A Flexible Traffic Shaper for High Speed Networks: Design and Comparative Study with Leaky Bucket*

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

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

Abstract

Maximizing bandwidth utilization and providing performance guarantees, in the context of multimedia networking, are two incompatible goals. Heterogeneity of the multimedia sources calls for effective traffic control schemes to satisfy their diverse Quality of Service(QoS) requirements. 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.*

Most multimedia sources are bursty in nature. Traffic shapers have been mainly studied hitherto from the point of view of their effectiveness in smoothing the burstiness. Leaky Bucket(LB) scheme, to cite an example, is a mean rate policer smoothing at the token generation rate. Studies on bursty sources show that burstiness promotes statistical multiplexing at the cost of possible congestion. Smoothing, on the other hand, helps in providing guarantees at the cost of utilization. Thus need for a flexible scheme which can provide a reasonable compromise between utilization and performance is imminent. Recent studies [10, 12] have also questioned the suitability of LB for policing real-time traffic due to the excessive delays. We argue for a policy which is less stringent on short term burstiness than the LB.

We propose a new traffic shaper which can adjust the burstiness of the input traffic to obtain reasonable bandwidth utilization while maintaining statistical service guarantees. The performance study is conducted in two parts. In the first part, we study the effect of varying the shaper parameters on the input characteristics. In the second part, we dimension our scheme and a LB equivalently and compare the mean and peak rate policing behavior with delay and loss as the performance parameters. Adopting a less stringent attitude towards short term burstiness is shown to result in considerable advantage while policing real-time traffic. Future research possibilities in this topic are explored.

Contents

*This is a revised version of the paper presented in JENC6 Conference
†Further correspondence may be addressed to Professor S. V. Raghavan

1 Introduction

Advances in optical transmission media and high speed switching 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, constituted of heterogeneous mix of video, voice and data, 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 meet the diverse Quality of Service (QoS) requirements of each application. These include admission control at connection set up, traffic enforcement and shaping at the edges of the network and multiclass scheduling schemes at the intermediate switches. Latency effects apparent at the gigabit speeds make the conventional feedback techniques ineffective. Thus the responsibility of preventing congestion lies with the admission control and traffic enforcement schemes.

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 [13]. Admission control restricts the number of connections that can be supported by the network. Admission control is decided by an algorithm which expects that the user provides an estimate of the traffic parameters and abides by their negotiated values. Resource reservation schemes manages the allocation of the resources at each of the nodes so that per-node QoS requirements can be met for each connection. Scheduling policies provide sharing of bandwidth among the various classes and the various streams within each class so that the individual requirements can be satisfied.

In a resource sharing packet network, admission control and scheduling schemes by themselves are not sufficient to provide guarantees. 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. Traffic policing schemes proposed in the literature include mainly Leaky Bucket (LB), Jumping Window (JW), Moving Window (MW), Exponential Weighted Moving Average (EWMA) and associated variations. A performance comparison among these schemes from the point of view of violation probability, sensitivity to overloads, dynamic reaction time and worst case traffic admitted into the network can be found in [11]. It has been shown that the LB and the EWMA are the most promising mechanisms to cope with short-term fluctuations and hence suited for policing bursty traffic. Several improvements of the LB has been proposed for increasing utilization in an ATM environment [3, 5, 15]. Traffic enforcement schemes police the source streams to check that their characteristics conform to the declared values throughout the life of the connection. The various schemes have been studied from the point of view of their capability to smooth the burstiness in the source. Traffic Shaping, on the other hand, conditions the input stream so that the characteristics are amenable to the scheduling mechanisms to provide the required QoS guarantees. Although, one may imply the other, there are subtle differences. The former checks the conformance to the declared values whereas the latter shapes it to be more agreeable to the scheduling policies.

Traffic shapers have been mainly studied hitherto from the point of view of their effectiveness in smoothing the burstiness. Leaky bucket scheme, to cite an example, is a mean rate policer smoothing at the token generation rate. Studies on bursty sources show that burstiness promotes statistical multiplexing at the cost of possible congestion. Smoothing, on the other hand, helps in providing guarantees at the cost of utilization. Thus need for a flexible scheme which can provide a reasonable compromise between utilization and performance is imminent. Recent studies [10, 12] have also questioned the suitability of LB for policing real-time traffic. LB, in its attempt to enforce smoothness often introduces excessive access delays thereby making it incapable of regulating real-time traffic. A policy which is less stringent on short term burstiness while bounding long term behavior with a LB-bound would be better suited for time critical traffic. This was the second motivation which led us to the new proposal.

We propose a new traffic shaper which can adjust the burstiness of the input traffic to obtain reasonable bandwidth utilization while maintaining statistical service guarantees. It uses a window based shaping policy which captures the essence of the LB scheme, permits short term burstiness in a more flexible manner and is inherently peak rate enforced. The decision to admit an arriving packet is based on the temporal image of the past data maintained in a shift register. We will refer to the new scheme as the SRTS (Shift Register Traffic Shaper). A single sliding window mechanism for traffic shaping was incorporated for traffic regulation by Rigolo and Fratta in [14]. In that paper, the shaper consisted of a sliding window followed by a server operating at a constant rate. Our scheme employs more than one window, which jointly provide a more general control over the burstiness of the input stream. The motivation for our scheme is derived by studying the characteristics of the traffic generated by the leaky bucket scheme.

The performance characteristics of SRTS is studied in this paper in two parts. In the first part, we investigate the controlling effect of shaper parameter variations on the input traffic characteristics. Delay, loss and burstiness behavior at the output is studied for different window parameters and input burstiness. The adjustable burstiness feature is demonstrated in this study. In the second part, we dimension the proposed SRTS shaper and a LB shaper equivalently and compare the mean and peak rate policing behavior with delay and loss as the performance parameters. Adopting a less stringent attitude towards short term burstiness is shown to result in considerable advantage for policing real-time traffic.

The rest of the paper is organized as follows:
Section 2 discusses qualitatively how burstiness of the source decides the bandwidth that needs to be allocated for specified QoS guarantees. A quantitative means of representing burstiness

bounds is defined. Section 3 presents the general requirements of a traffic shaper and briefly describes LB and EWMA schemes. Effect of shaping on delays and bandwidth requirement is discussed in Section 4. Sections 5 and 6 describe the proposed SRTS scheme and its variable burstiness feature. Section 7 presents the simulation results, observations and inferences. Finally Section 8 summarizes and 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. But we confine our discussion to mainly data, voice and video. 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. Most of the multimedia sources are bursty in nature.

A CBR source needs peak rate allocation of bandwidth for congestion-free transmission. For a VBR source, average rate of transmission λ_a can be a small fraction of the peak rate λ_p. Thus a peak rate allocation would result in gross under 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 λ_{eff}, whose value lies between the average and the peak rates is determined for the various sources [6, 7]. An allocation corresponding to the effective bandwidth optimizes the network utilization and performance guarantees. An allocation nearer to the peak rate allows providing 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 that is used for measuring performance is the ON-OFF bursty model [2, 17, 19]. 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 [8].

In this paper, we use an ON-OFF bursty model for the source. The burstiness can be varied by altering the T_{ON} or T_{OFF} keeping the other constant.

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:

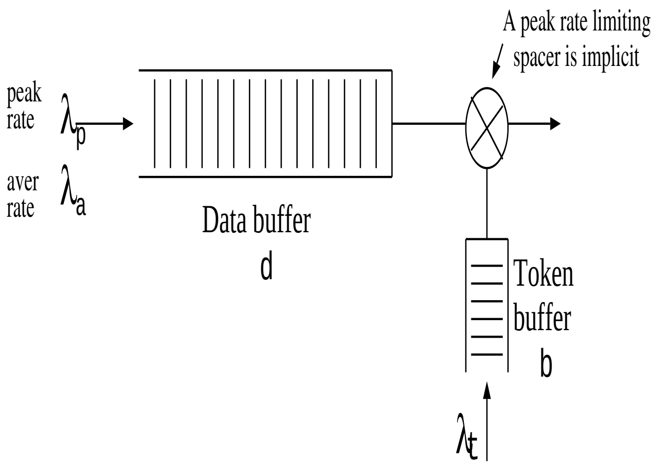

Figure 1: A Generalized Leaky Bucket scheme

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 Leaky Bucket with a Peak rate Policer(LBP). In the following sections, we first describe the characteristics of a LBP output traffic. These characteristics motivated the development of the scheme proposed in this paper. A brief description of EWMA, a window based policer is also given for comparison with the proposed scheme.

3.1 Leaky Bucket Scheme

Leaky Bucket [18] and its variant schemes are described in [3, 5, 11, 15]. 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 negotiated peak rate λ_p, where $\tau = 1/\lambda_p$. We will use LBP to designate the leaky bucket with peak rate policer.

For the leaky bucket parameters defined above, maximum burst size at the output is $b' = b/(1-\lambda_t/\lambda_p)$. This includes the new tokens that arrive during the transmission of the first b packets. 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.* It is assumed here that the token buffer is empty at t=0.

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) [4]. Here, σ represents the maximum burst size b' and ρ represents the token rate λ_t.

In terms of the smoothness definition given in Section 2.3, we can state that for any T starting from 0 (or from any epoch when token buffer is empty), LBP output is (n_t,T) smooth.

3.2 Exponentially Weighted Moving Average Scheme

EWMA is a window based scheme [11] where the maximum number of cells permitted within a fixed time window is limited. If we consider the connection time to consist of consecutive windows of same size, the maximum number of cells accepted in the ith window N_i is a function of the mean number of cells per window N and an exponentially weighted sum of the cells accepted in the preceding windows as given below
$N_i = \frac{(N-(1-\gamma)(\gamma X_{i-1}+....+\gamma^{i-1}X_1))-\gamma^i S_0}{(1-\gamma)}$ where S_0 is the initial value for the EWMA. The weight factor γ decides the number of relevant preceding windows which influence the number of packets permitted in the current window. A nonzero value of γ permits more burstiness. For a value of $\gamma = 0.8$, up to 5 times N number of packets can occur in the first window. Thus a large value of γ increases the reaction time and it is shown in [11] that the dynamic behavior of EWMA is the worst. Moreover, the implementation complexity of this scheme is higher than LB and other window based schemes.

4 Shaping and BW Allocation

The bandwidth that needs to be allocated to the shaped stream depends on the shaper parameters. For instance, a LB produces a stream which requires, at a minimum, bandwidth equal to the token arrival rate, to be allocated at the access multiplexer. A larger token arrival rate reduces the access delay at the policer but needs a larger bandwidth allocation. For a source characterized by a peak rate λ_p and burstiness $\hat{r}$, bandwidth allocation λ_{bw} is such that $\lambda_p/\hat{r} \leq \lambda_t \leq \lambda_{bw} \leq \lambda_p$. At the access multiplexer, the capacity of the output link $\lambda_o = \sum_{i=1}^{m} \lambda_{bw}(i)$ for m streams multiplexed to the same output. Since most multimedia traffic is bursty in nature, a large statistical multiplexing gain is possible only if λ_t is near the average arrival rate $\lambda_a = \lambda_p/\hat{r}$. On the other hand, smaller the λ_t, larger the access delay and/or violation probability incurred by the source. A lenient enforcement policy can increase the delay at the multiplexing/switching nodes due to buffer overflows. Thus there is a trade off between the access delay introduced by the policer and the network delay at the switches. From the end user's point of view, the delay incurred by the application includes the access delay and the network delay. For a constant bandwidth allocation, the effect of input rate control can be summarized by the following observations [10, 12].

1. The total delay experienced by a cell is the sum of the access delay due to queuing at the shaper and the network delay at the switch. The policer simply transfers the network delay on to the input side thereby avoiding overflow losses/delays within the network. Thus unless the source has a large buffer and can tolerate excess delay, the input rate control as performed by the LB can hardly improve the network performance [12]. For many real time applications, this access delay could be prohibitive.

2. A stringent input rate control may unnecessarily increase the user end-to-end delay by a significant amount [12].

3. The minimum total average delay is achieved when no traffic enforcement is invoked [10, 12]. This observation is applicable when the network bandwidth is considerably greater than the source transmission rate, in which case the effect of individual streams is smoothed by statistical multiplexing. Nevertheless, to check excessive burstiness and prolonged rate violations, input policer is practically needed.

It is evident from the aforementioned points that the access delay introduced by the traffic policer can be significant. One way of reducing the access delay would be to permit more short term burstiness subject to:

- the maximum burst size should be bounded and burst arrivals must be peak rate enforced.
- the number of arrivals over a larger time durations to be bounded at the average policing rate.

LB and the EWMA mechanisms perform the above two in different ways. The short term burstiness permitted by the LB is decided by the size of the token buffer b. As explained earlier, over any time duration T starting from 0 (or any epoch when the token buffer becomes empty), the number of packets admitted into the network are bounded by $\lambda_t * T$. With reference to Figure 2 which shows the number of admitted packets versus time, the operating region for LB operation is below the line OA corresponding to the average policing rate. A source is permitted to send a burst only if it remains inactive for a sufficient amount of time to gather enough number of tokens in the token buffer. Thus the operating point is always below the line OA. A well behaved source transmitting uniformly at the token arrival rate will operate along OA.

The short term burstiness in the EWMA mechanism is influenced by the factor γ as described earlier in section 3.2. The dynamic response for the EWMA is however poor for reasonable values of γ. EWMA output is not peak rate enforced and the implementation complexity is also considerable compared to the other schemes.

We describe in the next section a traffic shaper which has the following features:

1. permits short term burstiness but bounds long term behavior so that the number of packets admitted over a long time is same as that admitted by an equivalent leaky bucket.

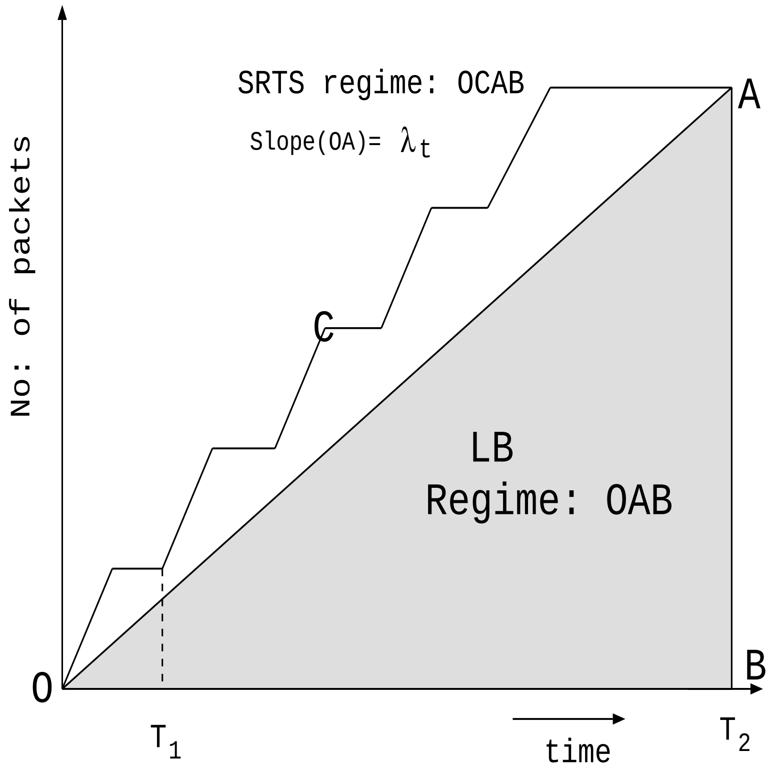

Figure 2: Permitted no: of packets vs time

2. variable burstiness easily incorporated.

3. it is inherently peak rate enforced.

4. it is a window based shaper consisting of two or more windows and the shaper behavior can be more flexibly set unlike the EWMA which has only one control parameter γ.

5. it is designed using a shift register and two counters and hence can easily be implemented in hardware.

5 Shift Register Traffic Shaper (SRTS)

5.1 Motivation for the new scheme

Two basic concepts motivated the development of the SRTS.

1. provide burstiness variation for possible multiplexing gain.

2. reduce the access delays by adopting a less stringent attitude towards short term burstiness following the observations made in Section 4.

These are elaborated below. We have seen that in a Leaky Bucket (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 LBP case.

- Over sub-durations within T_1, we can allow more burstiness, *of course*, with bounds.

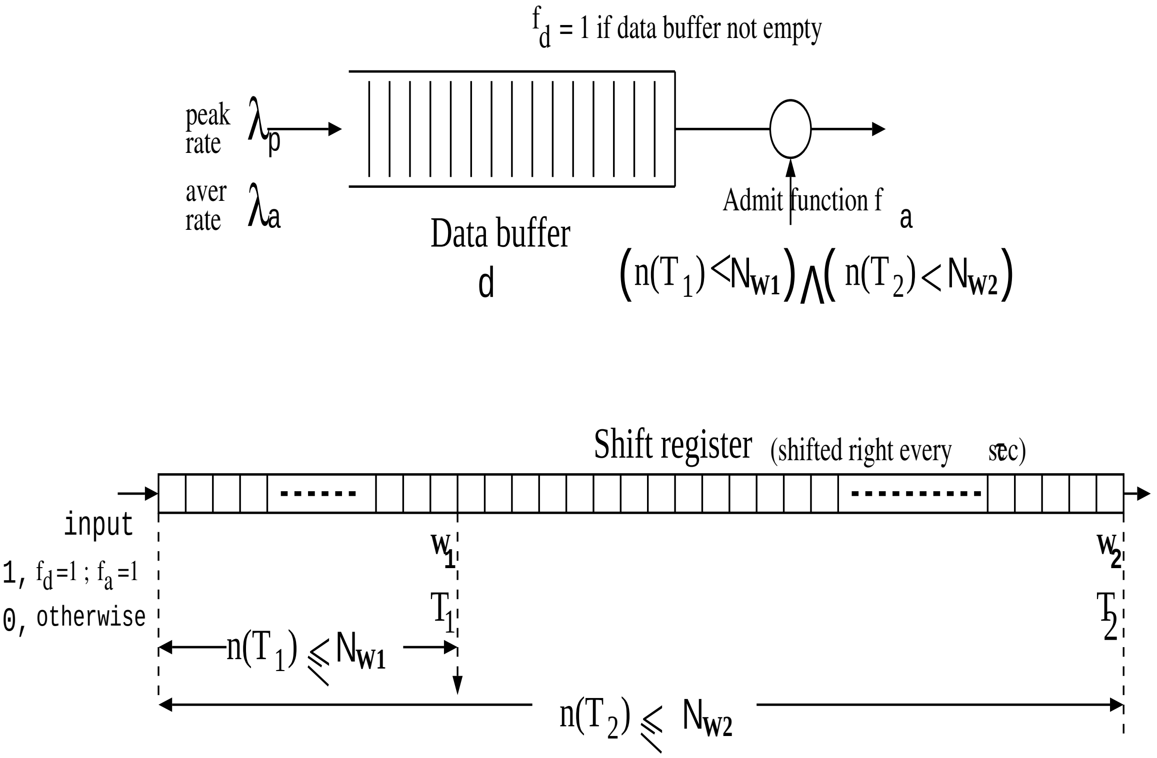

Figure 3: Shift Register Traffic Shaper (SRTS)

The advantage that is foreseen in *permitting controlled burstiness* is improvement of the statistical multiplexing gain at the switches. *This is of at most 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.

In Figure 2, the operating region of the LB was depicted. Previous section described how LB introduces access delays which can become prohibitive for real-time applications. With an aim to *reduce the access delays*, what we need is a traffic shaper which performs like the LB over longer durations, but allows short-term burstiness in a more liberal sense than is permitted by the LB. With reference to Figure 2, we attempt to operate above line OA over short durations while confining to the LB bound over a large interval (say OB). As mentioned in the previous section, OA is the upper boundary for LB operation. A typical upper boundary for the proposed shaper can be the piecewise linear line OCA. Thus by virtue of its short term operation above line OA, short term burstiness is more flexibly permitted by the proposed shaper. In the case of LB, a stream has to gather enough number of tokens by remaining inactive before it can afford to drive in a burst of data. On the contrary, a larger operating region of SRTS permits the source to have short term overdrafts as long as it confines within the operating region. A simple implementation of the scheme using 2 windows is outlined in the following section.

5.2 Description of the new scheme

The Shift Register Traffic Shaper (SRTS) makes use of the temporal profile [1] 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 $f_d = 1; f_a = 1$
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 3 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 $n(T_2) \leq n_2$.
Even 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 that arises in the above arrangement is due to the discretization of time into slots of τ. A slot 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 shown in Figure 4. 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 4.

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

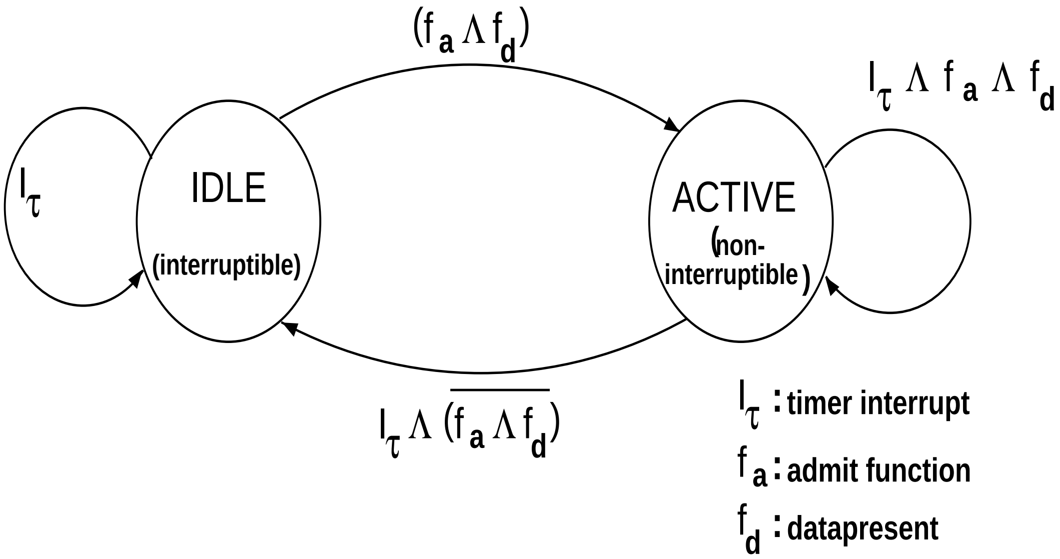

Figure 4: FSM describing the transitions between idle and active states

6 Providing adjustable burstiness

Using 3 counters and associated SRTS parameters N_{W1}, W_1, N_{W2}, W_2 and N_{W3}, W_3, it is possible to tune the burstiness at the output of the SRTS while complying with the LB bounds over a predecided time duration. Window parameters can be derived from the key observations about the characteristics of the LBP output.

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 2 LBP parameters. The third window, namely W_2, is the one for providing the variable burstiness feature. An adjustable burstiness feature can be provided in SRTS by the following choice of parameters.

1. The parameters of the smallest window T_1 are chosen as $W_1 = b'$ and $N_{W1} = b'$. This bounds the maximum burst size.

2. 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, $N_{W3} = \lambda_{eff} \cdot \tau \cdot W_3$.

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, $N_{W2} = b' + \lambda_t \cdot (W_2 - W_1) \cdot \tau$. The region of operation to permit higher burstiness is shown by the shaded arrow in Figure 5. The burstiness can be varied by adjusting N_{W2} or W_2. For instance, increasing N_{W2} or reducing W_2 increases the output burstiness.

Example For a bursty model with mean ON period of 200msec, minimum intercell time τ of 10 msec and burstiness 5,
$\lambda_p = 100$ and $\lambda_a = 20$.

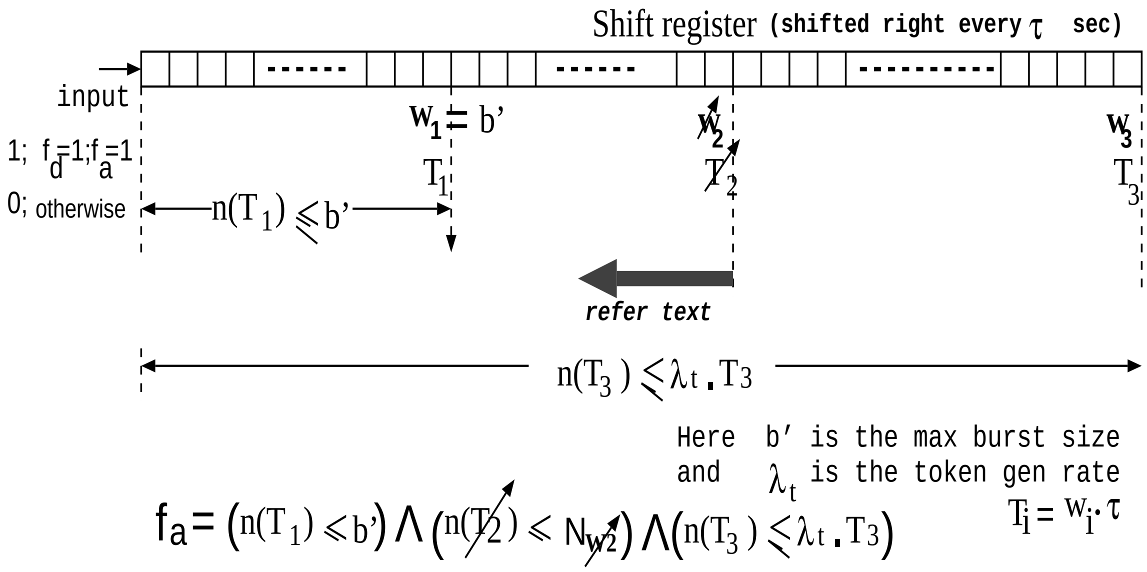

Figure 5: SRTS with variable burstiness

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 $W_3 = 450$, $N_{W3} = \lambda_{eff} \cdot \tau \cdot W_3 = 450{*}40/100 = 180$.

The exact choice of W_2 and W_3 is currently arbitrary and can be tailored to suit the specific 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.

7 Performance Study & Results

The performance characteristics of SRTS is studied in this paper in two parts. In the first part, we investigate the controlling effect of shaper parameters on the input traffic characteristics. Delay, loss and burstiness behavior at the output is studied for different window parameters and input burstiness. The adjustable burstiness feature is demonstrated in this study. In the second part, we dimension the proposed SRTS shaper and a LB shaper equivalently and compare the mean and peak rate policing behavior with delay and loss as the performance parameters.

7.1 SRTS Characteristics and Features

7.1.1 Simulation Experiments

The experiments performed to study the controlling effect of shaper parameters on input characteristics is described in this section. 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, $W_1 = N_{W1} = 30$; $W_3 = 450, N_{W3} = 180$; $N_{W2} = 48$; Size of control window W_2 is a variable parameter. Each simulation run is performed with 10^7 packets. These values are chosen based on the discussion in the previous section.

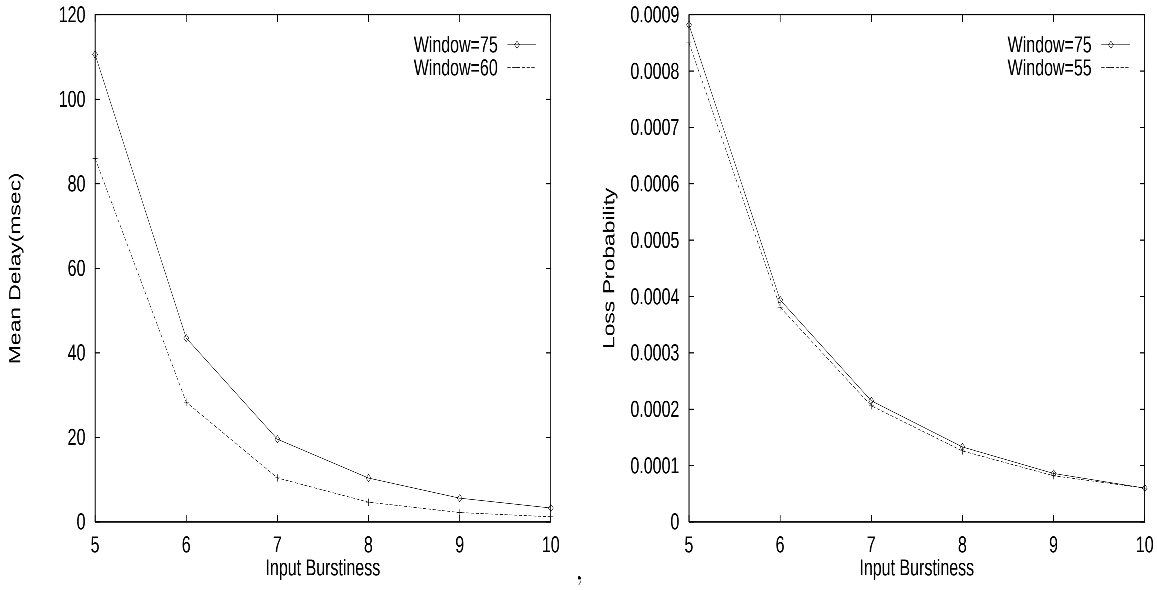

Figure 6: (a)Mean Delay and (b)Loss Characteristics Vs input burstiness

Experiment 1 In this experiment, we study the delay characteristics of the traffic shaper 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 6a 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 In this experiment, we study the loss characteristics incurred by SRTS shaping as a function of the input burstiness for different window parameters. 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.

Simulation is conducted for sufficient number of packets to yield loss probability values of up to 10^{-6} (See Figure 6b).

Experiment 3 In this experiment, we study the output burstiness as a function of window parameters, for the same source burstiness. 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 [9, 16] for characterizing the burstiness. We will use the term "burst factor" for this ratio to differentiate this definition of burstiness from the definition given in Section2.2. Figure 7a presents the result for 2 source ON-OFF characteristics. Keeping the ON time at 200msec, measurements are taken for two OFF period values, namely 800msec and 1800msec respectively.

Figure 7b 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.

7.1.2 Observations & Inference

Main observations in the simulation results and inferences drawn, thereof, are as follows.

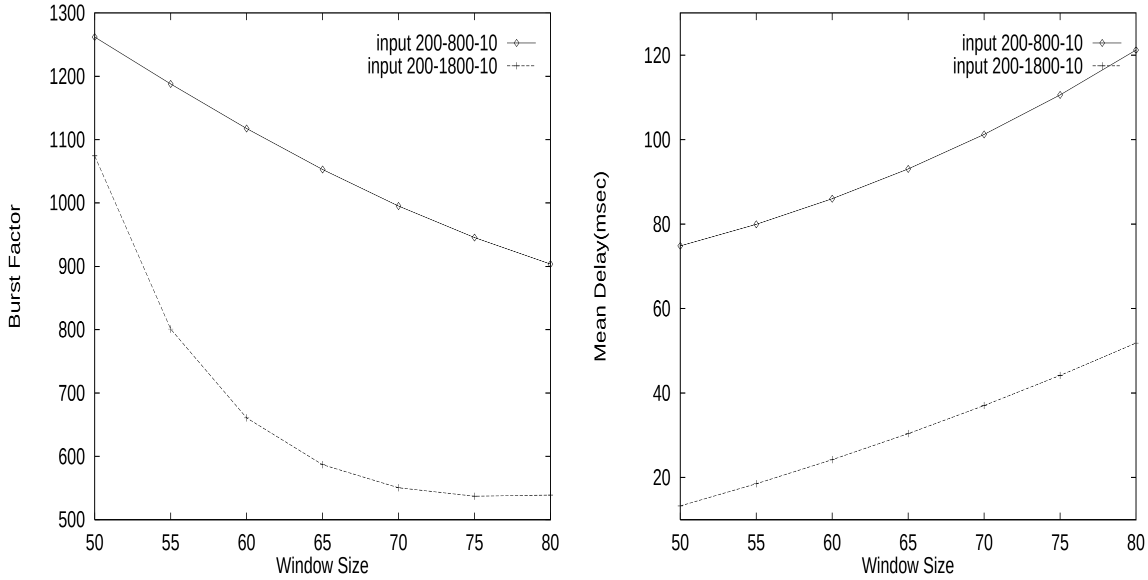

Figure 7: (a)Output Burst Factor and (b)Mean Delay Vs Window size

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 6a, 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 6b. 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 7a 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. The burstiness of the output can be tuned to provide higher bandwidth utilization at the switches.

4. The results of Figure 7b 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 practical sources.

7.2 Comparison of SRTS and LB Policing

7.2.1 Establishing Equivalence

For comparing the performance of SRTS with the LBP scheme, the parameters of the two schemes have to be chosen to establish a functional equivalence. In this paper, we use a SRTS with two windows. Our aim in this experiment is to obtain the transfer characteristics OCA depicted in Figure 2 and study its effects. The shaping parameters are the window sizes W_1, W_2 and the maximum number of packets permitted in each window N_{W1}, N_{W2}. The window parameters can be derived from the key observations made earlier regarding the LBP scheme.

The maximum burst size b' for the LBP is $b' = b/(1 - \lambda_t/\lambda_p)$. If we observe the number of packets within a window of size W(say), the maximum number of packets allowed N_W within

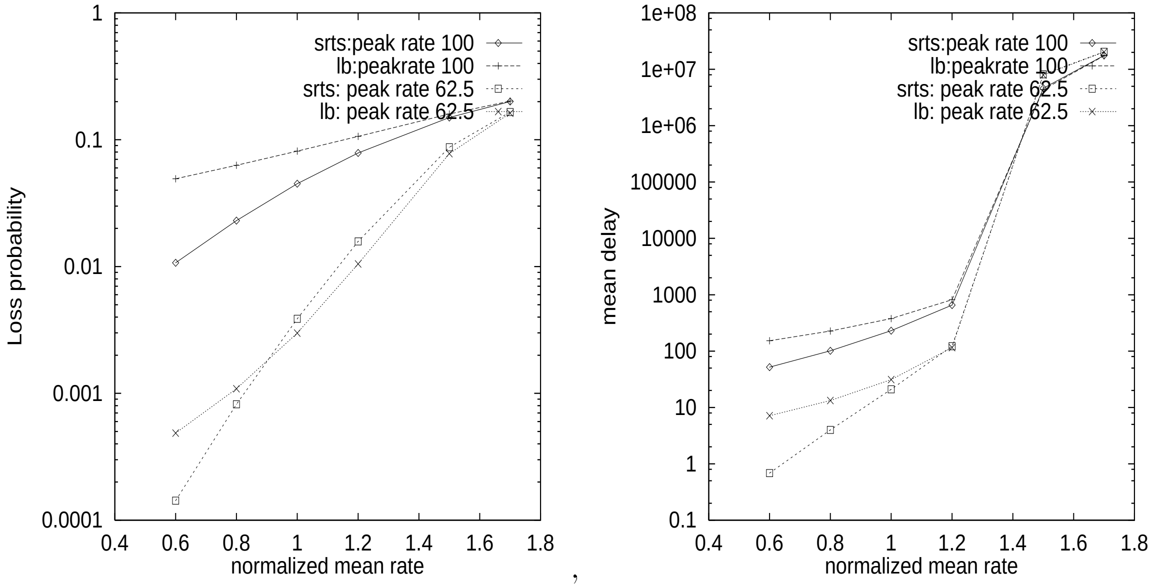

Figure 8: (a)Loss and (b)Delay characteristics vs normalized mean rate

W is:

$$\begin{aligned} &\text{for } W \leq b', \quad N_W = W; && (a)\\ &\text{for } W > b', \quad N_W = b' + \lambda_t \cdot (W - b') \cdot \tau; && (b)\\ &\text{for } W \gg b', \quad N_W \cong W \cdot \tau \cdot \lambda_t; && (c) \end{aligned}$$

The values assumed for the LBP in the current simulation are $b = 18$; $\lambda_p = 100$and $\lambda_t = 40$. Then max burst size $b' = b/(1 - \lambda_t/\lambda_p) = 30$. The first window W_1 is chosen as 50 to satisfy (b). Thus to admit more short term burstiness than what is permitted by the LBP, $N_{W1} shouldbe \geq 30 + 20 \cdot 40/100 = 38$. For the current simulation, we have chosen this value of 38 for N_{W1}. For the LBP, the distribution of these packets within W_1 should be subject to operation within the shaded region in Figure 2. Whereas, for the SRTS, they can be more flexibly distributed since the SRTS operating regime is bigger than that of LBP. Window-2 parameters can enforce the average policing characteristics exhibited by the LBP over large time durations. Hence the window size, in this case, follows (c). Consequently, the number of packets policed over a time duration $T_2 (= W_2 \cdot \tau)$ for the LBP and the SRTS are identical. For the current study, we have chosen $W_2 = 10 \cdot W_1 = 500$ and $N_{W2} = 500 \cdot \tau \cdot \lambda_t = 200$.

The exact choice of W_1 and W_2 is currently arbitrary and can be tailored to suit the application stream. The only criteria is that over W_1 , we assume the "equivalent" LBP to generate a LBAP stream whereas over the larger window W_2 , an averaging property is expected.

7.2.2 Simulation Experiments

In this section, we compare and study the effectiveness of SRTS and LBP as mean and peak rate policers. Two simulation experiments are performed. The source model is the bursty ON-OFF model explained in Section 2.2. Since we intend to vary the burstiness of the source, the mean ON time is kept at 200msec. The OFF times and τ_p are appropriately adjusted to obtain the required mean rate.

SRTS is a mean as well as a peak rate policer. In the two experiments, we assume an overdimensioning factor $C = 1.5$ relating the policed rate and the mean rate of the source (as in [11]). The peak enforced rate is 100 and hence the minimum delay between consecutive packets at the output of the shaper τ is 10 ms. Each simulation run is performed with 10^7 packets.

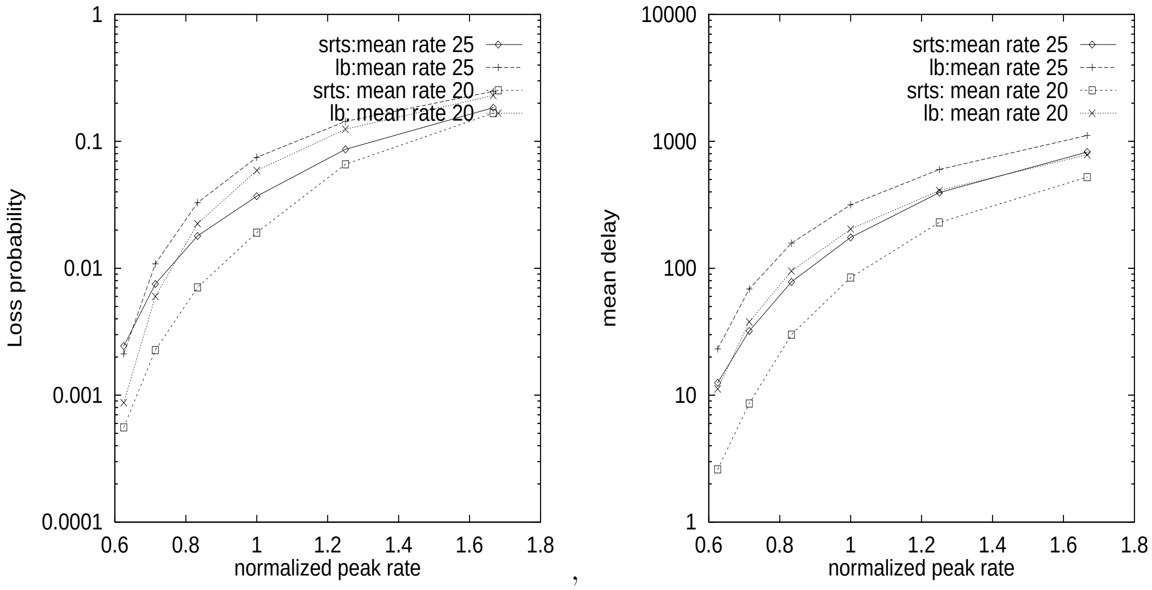

Figure 9: (a)Loss and (b)Delay characteristics vs normalized peak rate

Experiment 1 In this experiment, we study the loss and delay characteristics for different source mean rates. The mean rate variation is achieved by varying the OFF time keeping mean ON time = 200 and mean policed rate, $\lambda_t = 40$. With the overdimensioning factor of 1.5, the negotiated mean rate = 26.67. The OFF time is varied such that
$\lambda_p/(1 + \frac{(T_{OFF})}{(200)}) \leq 26.67$. Thus $T_{OFF} > 550$ for a well behaved source. X axis shows the normalized mean rate. For the first part which estimates the violation probability, a finite data buffer of size 20 is assumed. In a practical case, the size can be based on the maximum access delay that can be tolerated by a particular application. For the second part of the experiment which studies the access delay, size of the data buffer is kept very large so as to keep losses close to zero. The experiment is performed for two values of the peak rate, $100(\tau_p = 10)$ and $62.5(\tau_p = 16)$. The results are shown in Figure 8.

The number of simulation runs are such that the results are accurate to within 5% with 95% confidence level.

Experiment 2 In this experiment, we study the loss and delay characteristics for different source peak rates. Thus we compare the peak rate enforcement provided by the SRTS and the LBP. For each run, the peak rate and the OFF duration are adjusted to keep the mean rate constant. X axis plots the normalized peak rates. The experiment is repeated for two values of the mean rate, 25 and 20. Both these values are within the negotiated rate of 26.67. Other parameters are as in the previous experiment. The results are shown in Figure 9.

7.2.3 Observations & Inference

Main observations in the simulation results and inferences drawn, thereof, are as follows.

With reference to Figure 8a, for an input stream with peak rate 100 (corresponding to the peak rate limit built in the shaper, SRTS has much lesser loss probability for mean rates up to the policed rate (1.5 * source mean). Beyond this, both the curves converge quickly. At the lower peak rate of 62.5, however, there is a crossover between the SRTS and LBP loss curves. This we attribute to the fact that the source traffic is smooth in this region and the advantage of SRTS in favoring short term burstiness is not made use of. In both the cases, the steeper gradient of the SRTS curve is an indicator of its effectiveness as a mean rate policer. The flexible admission of short term burstiness results in a lower access delay for the SRTS. This fact is evident from

Figure 8b. For well behaved sources with mean rate below the negotiated value, lower the mean rate, better the performance of the SRTS. This is true from the point of view of loss probability as well as access delay. At 0.6 times the mean rate, the access delay introduced by the SRTS is one order less than that introduced by the equivalent LBP.

Figure 9 depicts the response of the shapers to peak rate violation. For the loss curves, violation is more gradual than in the mean rate case. For our simulation which assumed a data buffer of size 20, SRTS yields lower values of violation probability than the LBP for traffic conforming to the negotiated rate. This is due to the more liberal admission policy for burstiness existing over short durations. The access delay curves for the two shapers are almost parallel to each other. As in the previous case, SRTS shaped streams have a consistently smaller access delay compared to the LBP case. However, compared to the delay characteristics for mean rate violation behavior, peak rate violation curves for SRTS as well as LBP do not exhibit steep gradients.

The advantages of the SRTS policy in terms of lower violation probability and access delay for traffic *within the negotiated rates* is due to the the larger operating regime shown in Figure 2. The above advantage of the SRTS however comes at a cost. The SRTS output is burstier than its LB counterpart. This would necessitate a more careful buffering and scheduling design at the switches to prevent congestion at the intermediate nodes. Since the network link transmission rate is generally much higher than the maximum source transmission rate, we expect that the fluctuations at the SRTS output will be effectively smoothed by the statistical multiplexing effect at the switches. Also, since the maximum burst size is limited and the long term behavior is bounded, the buffers and the schedulers can be dimensioned appropriately at the switches to provide the required degree of loss and delay guarantees.

From the point of view of minimizing congestion within the network, the policy adopted by the LB is quite effective. LB reduces the delays within the network by transferring them on to the input side. However, the stringent enforcement increases the access delay and hence raises questions regarding the suitability of LB for real time traffic. We show through this study that the access delays can be reduced by adopting a more liberal attitude over shorter durations while maintaining the LB bounds over larger durations. For the same bandwidth allocation at the switches, such a policy is shown to perform better for real time source traffic.

8 Summary and Conclusion

In this paper, we proposed a flexible traffic shaper and compared its performance with a LBP. The motivation for the new scheme is derived from the output characteristics exhibited by the LBP. Two main goals were set. One is to provide an adjustable burstiness feature so that higher bandwidth utilization along with reasonable guarantees can be obtained. The second was to reduce the access delays for real-time traffic by being more liberal in permitting short term burstiness. The window based shaping policy adopted in the SRTS scheme can be used to achieve both the goals.

The performance of the proposed shaper is studied in two parts. In the first we study the effect of window parameters on input characteristics and demonstrate the adjustable burstiness feature. In the second part, we compare the loss and delay performance of a 2-window SRTS and a LBP. By adopting a more liberal, yet bounded attitude over short durations, SRTS reduces the access delays for time critical traffic.

For providing the desired utilization and guarantees, a traffic shaper must work in unison with the buffer management and scheduling schemes at the switches. A composite study involving the shaper and the scheduler is necessary to see the effect of SRTS shaping on end to end performance. Such a study will constitute our future research.

9 References

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

[2] 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.

[3] Krishna Bala, Israel Cidon, and K.Sohraby. Congestion Control for High Speed Packet Switched Networks. In *IEEE INFOCOM*, pages 520–526, 1990.

[4] 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.

[5] 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.

[6] A. L. Elwalid and D. Mitra. Effective Bandwidth of General Markovian Traffic Sources and Admission Control of High Speed Networks. In *Proceedings of IEEE INFOCOM'93*, pages 256–265, San Francisco, California, March 1993.

[7] 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.

[8] 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.

[9] 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.

[10] Masayuki Murata, Yoshihiro Ohba, and Hideo Miyahara. Analysis of Flow Enforcement Algorithm for Bursty Traffic in ATM Networks. In *Proceedings of IEEE INFOCOM'92*, pages 2453–2462, Firenze, Italy, May 1992.

[11] 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.

[12] San qi Li and Song Chong. Fundamental Limits of Input Rate Control in High Speed Network. In *Proceedings of IEEE INFOCOM'93*, pages 662–671, San Francisco, California, March 1993.

[13] 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.

[14] 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.

[15] 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.

[16] 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.

[17] 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.

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

[19] David Yates, James Kurose, Don Towsley, 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.

Author Information

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. Dr. Raghavan is a life member of the Computer Society of India, a member of the Institution of Engineers and a fellow of the Institute of Electronics and Telecommunication Engineers. 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. From 1968 to 1970, Dr. Agrawala was a Senior Engineer in the Applied Research Department, Honeywell, Inc. Waltham, MA. In 1970, he joined the Honeywell Information Systems as Principal Engineer in the Optical Character Recognition Department. Since 1971, he has been on the Faculty of the Department of Computer Science, University of Maryland, College Park. He has been actively involved in research on several aspects of computer systems. He introduceed the idea of clustering for characterizing the workload of computer systems. He developed techniques for transient analysis of queues. 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 operating in distributed environments and supporting fault tolerant operations. He has also developed deterministic analysis techniques for computer networks and applied them to flow and congestion control problems. In the past, he has published on microprogramming, computer architectures, local area networks and pattern recognition. He is the author of many papers and books.

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.