-12 August 1994.
- [10] Romanow, A. and Floyd, S., *Dynamics of TCP Traffic over ATM Networks*, ACM SIGCOMM Computer Communications Review, October 1994.
- [11] Stevens, R.W., *TCP/IP Illustrated, Volume 1: The Protocols*, Addison-Wesley Publishing Co., New York 1994.
- [12] Takine, T., Suda, T. and Hasegawa, T., *Cell Loss and Output process Analysis of a Finite-Buffer Discrete-Time ATM Queueing System with Correlated Arrivals*, Proc. IEEE INFOCOM '93, pp. 1259-1269, San Francisco, 1993.