Design & Performance Study of a Flexible Traffic Shaper for High Speed Networks

S. Radhakrishnan *S. V. Raghavan,*
Department of Computer Science and Engineering,
Indian Institute of Technology, Madras, India
Email: radha,svr@iitm.ernet.in

Ashok K. Agrawala
Department of Computer Science,
University of Maryland,
College Park, MD, USA
Email: agrawala@cs.umd.edu

Abstract

In networks supporting distributed multimedia, maximizing bandwidth utilization and providing performance guarantees are two incompatible goals. Heterogeneity of the multimedia sources calls for effective congestion control schemes to satisfy the diverse Quality of Service(QoS) requirements of each application. These include admission control at connection set up, traffic control at the source ends and efficient scheduling schemes at the switches. The emphasis in this paper is on traffic control at the source end.

Traffic control schemes have two functional roles. One is traffic *enforcement* as a supplement to the admission control policy. The other is *shaping* the input traffic so that it becomes amenable to the scheduling mechanism at the switches for providing the required QoS guarantees. Studies on bursty sources have shown that burstiness promotes statistical multiplexing at the cost of possible congestion. Smoothing the traffic helps in providing guarantees at the cost of bandwidth utilization. The need for a flexible scheme which can provide a reasonable compromise between the utilization and guarantees is imminent.

We present the design and performance study of a flexible traffic shaper which can adjust the burstiness of input traffic to obtain reasonable utilization while maintaining statistical service guarantees. The performance of the traffic shaper for bursty sources is studied using simulation.

1 INTRODUCTION

Gigabit speeds have paved the way for many exciting multimedia applications, such as teleconferencing and real-time distributed computing, to be supported on computer networks. Most of these new applications are characterized by stringent QoS requirements in terms of throughput, delay, jitter and loss guarantees. The heterogeneity of the sources calls for effective congestion control schemes to satisfy the diverse Quality of Service (QoS) requirements of each application. These include admission control, traffic enforcement and shaping at the edges of the network and multiclass scheduling schemes at the intermediate switches. Some of the admission control, resource reservation and scheduling schemes proposed for integrated broad band networks in the recent past and the related issues are surveyed in a previous paper [RR94].

Traffic enforcement schemes have a vital role in any resource sharing environment. This is due to the fact that the users may, inadvertently or otherwise, attempt to exceed the rates specified at the time of connection establishment. Many policing schemes viz., Leaky Bucket (LB), Jumping Window (JW), Moving Window (MW), Exponentially weighted moving average Window (EW) and associated variations have been proposed and analyzed [Tur86, SLCG89, ELL90, Rat91]. Studies on bursty sources have shown that burstiness promotes statistical multiplexing at the cost of possible congestion. Smoothing the traffic helps in providing guarantees at the cost of bandwidth utilization. The need for a flexible scheme which can provide a reasonable compromise between the utilization and guarantees is imminent.

This paper describes the design and performance results of a flexible traffic shaper which can provide a variable burstiness at its output. A preliminary version of the basic scheme was presented in [RRA95a]. In our scheme, *the decision to admit an arriving packet is based on the temporal image of the past data maintained in a shift register. To achieve this, a window based enforcement scheme is employed.* A single sliding window mechanism for traffic shaping was incorporated for traffic regulation by Rigolio and Fratta in [RF91]. In that paper, the shaper consisted of a sliding window followed by a server operating at a constant rate. Mukherjee et. al in [MLF92] describes a dynamic time window scheme for end to end congestion control for data traffic. *Our scheme employs more than one window, which jointly provide a more general control over the burstiness of the input stream. The effect of the window parameters on delay distribution and the loss probabilities for varying source burstiness is studied. The sensitivity of output burstiness, delay and loss to window parameters is demonstrated.*

The rest of the paper is organized as follows:
Section 2 describes the relationship between source burstiness and bandwidth requirement for specified QoS requirements. A quantitative means of representing the smoothness of a general packet stream is proposed. The motivation leading to the new scheme is derived by observing the smoothness provided by the leaky bucket with a peak rate policer (LBP). LBP smoothness is characterized in Section 3. The next two sections describe our shaping scheme with the adjustable burstiness feature. Performance is studied through simulation and the results and inferences are presented in Section 6. Section 7 concludes this paper.

2 BURSTINESS AND BANDWIDTH ALLOCATION

2.1 Introduction

Traffic sources in multimedia applications can be basically classified into five categories, *viz.*, data, voice, video, image and graphics. Data sources are generally bursty in nature whereas voice and video sources can be continuous or bursty, depending on the compression and coding techniques used. Continuous sources are said to generate constant bit rate (CBR) traffic and bursty sources are said to generate variable bit rate (VBR) traffic.

A CBR source needs peak rate allocation of bandwidth for congestion-free transmission. For a VBR source, average rate of transmission is generally a small fraction of the peak rate. Thus a peak rate allocation would result in reduced utilization of the system resources. With peak rate allotment, providing performance guarantees is easy. On the other extreme, average allotment may lead to buffer overflows and consequent losses/delays. No meaningful guarantees can be offered in such cases. An effective bandwidth, whose value lies between the average and the peak rate is determined for the various sources [GAN91]. An allocation corresponding to the effective bandwidth optimizes the network utilization and performance guarantees. An allocation nearer to the peak rate allows for tighter probabilistic guarantees. In the extreme, with peak rate allotment, the guarantees can be deterministic.

2.2 Bursty Model and Bandwidth Requirement

The source model used in this paper for measuring performance is the ON-OFF bursty model [SAG94, DYH93, BS91]. On-Off model is characterized by interspersed ON and OFF periods each exponentially distributed with mean T_{ON} and T_{OFF} respectively. During an ON period, cells are periodically transmitted at peak rate λ_p(intercell time during an ON period is $\tau_p = 1/\lambda_p$). The average rate λ_a for this model is $\lambda_p \cdot T_{ON}/(T_{ON} + T_{OFF})$ and the burstiness $\hat{r} = (T_{ON} + T_{OFF})/T_{ON}$. The effective bandwidth requirement for this source λ_{eff} is such that $\lambda_a \leq \lambda_{eff} \leq \lambda_p$.

The ON-OFF bursty model can be justifiably used in modeling many of the sources, currently of interest in multimedia networks. For example, voice sources using talkspurt and video sources after compression and coding, generate bursty streams. Since voice and video sources are *basically* of the CBR type, cell generation during ON period is periodic in nature. To model a generalized data source, as in the case of a large data file transfer application, the ON-OFF model can be modified to make the ON period intercell times exponentially distributed. This assumption will result in an Interrupted Poisson Process(IPP). Further generalizations will lead to 2-state and n-state Markov Modulated Poisson Process(MMPP) models [HL86].

2.3 Defining smoothness for a general stream

In order to compare the proposed scheme with other enforcement schemes, we define the smoothness of a traffic stream as follows:

Definition A generalized packet stream is defined to be $< n_1, T_1; n_2, T_2; ..; n_k, T_k >$ smooth if,
over any time window of duration T_1, no: of packets $\leq n_1$ *and,*
over any time window of duration T_2, no: of packets $\leq n_2$ *and*
$\vdots$
over any time window of duration T_k, no: of packets $\leq n_k$,

where, k denotes the number of windows for characterizing the smoothness of the stream. A larger k can provide a more flexible description of the stream.

3 GENERAL MODEL FOR TRAFFIC SHAPING

A general framework for studying the performance of a traffic shaper is presented in this section. Source is characterized by a peak rate λ_p, an average rate λ_a and mean ON duration T_{ON}. We assume that the network access link at the output of the traffic shaper has a capacity equal to the peak rate of the source stream. Thus any burst arrival is serviced fastest at the peak rate. A traffic shaper which closely fits the model above is the

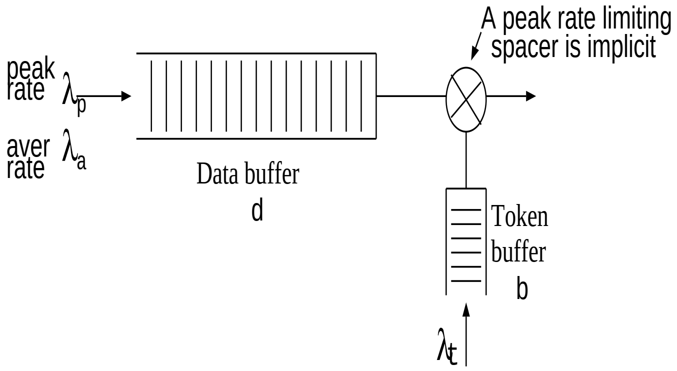

Figure 1: Leaky Bucket with Peak Rate Policer(LBP).

Leaky Bucket with a Peak rate Policer(LBP). In the following sections, we first describe the characteristics of a LBP and then state how possible modifications of the characteristics motivated the development of our scheme.

3.1 Leaky Bucket scheme

Leaky Bucket [Tur86] and its variant schemes are described in [SLCG89, ELL90, Rat91]. In a generalized model of the leaky bucket shown in Figure 1, tokens are generated at a fixed rate as long as the token buffer of size b is not full.

When a packet arrives from the source, it is released into the network only if there is at least one token in the token buffer. This scheme enforces the token arrival rate λ_t on the input stream. Clearly, λ_t should be greater than the average arrival rate λ_a for stability and less than the peak arrival rate λ_p for achieving bandwidth utilization. An input data buffer of size d permits statistical variations. An arriving packet finding the input buffer full is said to be a violating packet and can be dropped or tagged for a preferential treatment at the switching nodes. In this paper, we assume that a peak-rate limiting spacer is an integral part of the leaky bucket mechanism. When a burst of data arrives at the input, even if enough tokens are present, the packets are not instantaneously released into the network. Successive packets are delayed by τ, the transmission time at the negotiated peak rate λ_p, where $\tau = 1/\lambda_p$.

The output of the leaky bucket is characterized as follows:

1. **maximum burst size**: For the LBP, maximum burst size at the output is $b' = b/(1 - \lambda_t/\lambda_p)$, obtained as follows. If we assume the largest burst starts at t_1, the token buffer should be full at t_1. This would be possible only if the source generated an input burst after a prolonged OFF period of b/λ_t, where b is the token buffer size. Since the burst service is not instantaneous due to peak rate policer, more tokens may arrive during the consumption of the existing tokens. Since tokens are removed at λ_p and arrive at λ_t, the instantaneous token count in TB will be $b(t) = b + (\lambda_t - \lambda_p) \cdot t$ and hence TB empties at time $b/(\lambda_p - \lambda_t)$. The maximum burst size b' then becomes $b/(1 - \lambda_t/\lambda_p)$.

2. **long term output smoothness**: over a large time duration T, no: of packets sent out by the leaky bucket, N(T) is $\leq \lambda_t \cdot T = n_t$.

   This relationship is also true for any time duration T' *starting from zero or any epoch when token buffer becomes empty*, if the token buffer is assumed to be empty at $t = 0$.

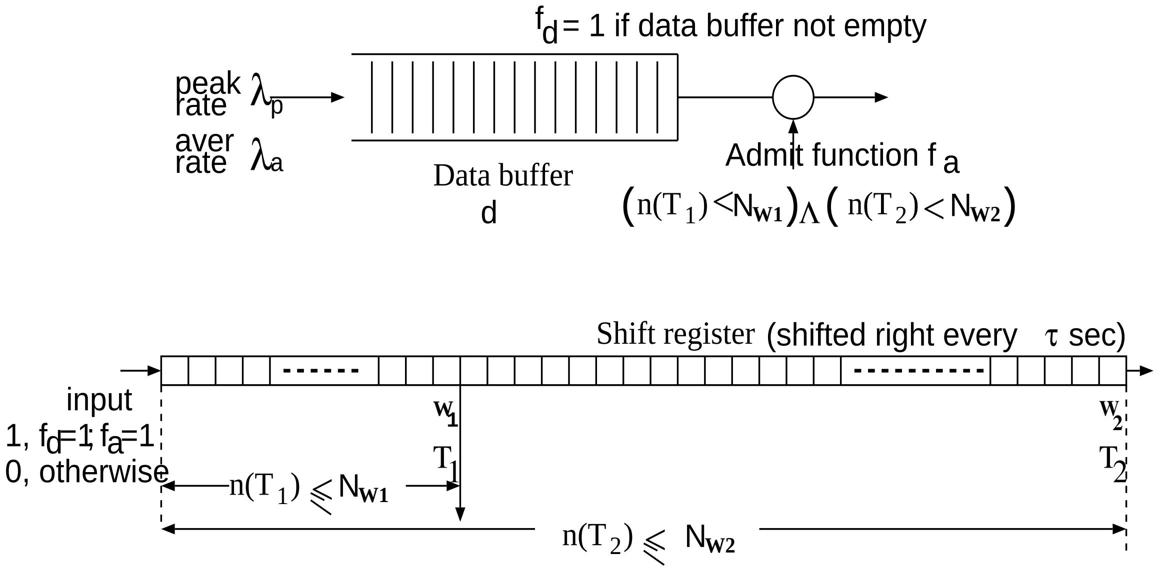

Figure 2: Shift Register Traffic Shaper(SRTS).

3. **short term burstiness**: Over durations smaller than T mentioned in the previous item and exceeding the maximum burst size, leaky bucket output can be modeled as a Linear Bounded Arrival Process(LBAP) with parameters (σ, ρ) [Cru91]. Here, σ represents the maximum burst size b' and ρ represents the token rate λ_t.

Thus the leaky bucket output is (n_t, T) smooth for any time duration T *starting from time 0* .

4 PROPOSED SCHEME

4.1 Motivation for the new scheme

We have seen that in LBP policer, no: of packets over any time duration T starting from 0 is bounded by $\lambda_t \cdot T$. One possible modification to this boundedness is as follows.

- Over any predecided time duration of value T_1 (constant), we can bound the number of packets as in the LB case.
- Over sub-durations within T_1, we can allow more burstiness, *of course*, controlled and within bounds.

The advantage of permitting controlled burstiness is in improving the statistical multiplexing gain at the switches. *This is of atmost relevance in the current scenario since most of the multimedia traffic sources are bursty in nature.* These include naturally stream based sources which are also rendered bursty by the efficient compression and coding mechanisms employed.

4.2 Description of the Scheme

The proposed traffic shaper which we will call Shift Register Traffic Shaper (SRTS) makes use of the temporal profile [Agr94] of the packet stream admitted by the shaper over the

immediate past N time slots, where a time slot τ refers to the reciprocal of the peak rate. This temporal history can be maintained by a shift register with 1 bit corresponding to every packet sent. The shift register is shifted right every time slot τ. The entry of the bits into the shift register is as per the following;
Let $f_d = 1$ if data buffer is not empty and 0 otherwise;
Similarly, let f_a denote the admit control function defined as
$f_a = (n(T_1) < n_1)$ and $(n(T_2) < n_2)$ and $(n(T_3) < n_3) \cdots$ depending on the number of windows. Here T_i refers to a time window. The size of the corresponding window is denoted by W_i and maximum number of packets permitted in W_i by N_{Wi} (note that $N_{Wi}{=}n_i$).

The data bit shifted in is 1 if $f_d = 1, f_a = 1$; and 0 otherwise.

Thus the bit contents of the shift register at any instant, provides an image of the history of the packets sent. All the time durations mentioned with reference to the shift register start from the time point corresponding to the entry point of the shift register. To determine the number of packets in any time duration, a counter is used. It increments whenever a '1' enters the shift register and decrements when a '1' shifts out of the right edge of the corresponding window monitored by the counter.

Figure 2 describes an enforcement scheme using two windows. This scheme generates an $(n_1, T_1; n_2, T_2)$ smooth traffic, which means that over any period of duration T_1, the number of packets $n(T_1) \leq n_1$ and over any period of duration T_2, the number of packets $n(T_2) \leq n_2$. Though we have described the scheme with two windows, further flexibility in moulding the burstiness is possible using the appropriate number of windows. Since the restriction on the number of packets permitted in a time window is enforced at the entry point of the shift register and the window shifts to the right every τ seconds, the smoothness is guaranteed over *any time window over the entire duration of the connection.*

One limitation of the above scheme is caused by the discretization of time into slots of τ. A slot of is termed active if a cell is transmitted during that slot and idle, otherwise. Since the cell arrival instant need not synchronize with the output slots, a cell arriving during an idle slot will have to wait till the end of that slot for transmission. This limitation is removed in our current scheme by using 'soft' discretization. If a cell arrives during an idle slot, say after τ' elapses (out of τ), idle slot is frozen and an active slot is initiated immediately. At the termination of this active slot, if either data is absent or the admit function is false, the residual idle slot of duration $(\tau - \tau')$ commences. The end of a slot is indicated by the timer interrupt in Figure 3. The shift register is shifted right at the end of every slot, active or passive. The essence of the above arrangement is that an idle slot is interruptible whereas an active slot is not. Every time an idle slot is interrupted, the residual idle time is saved for future use up.

The modification described above is illustrated as an FSM in Figure 3.

The key features are:

- Idle to Active state transition is fired by the event $(f_a \wedge f_d)$ where f_a: admit function and f_d : data present flag.
  The following actions ensue:

  1. save residual time by freezing the counter.
  2. initiate transmission and go to active state.
  3. every slot timer interrupt in idle state will cause transition to itself after resetting the counter.

- Active to Idle state transition is fired by the timer interrupt.

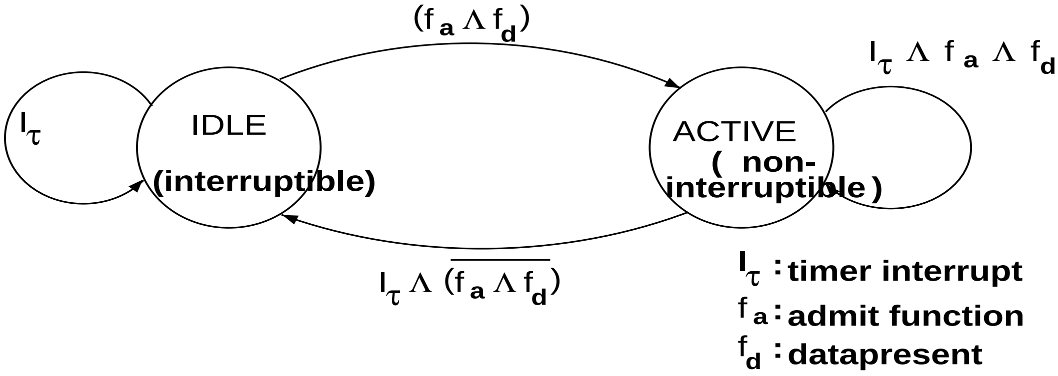

Figure 3: FSM describing the transitions between idle and active states.

1. if $((f_a \wedge f_d) = 1$, initiate another active slot.
2. else initiate an idle slot and go to idle state.

4.3 Choice of Windows

The shaping parameters of the proposed scheme are the window sizes W_1, W_2, W_3 (for a 3 window case) and the maximum number of packets permitted in each window N_{W1}, N_{W2} and N_{W3}. The window parameters can be derived from the key observations made earlier regarding the LBP scheme. For restricting the size of the maximum burst at the output, Window-1 parameters are chosen as $W_1 = N_{W1} = b'$ where b' is the maximum burst size.

Window-3 parameters can enforce the average policing characteristics exhibited by the LBP over large time durations. If λ_{eff} is the effective bandwidth allotted for the bursty source(λ_p, λ_a), then the token arrival rate λ_t of the equivalent leaky bucket should be equal to the effective bandwidth. Thus the window parameters are chosen as follows:
for W_3 = large value T, $N_{W3} = \lambda_{eff} \cdot \tau \cdot W_3$.

Window-2, the main control parameter of the shaper can be suitably tuned to incorporate the burstiness control feature. If we assume a LBAP(σ, ρ) for the output of the LBP over durations larger than and of the order of maximum burst size, σ will be b' and ρ equals λ_t. Then for a chosen value of W_2, $N_{W2} = b' + \lambda_t \cdot (W_2 - W_1) \cdot \tau$.

Example For a bursty model with mean ON period of 200msec, intercell time τ of 10 msec and burstiness 5, $\lambda_p = 100$ and $\lambda_a = 20$.
If we choose λ_{eff} to be 40, for a bucket size(of an equivalent LBP) of 18,
max burst size $b' = b/(1 - \lambda_t/\lambda_p) = 30$. Thus $W_1 = N_{W1} = 30$. For $W_2 = 75$, $N_{W2} = 30 + 45 \cdot 40/100 = 48$. W_3 corresponds to the large duration over which the average policing is enforced. For a choice of $W_3 = 450$, $N_{W3} = 450*40/100 = 180$.
The exact choice of W_2 and W_3 is currently arbitrary and can be tailored to suit the application stream. The only criteria is that over W_2 , we assume the equivalent LBP to generate a LBAP stream whereas over the larger window W_3 , an averaging property is expected. The influence of the source leading to a judicious choice of W_2 and W_3 is yet to be investigated.

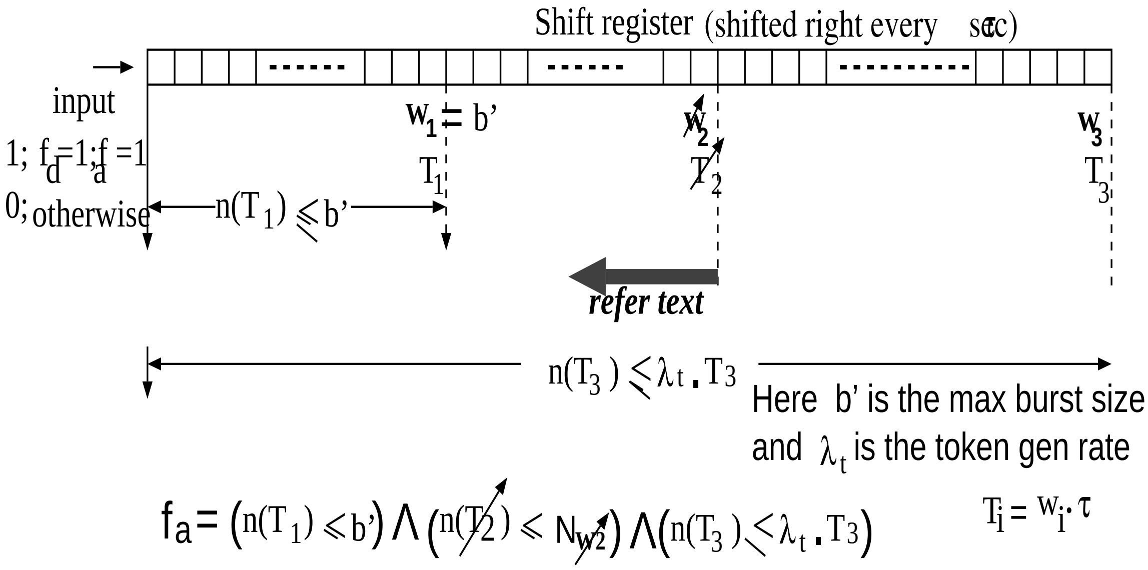

Figure 4: The Proposed Scheme(SRTS) with variable burstiness.

5 PROVIDING ADJUSTABLE BURSTINESS

A traffic shaper with controllable burstiness can facilitate statistical multiplexing along with reasonable performance guarantees. The manner in which the variable burstiness feature can be incorporated in SRTS is described in this section.

A LBP has essentially 2 parameters. The bucket size b which decides the maximum burst size and the token arrival rate λ_t which provides a measure of the effective bandwidth allotted to the source. The model proposed in this paper has 3 parameters. One window , W_1 which limits the maximum burst size and a second window (W_3) for long term average policing correspond conceptually to the two LBP parameters. The third window, namely W_2, is the one for providing the variable burstiness feature.

The proposed scheme with the variable burstiness feature is schematically illustrated in Figure 4. The region of operation to permit higher burstiness is shown by the shaded arrow. An adjustable burstiness can be provided by the following choice of SRTS parameters.

1. The parameters of the smallest window T_1 are chosen as $N_{W1} = b'$ and $W_1 = b'$. This bounds the maximum burst size.
2. Over the largest window W_3, we enforce the LB smoothness, namely $N_{W3} \leq \lambda_t \cdot T_3$, where λ_t, the token generation rate of the analogous LB, is the long term policed rate of the stream.
3. The intermediate window W_2 can be adjusted to permit more burstiness than that allowed by LB strategy as described in Section 4.3.

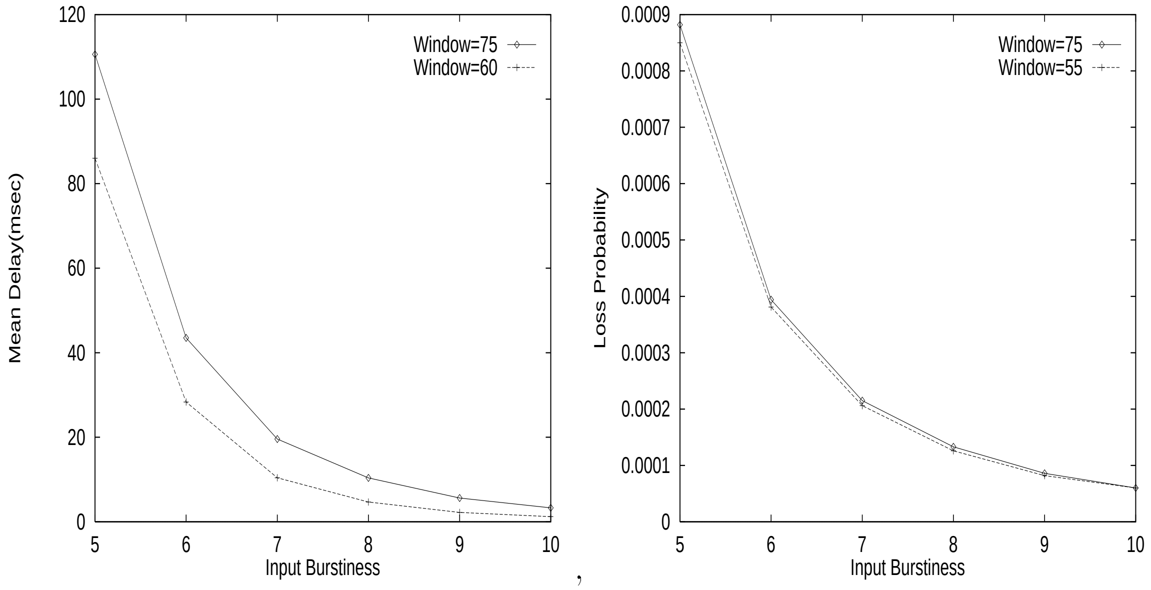

Figure 5: (a) Mean Delay and (b) Loss vs input burstiness.

6 PERFORMANCE STUDY AND RESULTS

6.1 Simulation Experiments

The performance of the proposed traffic shaper is analyzed through simulation. As mentioned in Section 2.2, the source is assumed to be of ON-OFF bursty type. Three simulation experiments are performed as detailed below. In all the cases, $W_1 = N_{W1} = 30; W_3 = 450, N_{W3} = 180$; $N_{W2} = 48$; Each simulation run is performed for 10^7 packets.

Experiment 1 The delay characteristics of the traffic shaper is studied as a function of the input burstiness for different window parameters. Size of data buffer is very large to keep losses close to zero. The input burstiness is varied by adjusting the ON period, keeping the OFF period constant. Intercell time is 10msec and hence $\lambda_p = 100$. Since the long term average policed rate is λ_t, the range of ON period variation is such that λ_a remains $\leq \lambda_t$ for stability. Thus $(T_{ON}/(T_{ON}+T_{OFF}) \cdot 100) < \lambda_t$, which is fixed at 40. Input burstiness is varied from 5 to 10 by keeping the OFF period constant at 800 msec and adjusting the ON period. Figure 5a gives the delay distribution for window sizes of 75 and 60. The number of simulation runs are such that the results are accurate to within 5% with 95% confidence level.

Experiment 2 The loss characteristics incurred by the SRTS is studied in this experiment. Data buffer size is finite. In this case, the input burstiness is varied by keeping the ON period constant at 200 msec and varying the OFF period.

Experiment 3 For the same source burstiness, we study the sensitivity of output burstiness variation at the output stream for different window parameters. Since the output stream is of an arbitrary nature unlike the input stream which is described by a bursty ON-OFF model parameters, we use ratio of Variance to Mean of cell interarrival times [SW86, HA87] for characterizing the burstiness. We will use the term 'burst factor' for this ratio to

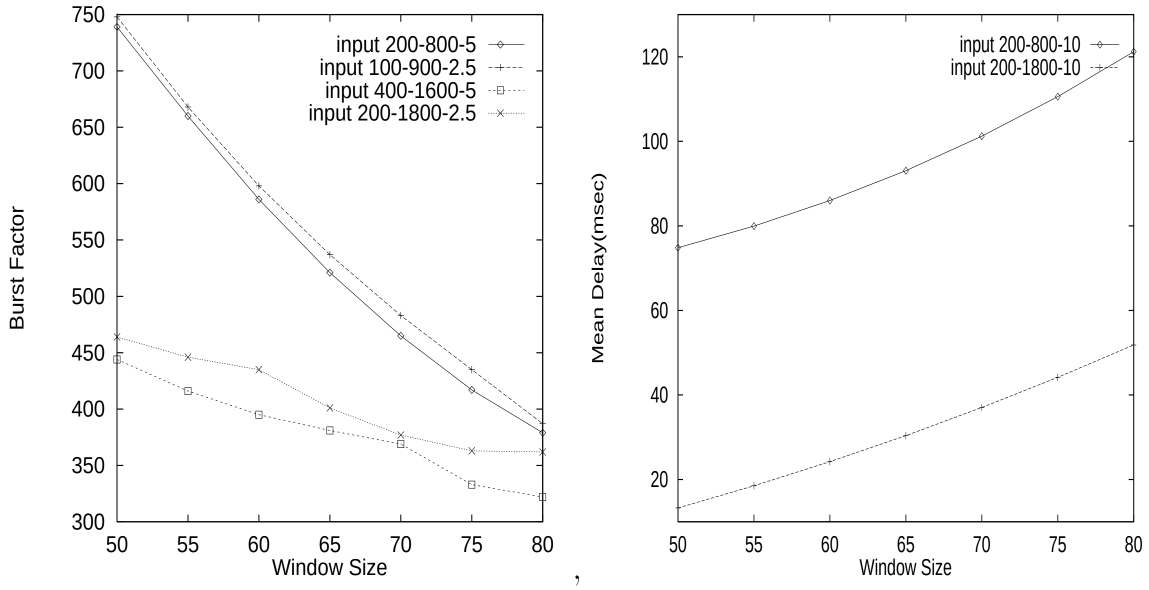

Figure 6: (a) Output burst factor and (b) Mean Delay vs window size.

differentiate this definition of burstiness from the definition given in Section2.2. Figure 6a presents the result for different source ON-OFF characteristics. In each case, the source parameters are such that the average rate is $< \lambda_t$, which is fixed as 40.

Figure 6b illustrates the effect of window size on mean delay. The number of simulation runs are such that the results are accurate to within 5% with 95% confidence level.

6.2 RESULTS AND INFERENCE

Results of the simulation and inferences drawn, thereof, are as follows.

1. Increase in input burstiness (as defined in Section 2.2) causes a reduction in the mean delay. This is expected since a larger burstiness implies a shorter source active period for a constant OFF period. As can be seen in Figure 5a, a smaller window size W_2 for the same N_{W2} admits burstier streams than would be admitted by a correspondingly larger window size for the same N_{W2}.

2. For the finite buffer case, the loss characteristics are presented in Figure 5b. For reasons similar to the results in the previous experiment, a smaller window reduces the losses. The difference is however not as much pronounced as in the previous case.

3. The output burst factor variation demonstrated in Figure 6a is a significant result in concurrence with our concept of a 'controllable' burstiness. A shaper with a larger control window size generates a smoother output stream. We believe that the burstiness of the output can be tuned to provide statistical multiplexing gains at the switches.

4. The results of Figure 6b provide a means of selecting the window parameters suitable for the delay requirements of the application. By judiciously selecting the window-2 parameters, namely W_2 and N_{W2}, it is possible to tune the shaper behavior based on the application characteristics and the performance requirements. Although the general influence of the parameters is apparent, the precise correspondence between

the source behavior and the window parameters needs to be established for different sources.

7 CONCLUSION

In this paper, we have presented the design and performance analysis of a traffic shaper for high speed networks. The motivation for incorporating the features in the proposed scheme is derived from the leaky bucket characteristics. The major advantage of this approach over the leaky bucket is in permitting controlled burstiness at the output. If the burstiness can be characterized and controlled, we feel that higher statistical multiplexing gain can be derived along with improved QoS guarantees. Simulation results are encouraging.

The proposed scheme can be easily implemented in hardware using a shift register, two counters and the control gates. For highly bursty sources, we believe that the controlled burstiness feature of this scheme can be put to use effectively for increasing the statistical multiplexing gain. Results of the comparison between the SRTS and LBP can be seen in [RRA95b]. A traffic shaper must work in unison with the scheduling schemes at the switches for providing the required utilization and guarantees. We have demonstrated that a controlled burstiness feature can be provided. Our future work will investigate how the characteristics of the shaper output can be fruitfully exploited at the switches to provide utilization and guarantees.

8 REFERENCES

[Agr94] Ashok K. Agrawala. Temporal Profile Capture, 1994. Personal Communication.

[BS91] Jaime Jungok Bae and Tatsuya Suda. Survey of Traffic Control Schemes and Protocols in ATM Networks. *Proceedings of the IEEE*, 79(2):170 – 189, February 1991.

[Cru91] R. L. Cruz. A Calculus for Network Delay, Part I: Network Elements in Isolation. *IEEE Transactions on Information Theory*, 37(1):114–131, January 1991.

[DYH93] Don Towsley David Yates, James Kurose and Michael G. Hluchyj. On per-session end-to-end delay distributions and the call admission problem for real-time applications with QoS requirements. In *Proceedings of ACM SIGCOMM'93*, pages 2–12, Ithaca,N.Y, USA, September 1993.

[ELL90] A.E. Eckberg, D.T. Luan, and D.M. Lucantoni. Bandwidth Management: A Congestion Control Strategy for Broadband Packet Networks-Characterizing the Throughput-burstiness Filter. *Computer Networks and ISDN systems*, 20:415–423, 1990.

[GAN91] R. Guerin, H. Ahmadi, and M. Naghshineh. Equivalent Capacity and its Application to Bandwidth Allocation in High Speed Networks. *IEEE Journal on Selected Areas in Communications*, 9(7):968–981, September 1991.

[HA87] J. Y. Hui and E. Arthurs. A Broadband Packet Switch for Integrated Transport. *IEEE Journal on Selected Areas in Communications*, 5:1264–1273, October 1987.

[HL86] H. Heffes and D. M. Lucantoni. A Markov Modulated Characterization of Packetized Voice and Data Traffic and Related Statistical Multiplexer Performance. *IEEE Journal on Selected Areas in Communications*, 4(6):856–868, September 1986.

[MLF92] A. Mukherjee, L. H. Landweber, and Theordore Faber. Dynamic Time Windows and Generalized Virtual Clock: Combined Closed-Loop/Open-Loop Congestion Control. In *Proceedings of IEEE INFOCOM'92*, pages 322–332, Firenze, Italy, May 1992.

[Rat91] Erwin P. Rathgeb. Modelling and Performance Comparison Of Policing Mechanisms for ATM Networks. *IEEE Journal on Selected Areas in Communications*, 9(3):325–334, April 1991.

[RF91] G. Rigolio and L. Fratta. Input Rate Regulation and Bandwidth Assignment in ATM Networks" an Integrated Approach. In *Proceedings of the ITC-13: Queueing, Performance and Control in ATM*, pages 123–128, Copenhagen, Denmark, June 1991.

[RR94] S. Radhakrishnan and S. V. Raghavan. Network Support for Distributed Multimedia - Issues and Trends. In *Proceedings of SEACOMM'94, International Conference on Communications and Computer Networks*, Kuala Lumpur, Malaysia, October 1994.

[RRA95a] S. Radhakrishnan, S. V. Raghavan, and Ashok K. Agrawala. A Flexible Traffic Shaper with Variable Burstiness for High Speed Networks. In *Proceedings of JENC6 International Conference* , Tel Aviv, Israel, May 1995.

[RRA95b] S. Radhakrishnan, S. V. Raghavan, and Ashok K. Agrawala. Performance Comparison of a New Traffic Shaper and Leaky Bucket for Bursty Real-Time Traffic. Communicated to MMM'95 International Conference on Multi-Media Modelling to be held in Singapore, November 1995.

[SAG94] G. D. Stamoulis, M . E. Anagnoustou, and A. D. Georgantas. Traffic Source Models for ATM Networks: A Survey. *Computer Communications*, 17(6):428–438, June 1994.

[SLCG89] M. Sidi, W. Liu, I. Cidon, and I. Gopal. Congestion Control through Input Rate Regulation. In *Proceedings of GLOBECOM'89*, pages 1764–1768, Dallas, Texas, November 1989.

[SW86] K. Sriram and W. Whitt. Characterizing Superposition Arrival Processes in Packet Multiplexers for Voice and Data. *IEEE Journal on Selected Areas in Communications*, 4:833–846, September 1986.

[Tur86] J. S. Turner. New Directions in Communications (or Which way to the Information Age? *IEEE Communications*, 24(10):8–15, October 1986.

9 BIOGRAPHY

S.Radhakrishnan received B.Sc(Engg) degree in Electronics & Communication with first rank and distinction from University of Kerala in 1982. During 1983, he worked as an Engineer in Vikram Sarabhai Space Center, Trivandrum after which he joined College of Engineering, Trivandrum as a lecturer. In 1988, he received an M.Tech in Computer Science and Engineering from Indian Institute of Technology, Madras. Since 1988, he has been a member of the faculty of Government Engineering College, Trichur. He is currently pursuing his doctoral studies in the Department of Computer Science and Engg, I.I.T Madras. His areas of interest are high speed networks, congestion control and performance evaluation.

S.V.Raghavan is a Professor in the Department of Computer Science and Engineering, Indian Institute of Technology, Madras. He is also the Chief Investigator of the project on Education and Research in Computer Networking jointly sponsored by the Department of Electronics, Government of India, and the United Nations Development Programme. He is presently serving on the Board of Editors of the journal of Institute of Electronics and Telecommunication Engineers(IETE) for Computers and Control. He is also a member of the Editorial Advisory Board for Computer Communications, Butterworth-Heinemann Ltd. His research interests are networks, protocols, multimedia systems and performance.

Ashok K. Agrawala is a Professor of Computer Science at the University of Maryland, College Park. He earned the A.M. and Ph.D. degrees in Applied Mathematics from Harvard University, Cambridge, MA in 1970. He has published on microprogramming, computer architectures, local area networks and pattern recognition. Recently, he has been working on the design problems for hard real-time systems and has developed the system, MARUTI, which addresses the needs of next generation real-time systems. Dr. Agrawala is a fellow of the IEEE and a member of the Association for Computing Machinery, American Association for Advancement of Science and Sigma Xi.