The Partially Stopped Leaky Bucket: An Efficient Traffic Regulator with Constant Time Implementation

Florin Ciucu, Brian L. Mark*, and Robert Simon
Lab. of Cooperative Work Environment
Dept. of Computer Science
*** Dept. of Electrical and Computer Engineering**
School of Information Technology and Engineering
George Mason University
Fairfax, VA 22030
fciucu@gmu.edu, bmark@gmu.edu, simon@cs.gmu.edu

Abstract— ***Traffic regulation*** **is a key component in providing quality-of-service guarantees in high performance networks. In this paper, we propose a novel traffic regulator called "Partially Stopped Leaky Bucket" (PSLB) that can accurately capture a general class of traffic patterns, yet has a constant time implementation. PSLB is based on the well-known leaky bucket regulator, but is able to capture the important property of decreasing traffic rate over sufficiently large intervals. We derive a number of key properties of the PSLB and compare its performance to the general D-BIND model, which has a polynomial time implementation. We show that the PSLB model retains the important properties of the D-BIND model, yet has a much more efficient implementation.**

Keywords— **Network Calculus, envelope functions, traffic regulators, leaky bucket, quality of service.**

I. Introduction

Providing deterministic quality of service (QoS), especially to critical real-time network applications, is an important issue that continues to receive a lot of attention. The difficulty in providing deterministic QoS stems from application-level heterogeneity and the resulting complexity in network support mechanisms. Two major architectural components used to provide deterministic QoS are the traffic constraint functions, also called envelopes, and the scheduling algorithms. Constraint functions specify an upper bound on input arrival traffic during each interval of time. Usually, the specification is time-invariant, in the sense that for each interval of length t, an upper-bound of $f(t)$ is imposed [17]. Time-varying constraint functions are also considered in the literature [16]. The scheduling algorithms specify the order and timeliness of how the input arrivals will depart from a network element. Many such algorithms have been proposed and extensively studied. An excellent review of these may be found in [10].

In this paper we propose a new traffic regulator, the Partially Stopped Leaky Bucket (PSLB), that captures the characteristics of a general class of traffic patterns, yet has an efficient constant time implementation. The intuitive idea of our model is a leaky bucket whose token generation rate goes to 0 during specific intervals of time. The name of the model expresses this idea.

We will discuss the PSLB model in the context of both the properties of traffic constraint functions and the implementation of the corresponding traffic regulators. With respect to the first issue, in terms of description and implementation, there are simple models like (σ, ρ) and $(X_{\min}, X_{ave}, I, S_{\max})$ which capture traffic information such as temporary burst and long-term rate. More complex models like $(\overrightarrow{\sigma}, \overrightarrow{\rho})$ or D-BIND (deterministic bounding interval-length dependent) describe the traffic such that the traffic rate may decrease over long intervals [9]. The D-BIND model is more general than the $(\overrightarrow{\sigma}, \overrightarrow{\rho})$ model because it can capture properties such as the rate being greater over a longer time interval than a shorter one. Despite the generality of D-BIND, its shortcoming comes from the polynomial implementation of the corresponding traffic regulator. Our primary goal herein is to give a description of a new envelope model, as general as possible, yet having an O(1) implementation of the corresponding traffic regulator.

The rest of the paper is structured as follows. First, we review some concepts of deterministic traffic models. Then, we present the PSLB model and derive its key properties. Finally, we will give a constant time implementation of PSLB and compare it to the implementation of D-BIND.

II. Background

Let us consider a network element E, usually a switch, an equidistant division of time with increment 1 and a discrete-time input arrival process $a(t), t = 0, 1, ...,$ with the convention that $a(0) = 0$. The incoming packets will have a length of 1 and will be served at a rate of 1. The non-negative, non-decreasing function $A : \mathbf{N} \rightarrow \mathbf{N}$ with

$A(t) = \sum_{s=0}^{t} a(s)$ will denote the cumulative arrival process. For simplicity, we will simply denote by A the arrival process. The output process will be denoted by B and the relation between A and B will be considered causal; also, we will consider a perfect cut-through transmission through each network element.

We give now some useful definitions from the functional analysis:

Definition 1: A function $f : \mathbf{R} \to \mathbf{R}$ is concave if the following inequality holds:

$$f(\alpha s + (1-\alpha)t) \geq \alpha f(s) + (1-\alpha) f(t) \ (\forall)\ s \leq t \text{ and } (\forall)\ 0 \leq \alpha \leq 1.$$

Definition 2: A function $f : \mathbf{R} \to \mathbf{R}$ is subadditive if the following holds:

$$f(s+t) \leq f(s) + f(t) \ (\forall)\ s, t.$$

Definition 3: A function $f : \mathbf{R} \to \mathbf{R}$ is superadditive if the following holds:

$$f(s+t) \geq f(s) + f(t) \ (\forall)\ s, t.$$

For the purpose of our paper, the following definition is crucial:

Definition 4: A function $f : \mathbf{N} \to \mathbf{N}$ with the property that

$$A(t) - A(s) \leq f(t-s) \ (\forall)\ 0 \leq s \leq t, \tag{1}$$

will denote a traffic constraint function or an *envelope* of A.

The relation between A and f will be denoted by $A \prec f$. Without loss of generality, f is assumed to be subadditive in the sense that if $A \prec f$, then $A \prec f^*$, where f^* represents the subadditive closure of f. For more details about the subadditivity of envelope functions, see [17].

The tightest envelope of A is given by the so-called *empirical envelope* [11] or *minimal envelope* process [6] and is defined as

$$\varepsilon_A(t) = \sup_{s \geq 0} (A(s+t) - A(s)). \tag{2}$$

In practice, all the traffic constraint models try to approximate ε_A. Of future help will the following:

Lemma 1: If f is an envelope for a given arrival system and there is an input pattern A such that $f = A$, then f is the minimal envelope of that system.

Proof: Let us suppose by contradiction that there exists a minimal envelope g with $g < f$. We have that: $f(t) = A(t) = A(t) - A(0) \leq g(t-0) = g(t)$ which is false. ■

In the following, we briefly review several well-known traffic envelope functions that have been discussed in the literature. The affine model (σ, ρ) is the simplest among the traffic constraint functions [3]. It is defined as $f(t) = \sigma + \rho t$, where $\sigma \geq 0$ represents the maximal allowable burst of A, while $\rho > 0$ represents the long term average rate of A. An immediate observation is that the envelope function f is subadditive. The $(X_{\min}, X_{ave}, I, S_{\max})$ model, also called the $X_{\min}$ model, is defined in [1] and specifies that the minimum inter-arrival time of two consecutive packets is delimited by $X_{\min}$, the maximum packet length is $S_{\max}$ and the average arrival traffic over a period of I is upper-bounded by $\frac{I}{X_{ave}}$. The corresponding minimal envelope may be expressed by:

$$f(t) = \left(\min \left(\left\lceil \frac{t \bmod I}{X_{\min}} \right\rceil, \left\lceil \frac{I}{X_{ave}} \right\rceil \right) + \left\lfloor \frac{t}{I} \right\rfloor \left\lceil \frac{I}{X_{ave}} \right\rceil \right) S_{\max} \tag{3}$$

In Fig. 1 we present an example of the constraint function (3). Based on this, we will see later the analogy between the $X_{\min}$ and PSLB models.

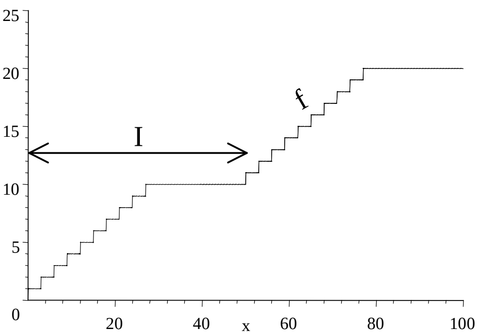

Fig. 1. X-min model

We point out that the original expression for f given in Eq. (1) in [7]:

$$f(t) = \left(\min \left(\left\lceil \frac{t \bmod I}{X_{\min}} \right\rceil, \left\lceil \frac{I}{X_{ave}} \right\rceil \right) + \left\lceil \frac{t}{I} \right\rceil \left\lceil \frac{I}{X_{ave}} \right\rceil \right) S_{\max} \tag{4}$$

is not a true envelope function, since the function f decreases at integer multiples of I. In the next proposition, we state that a corrected minimal envelope for the $X_{\min}$ is f given by (3).

Proposition 2: For the $X_{\min}$ model, the constraint function f given in (3) represents the minimal envelope. The proof for this is given in Appendix I.

The $(\overrightarrow{\sigma}, \overrightarrow{\rho})_n$ model [11] has an envelope function

$$f(t) = \min_{i=\overline{1,n}} (\sigma_i + \rho_i t), \tag{5}$$

where $0 \leq \sigma_1 < \ldots < \sigma_n$ and $0 < \rho_n < \ldots < \rho_1$. It is worth noticing that f is a piece-wise linear concave function. The model has its roots in the work of [3] where the author studied some properties of concave traffic constraint functions.

The last model we discuss here is the D-BIND model, described in [9], [11] and [13]. This model generalizes the

$(\overrightarrow{\sigma}, \overrightarrow{\rho})_n$ model by allowing non-concave traffic constraint functions. Formally, the condition $0 \le \rho_n \le ... \le \rho_1$ is replaced by the weaker condition $\rho_i \ge 0$ $(\forall) i = \overline{1,n}$. The motivation for this is the behavior of MPEG-compressed video traces, where average rate ratios may be increasing as functions of time. Also, using non-concave envelope functions is beneficial when heterogeneous VBR traffic streams are multiplexed [13]. The original description of D-BIND was through n pairs (R_i, I_i), where $0 < I_1 < ... < I_n$, that constrain the arrival traffic as having an average rate of $R_i \ge 0$ over the corresponding interval of time I_i; a similar model has been considered also in [14]. The envelope function of D-BIND is simply obtained by interpolation as:

$$f(t) = \frac{R_i I_i - R_{i-1} I_{i-1}}{I_i - I_{i-1}}(t - I_i) + R_i I_i, \quad I_{i-1} \le t \le I_i,\ (\forall)\ i = \overline{1,n} \tag{6}$$

and the convention that $I_0 = 0$.

Besides the concept of traffic envelope, the concept of *service curve* is necessary to derive performance results in the deterministic network calculus. We will adopt here the definition from [6], where the author puts the service curve's definition in the context of using min-plus algebra in the filtering theory [17], [15], as a fundamental framework for the deterministic network calculus. Specifically, S denotes a service curve if $B(t) \ge \min_{0 \le s \le t} (A(s) + S(t-s))$.

Deterministic network calculus has been developed to provide QoS to critical application. Two important QoS parameters are represented by the *maximum delay* encountered by the packets flowing through the network element and the necessary *buffer length* in order not to be forced to drop packets arbitrarily. Considering a FIFO scheduling policy and that $A \prec f$, the virtual delay $d(t) = \inf\{d \ge 0 \mid A(t) \le B(t+d)\}$ encountered by a packet arriving at a time t, is bounded by

$$d(t) \le \inf\left\{d \ge 0 \mid f(s) \le S(s+d)\ (\forall)\ s = \overline{1,t}\right\}. \tag{7}$$

Also, the virtual backlog $b(t)$ is bounded by $b(t) \le \sup_{0 \le s \le t} (f(s) - S(s))$ [17].

III. Partially Stopped Leaky Bucket Model

The main shortcoming of the $X_{\min}$ model is represented by the fact that the arrival rate of the traffic is almost constant as a function of time. The relaxation of this constraint lead to more sophisticated constraint models like $(\overrightarrow{\sigma}, \overrightarrow{\rho})_n$, in which the arrival rate decreases with time, or the D-BIND model, in which, as we have already noted, the rate may increase/decrease as a function of time.

The PSLB is based on the leaky bucket or (σ, ρ) model. Also, a subclass of PSLB may be viewed as a variant of the $X_{\min}$ model for which $S_{\max} = 1$ and more than one packet is allowed to arrive during the same time slot.

Informally, the intuition behind a PSLB envelope is a (σ, ρ) envelope function modified as follows. The PSLB curve starts at time 0 with a value of σ and then increases with a slope of ρ. At one point in time, x_1, it will remain constant for a given time t_1. Then, again, it will increase with a slope of ρ for a time of x_1, followed by a temporarily stopping of time t_2 and so on. Note that the intervals in which the PSLB envelope increases are equal, while the intervals t_i in which the PSLB envelope remains constant must be increasing. We now give the formal definition of the PSLB envelope.

Let us consider an affine model (σ, ρ), together with a non-decreasing function $g : \mathbf{N} \to \mathbf{N}$ such that the corresponding envelope will be

$$f(t) = \sigma + \rho(t - g(t)) \tag{8}$$

An example is presented in Fig. 2 for

$$g(t) = \begin{cases} 40j, & 50j \le t \le 50j + 10 \\ t - 10(j+1), & 50j + 10 \le t \le 50(j+1), \quad j = 0, 1, 2, ... \end{cases}$$

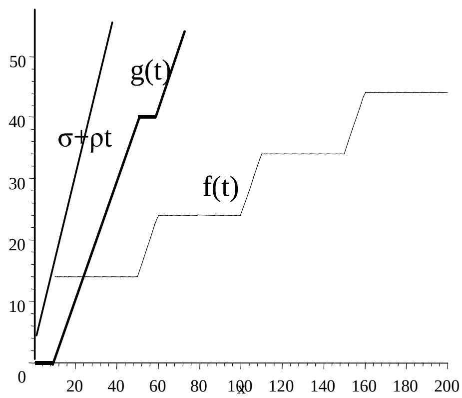

Fig. 2. PSLB envelope function derivation

When $X_{\min} = 1$, the $(X_{\min}, X_{ave}, I, S_{\max})$ model can be represented by the envelope (8) with $\sigma = 0$, $\rho = S_{\max}$ and g given as follows:

$$g(t) = t - \left\lfloor \frac{t}{p_2} \right\rfloor p_1 - \min(p_1, t \bmod p_2), \tag{9}$$

where $p_2 = I$ and $p_1 = \left\lceil \frac{I}{X_{ave}} \right\rceil$. This shows the PSLB model is more general than the $X_{\min}$ model for which $X_{\min} = 1$. Generalization of the PSLB model may also be given in order to cover the general $X_{\min}$ model, where $X_{\min}$ may be greater than 1.

For the PSLB model, we define the function g as follows: let $\{x_n\}_n$ be a sequence of positive numbers such that $x_0 = 0, x_i < x_{i+1}$ and $x_{i+1} - x_i \ge x_i - x_{i-1} \ge x_1$

$(\forall)\, i \geq 1$. The intuition behind this sequence stems from the informal description of PSLB given above.

Then g is defined, recursively, by $g(0) = 0$ and:

$$g(t) = \begin{cases} g(t-1)+1, t \in (x_i, x_{i+1} - x_1] \;\; (\forall) \;\; i \geq 1 \\ g(t-1), otherwise \end{cases} \tag{10}$$

Note that $\{x_n\}_n$ may be finite. In such a case, the last term $x_N = \infty$. With this definition, we can see that the expression (9) may be obtained for $x_i = (i-1)p_2 + p_1$. Also, we have the following important fact:

Proposition 3: The function g defined by (10) is superadditive.

Proof: We will prove by induction on t, that $g(t) \geq g(t-s) + g(s)$ $(\forall)$ $0 < s < t$. Since $g(0) = 0$, the first step of the induction is immediate. We suppose now that $g(t') \geq g(t'-s) + g(s)$ $(\forall)$ $0 < s < t' < t$.

If $t \in (x_i, x_{i+1} - x_1]$, where $i \geq 1$, we have that $g(t) = g(t-1) + 1 \geq g(s) + g(t-1-s) + 1 \geq g(s) + g(t-s)$.

Now let $t = x_i$, where $i \geq 1$, and $0 < s < t$. Suppose that $s \in (x_k, x_{k+1} - x_1]$. We have that: $g(x_i) = g(s) + x_k - s - x_1 + x_{k+1} - x_k - x_1 + ... + x_{i+1} - x_i - x_1 = g(s) + x_i - s - (i-k+1)x_1$. It follows that $g(x_i) - g(s) \geq x_i - s - (i-k+1)$. On the other hand, $g(x_i - s) = x_i - s - jx_1$, where j may be determined. However, since $x_{i+1} - x_i - x_1 \geq x_i - x_{i-1} - x_1$ $(\forall)\, i \geq 1$, we have that $j \geq i-k+1$. It follows immediately that $g(x_i - s) \leq g(x_i) - g(s)$. The case $s \in (x_k - x_1, x_k]$ may be treated similarly.

Finally, if $t \in (x_i - x_1, x_i]$ and $0 < s < t$, we have that: $g(t) = g(x_i) = g(x_i - s + s) \geq g(x_i - s) + g(s) \geq g(t-s) + g(s)$. ■

Proposition 4: If g is superadditive, then f in (8) is subadditive.

Proof: Let $s, t \geq 0$. It follows immediately that $f(s+t) = \sigma + \rho(t+s-g(t+s)) \leq \sigma + \rho(t+s-g(t)-g(s)) \leq \sigma + \rho(t - g(t)) + \sigma + \rho(s - g(s)) = f(s) + f(t)$. ■

Observe that other conditions on g may be specified for which f is subadditive. Fig. 3 shows an example of a PSLB envelope function.

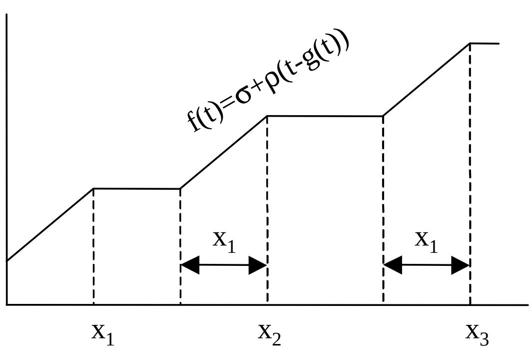

Fig. 3. PSLB envelope function

From Fig. 3, we can see that

$$\frac{\sum_{t=0}^{x_i - x_1} f(t)}{x_i - x_1} \geq \frac{\sum_{t=0}^{x_{i+1} - x_1} f(t)}{x_{i+1} - x_1} \;\; (\forall) \;\; i > 0,$$

which demonstrates that the PSLB model can capture the important property of decreasing rate for sufficiently large intervals. Also, it may be noticed that the PSLB model allows the arrival rate to increase temporarily as function of time. Recall that this is a key feature of the D-BIND model.

To explain how the function g may be obtained, note first that the average rate of a PSLB envelope is $\lim_{t\to\infty} \frac{\sigma + \rho(t - g(t))}{t} = \rho(1 - \lim_{t\to\infty} \frac{g(t)}{t})$. For the example from Fig. 2, we have that $\lim_{t\to\infty} \frac{g(t)}{t} = \frac{4}{5}$. This limit represents a design parameter. In order to attain a long run average rate of α, we need to have $\lim_{t\to\infty} \frac{g(t)}{t} = 1 - \frac{\alpha}{\rho}$.

Referring to the Proposition 2, we can say that (8), together with (10), implements a minimal envelope for a model that limits the number of packets over all intervals of length $x_i - x_1$ to $\sigma + i\rho x_1$ $(\forall)$ $i > 1$.

IV. Envelope implementation

We now address the issue of real-time implementation of the traffic constraint functions. After discussing on implementations of general traffic regulators, we will give an efficient algorithm for implementing a PSLB regulator. Here, we assume that the input A is given to us and is not otherwise regulated.

The leaky bucket regulator (σ, ρ) has an intuitive implementation as follows. There is a bucket T, of fixed length σ, where tokens arrive from an infinite pool P at a fixed rate of ρ. Initially, T is full. As packets arrive, they leave the regulator as soon as possible, with the restriction that the maximum number of packets to depart at a given time is bounded by the bucket's level. Each packet that leaves the regulator consumes one token. It is assumed that there is sufficient buffer space to hold all outstanding packets. Also, tokens that arrive when T is full will be dropped. In [17], the following result is presented:

Theorem 5: The leaky bucket regulator (σ, ρ) generates an output B that satisfies:

$$B(t) = \min\left(A(t), \min_{0 \leq s \leq t-1} (A(s) + f(t-s))\right)$$

where $f(t) = \sigma + \rho t$ $(\forall) t > 0$ and $f(0) = 0$.

The above construction implements a maximal f-regulator (under min-plus algebra terminology), meaning that for any other regulator that will generate an output $\tilde{B}$ with $\tilde{B} \prec f$, then $\tilde{B} \leq B$ holds. However, we point out that the statement of Theorem 5 has a small, but noteworthy error. To show this, let us consider the case where $A(t) = f(t)$ $(\forall) t \geq 0$. The output B generated by

(σ,ρ) will be $B(t)=A(t)-\rho$ $(\forall)t>0$ and $B(0)=0$. On the other hand, if we modify the (σ,ρ) model by considering the token bucket's level to be $f(1)$, instead of σ, an output $\tilde{B}=A$ will be obtained. Moreover, one has that $\tilde{B}>B$, meaning that the original (σ,ρ) mechanism does not implement a maximal f-regulator.

Remark 6: (on the proof of Theorem 5 given in [17])

The error appears in "Case 2". More exactly, instead of "The maximum number of departures ... ,i.e., $q_2(t)+\rho$.", it should be "... $\max(q_2(t)+\rho,\sigma)$.

Intuitively, the discrepancy comes from the fact that there are no arrivals (hence, no output) at time 0, while the considered model is in discrete-time. We will call the modified (σ,ρ) mechanism as $(\sigma,\rho)^{MR}$ (MR stands for "maximal regulator"). Following the original proof of the previous theorem, we conclude that the theorem will be valid for the $(\sigma,\rho)^{MR}$ model instead of the (σ,ρ) model. Also, it is worth mentioning that f represents the minimal envelope for the $(\sigma,\rho)^{MR}$ model (see Lemma 1).

The implementation of the $X_{\min}$ regulator may be done in a greedy fashion as in [2]: for any $\left\lceil \frac{I}{X_{ave}} \right\rceil + 1$ packets, the following condition must hold:

$$\sum_{i=1}^{\left\lceil \frac{I}{X_{ave}} \right\rceil} X_i \geq I,$$

where X_i represents the interarrival times between the i^{th} and $(i+1)^{th}$ packets.

The existing suggestion [3] to implement the $(\overrightarrow{\sigma},\overrightarrow{\rho})_n$ model ($f(t)=\min\limits_{i=\overline{1,n}}(\sigma_i+\rho_i t)$) is through n-parallel leaky buckets. The maximum number of packets to depart will be bounded by the minimum allowed to depart among all the leaky buckets. Because this operation requires $O(n)$ time at each instant of time, there is a trade-off between the accuracy of f (the larger is n, the more accurate is f relative to the empirical envelope ε_A) and the performance of the regulator in terms of time complexity.

As already noted, the subadditivity property of a constraint function plays an important role with respect to a model's performance.

Proposition 7: Let $f:\mathbf{N}\to\mathbf{N}$ be a concave function with $f(0)\geq 0$. Then f is subadditive.

Proof: Let $0<x\leq y$. As f is concave, we have $\frac{f(x)-f(0)}{x}\geq\frac{f(x+y)-f(x)}{y}$. It follows that:

$$f(x)(\frac{1}{x}+\frac{1}{y})\geq\frac{f(x+y)}{y}+\frac{f(0)}{x}. \tag{11}$$

Similarly,

$$f(y)(\frac{1}{x}+\frac{1}{y})\geq\frac{f(x+y)}{x}+\frac{f(0)}{y} \tag{12}$$

From (11) and (12) we have immediately that $f(x)+f(y)\geq f(x+y)+f(0)$. Since $f(0)\geq 0$ the subadditivity of f follows. ∎

From this fact, the subadditivity of $f(t)=\min\limits_{i=\overline{1,n}}(\sigma_i+\rho_i t)$ follows immediately by noting the concavity of f. Besides this property we have the following extension of Theorem 5 [17]:

Theorem 8: If $f(t)=\min\limits_{i=\overline{1,n}}(\sigma_i+\rho_i t)$ $(\forall)t>0$ and $f(0)=0$, then Theorem 5 holds.

Again, we stress the fact that the theorem is valid if the basic regulators work according to the $(\sigma_i,\rho_i)^{MR}$ model defined previously.

With respect to the more general D-BIND model, we note that the implementation of $(\overrightarrow{\sigma},\overrightarrow{\rho})_n$ cannot be used because it is restricted to concave envelopes. Moreover, a D-BIND envelope is generally not a subadditive function and finding its subadditive closure may be a very complex task. There are two current ideas for the implementation of D-BIND, both proposed in [9]. First, it is suggested to use the $(\overrightarrow{\sigma},\overrightarrow{\rho})_n$ implementation for the concave closure of D-BIND's envelope. The direct consequence of this approach is loss of accuracy. The second idea is to implement a cascade of leaky buckets with state dependent token generating rates, depending on the number of transmitted packets over the last interval. It seems that the last method provides a "pseudo"-$O(1)$ time complexity (meaning that the constant may be too large) for D-BIND and also for the $(\overrightarrow{\sigma},\overrightarrow{\rho})_n$ model. However, as we discuss below, this is not the case here. Consider the example of traffic input A and D-BIND envelope function f shown in Fig. 4:

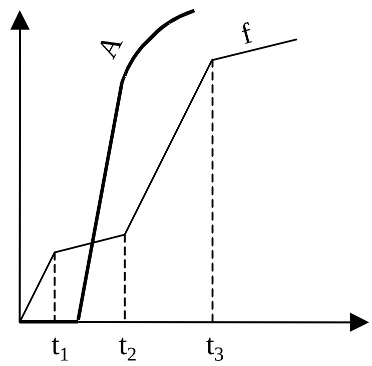

Fig. 4. D-BIND envelope

If we assume in Fig. 4 that $t_3-t_2>t_2$, $A(t_2-1)=0$ and $A(t_2)\gg 0$, it is clear that for any point t_0 during the interval $[t_2,t_3]$ we need the number of packets transmitted during the intervals $[t_0-t_1+1,t_0)$ and $[t_0-t_2+1,t_0]$. The idea behind this example is that at any time we have to guarantee the time-invariance property of the regulator. Moreover, this property must be generally satisfied for an $O(n)$ number of intervals. Therefore, without a property that will simplify the procedure from $O(n)$ intervals to a polynomially smaller number of intervals, it appears that the required time of D-BIND is actually $O(n)$, as in the case of the $(\overrightarrow{\sigma}, \overrightarrow{\rho})_n$ model.

From the implementation point of view of PSLB, one may notice that if we have modified the leaky bucket algorithm merely by temporarily stopping the token generation rate, there may be some unnecessary delays in the regulator. Although our algorithm is not able to implement a maximal regulator in O(1) time for the general case, it significantly outperforms the previous "first glance" algorithm. We now give an $O(1)$ time algorithm for implementing a PSLB regulator.

Variables:
P - pool of tokens, $|P| = \rho x_1$
T - leaky bucket, $|T| = \sigma + \rho$
B - temporary packets' buffer
d_p - number of packets which depart during $(x_i, x_{i+1} - x_1]$
Initialization:
$i = 1$
$d\_p = 0$
initialize T with $|T| = \sigma$
Repeat{
fill P and mark its last d_p tokens as invalid
run the leaky bucket regulator during $(x_i - x_1, x_i]$ as follows: {
- when an invalid token from P leaks to T, it is immediately switched to the valid state
and returned to the top of P. Further, no packets from B will depart on behalf of this
token
- when T is full and a valid token (tok) from P leaks to T, the first invalid token from
P, if there is one, is flipped to the valid state. Also, the token (tok) returns to P,
without being dropped as in the standard leaky bucket regulator
}
run the leaky bucket regulator during $(x_i, x_{i+1} - x_1]$ as follows: {
- count the number of departing packets during $(x_{i+1} - 2x_1, x_{i+1} - x_1]$ as d_p
- as before, tokens that are usually dropped by a leaky bucket regulator return to P
- when $i > 1$, enforce that $B(x_i - x_1 - (x_2 - 2x_1 - l), x_i + l] \leq \rho x_1$, for any
$l \leq x_2 - 2x_1$. Note that one simple way of doing this is by implementing a queue
to record the departures from B during $(x_i - x_1 - (x_2 - 2x_1), x_i - x_1]$
}
$i{+}{+}$
}

It is worth mentioning that if $\{x_n\}_n$ is finite (hence $x_N = \infty$), then the algorithm will loop forever, at iteration $i = N$, inside the first leaky bucket procedure. The complexity $O(1)$ of the PSLB regulator is simply due to the same complexity of the standard leaky bucket regulator. Also, depending on the expression of $\{x_n\}_n$, the complexity could be pseudo-$O(1)$.

Theorem 9: The previous algorithm implements a regulator for the PSLB model. Also, a maximal regulator is obtained in the case when $x_{i+1} - x_i$ is constant for any i. The proof for this is given in the Appendix.

Let us take a look now at the unnecessary delays which may appear in the PSLB regulator. These are due to the fact that the algorithm does not implement a maximal regulator for the general case. As one may visualize in the system, the delays could be significant when $x_i - x_1 \gg x_2 - x_1$ as i grows, and the input is idle for a sufficiently long interval. One way to alleviate this phenomenon would be not to count the time while the packet buffer is empty. For example, if the algorithm runs up to step i, and then discovers that during the step $i + 1$ the buffer is empty, it can return safely to step $i + 1$ without proceeding to step $i + 2$. One may argue that there will still be certain cases when a lot of unnecessary delays could exist. However, for regular traffic conforming with the original empirical envelope, the PSLB algorithm performs very well in the sense of minimizing the amount of unnecessary delay introduced by the regulator.

V. Conclusions

In this paper, we introduced a new traffic regulator called the Partially Stopped Leaky Bucket. The PSLB model retains the key properties of the general D-BIND model that allow the input traffic rate to decrease over longer time intervals and the rate to increase temporarily as a function of time. The main advantage of the PSLB model over D-BIND, is that a PSLB traffic regulator can be implemented in constant-time. This is a critical feature for real-time applications requiring quality-of-service.

Currently, we are performing simulation studies to evaluate the performance of the PSLB model with real network traffic and compare it empirically with the other traffic models described above. Also, we are investigating the consequences of implementing regulators that does not possess the maximality property.

References

[1] D. Ferrari and D. Verma. 1990. "A scheme for real-time channel establishment in wide-area networks." *IEEE J. Select. Areas Commun.*, vol. 8, no. 3, 368-379.

[2] D. Verma. 1991. "Guaranteed performance communication in high-speed networks." Ph.D. dissertation, University of California at Berkeley.

[3] R. L. Cruz. 1991. "A calculus for network delay, part I: network elements in isolation." *IEEE Trans. Info. Theory*, vol. 37, no. 1, Jan.:113-131.

[4] R. L. Cruz. 1991. "A calculus for network delay, part II: network analysis." *IEEE Trans. Info. Theory*, vol. 37, no.1, Jan.:132-141.

[5] A. K. Parekh and R. G. Gallager. 1993. "A generalized processor sharing approach to flow control in integrated services networks: the single-node case." *IEEE/ACM Trans. Networking*, vol. 1, no. 3, June:344-357.

[6] C.-S. Chang. 1994. "Stability, queue length and delay of deter-

ministic and stochastic queueing networks." *IEEE Trans. Automat. Contr.*, vol. 39, no. 5, May:913-931.

[7] E. Knightly, R. Mines and H. Zhang. 1994. "Deterministic Characterization and network utilizations for several distributed real-time applications." *Proc. IEEE WORDS '94*, Dana Point, CA.

[8] R. L. Cruz. 1995. "Quality of service guarantees in virtual circuit switched networks." *IEEE J. of Select. Areas in Commun.*, vol. 13, no. 6, Aug.:1048-1056.

[9] E.Knightly and H. Zhang. 1995. "Traffic characterization and switch utilization using a deterministic bounding interval dependent traffic model." *Proc. IEEE INFOCOM'95*, Apr.:1137-1145.

[10] H. Zhang. 1995. "Service disciplines for guaranteed performance service in packet-switching networks." *Proc. IEEE*, vol. 83, no. 10, Oct.:1374-1396.

[11] D. E. Wrege, E. W. Knightly, H. Zhang and J. Liebeherr. 1996. "Deterministic delay bounds for VBR video in packet-switching networks: fundamental limits and practical trade-offs." *IEEE/ACM Trans. Networking*, vol. 4, no. 3, June:352-362.

[12] H. Sariowan. 1996. "A service curve approach to performance guarantees in integrated-service networks." Ph.D. dissertation, Dept. Elect. Comput. Eng., Univ. California at San Diego.

[13] E. Knightly and H. Zhang. 1997. "D-BIND: An accurate traffic model for providing QoS guarantees to VBR traffic." *IEEE/ACM Trans. Networking*, vol. 5, no. 2, Apr.:219-231.

[14] V. Firoiu, J. Kurose and D. Towsley. 1998. "Efficient admission control of piecewise linear traffic envelopes at EDF schedulers." *IEEE/ACM Trans. Networking*, vol. 6, no. 5, Oct.:558-570.

[15] J. Y. Le Boudec. 1998. "Application of network calculus to guarantee service networks." *IEEE Trans. on Info. Theory*, vol. 44, 1087-1096.

[16] C.-S. Chang and R. L. Cruz. 1999. "A time varying filtering theory for constrained traffic regulation and dynamic service guarantees." *Proc. of IEEE INFOCOM* (New York), vol. 1, 63-70.

[17] C.-S. Chang. 2000. Performance guarantees in communication networks. Springer-Verlag, England.

Appendix

I.

Proof: (of Proposition 2)

Let A be an input process such that $A = f$. It is clear that A satisfies the $X_{\min}$ specifications. We will show now that f is subadditive. For simplicity, we will prove only the case when $X_{\min} = 1$ and $S_{\max} = 1$. Let $t, s \geq 0$ and $\alpha = \left\lceil \frac{I}{X_{ave}} \right\rceil$. We can write $t = Iq_1 + r_1, s = Iq_2 + r_2$ and $r_1 + r_2 = Iq_3 + r_3$ with $r_1, r_2 < I$. It follows immediately that $q_3 < 2$. We now have: $f(t+s) \leq f(t) + f(s) \iff \min(\lceil (t+s) \bmod I \rceil, \alpha) + \left\lfloor \frac{t+s}{I} \right\rfloor \alpha \leq \min(\lceil t \bmod I \rceil, \alpha) + \left\lfloor \frac{t}{I} \right\rfloor \alpha + \min(\lceil s \bmod I \rceil, \alpha) + \left\lfloor \frac{s}{I} \right\rfloor \alpha \iff \min(r_3, \alpha) + q_3\alpha \leq \min(r_1, \alpha) + \min(r_2, \alpha)$.

Case 1: $q_3 = 0 \Rightarrow r_1 + r_2 \leq I$. We have to show that: $\min(r_1 + r_2, \alpha) \leq \min(r_1, \alpha) + \min(r_2, \alpha)$. This is true by applying the Proposition 7 for the function $\min(x, \alpha)$ which is both positive and concave.

Case 2: $q_3 = 1$. Without loss of generality let us suppose that $r_1 < r_2$. It is clear that $r_3 \leq r_1 \Rightarrow \min(r_3, \alpha) \leq \min(r_1, \alpha)$ and $\alpha \leq r_2$. It follows that $\min(r_3, \alpha) + \alpha \leq \min(r_1, \alpha) + \min(r_2, \alpha)$ that is exactly our goal.

Now, for $s, t \geq 0$ it follows that $f(t) \leq f(t-s) + f(s) \iff A(t) - A(s) \leq f(t-s)$. It follows that f is an envelope for A.

Finally, simply apply Lemma 1 to complete the proof. ∎

II.

Proof: (of Theorem 9)

Due to the complexity of the proof, we will give the proof just for the particular case when $x_{i+1} - x_i$ is constant for each i. Let us denote $x_{i+1} - x_i = p_2$ and $x_1 = p_1$. We have to prove that the output B of the regulator satisfies (1). The key observation is that for any $ip_2 \leq s \leq (i+1)p_2 + p_1$, where $i \geq 0$, we have $B(ip_2, s] \leq f(s - ip_2)$ and $B(s, (i+1)p_2 + p_1] \leq f((i+1)p_2 + p_1 - s)$. Let now be $0 \leq s < t$ with $s = ip_2 + r_1$ and $t = jp_2 + r_2$ $(r_1, r_2 < p_2)$. If $i = j$, it results immediately that $B(t) - B(s) \leq f(t-s)$. Let us suppose in the following that $i < j$.

Case 1 : $s \geq ip_2 + p_1$ and $t \geq jp_2 + p1$. We have that: $B(s, (i+2)p_2] \leq f(p_1)$, $B((i+2)p_2, (i+3)p_2] \leq f(p_1)$, ..., $B((i+k-1)p_2, (i+k)p_2] \leq f(p_1)$ and $B((i+k)p_2, t] \leq f(p_1)$, where $i + k = j$ (note that k may be 1; in such a case we will have just one inequality: $B(s, t] \leq f(p_1)$). It results that $B(s, t] \leq kf(p_1) = f((k-1)p_2 + p_1)$. Note that the last equality holds because of the general expression of f. Further, $f((k-1)p_2 + p_1) = f(kp_2 - p_2 + p_1) = f((j-i)p_2 - p_2 + p_1) = f(t - s + r_1 - r_2 - p_2 + p_1) \leq f(t-s)$, because $r_1 < p_2$ and $p_1 \leq r_2$.

Case 2 : $s \geq ip_2 + p_1$ and $t < jp_2 + p_1$. We have that: $B(s, (i+2)p_2] \leq f(p_1)$, $B((i+2)p_2, (i+3)p_2] \leq f(p_1)$, ..., $B((i+k-1)p_2, (i+k)p_2] \leq f(p_1)$ and $B((i+k)p2, t] \leq f(t - (i+k)p_2)$, where $i + k = j$. Again, k may be 1, and in such a case we would have just one inequality: $B(s, t] \leq f(t-s)$ which is satisfied. It follows that $B(s, t] \leq (k-1)f(p_1) + f(t - (i+k)p_2) = f((k-1)p_2 + t - (i+k)p_2) = f(t - (i+1)p_2) = f(t - s + r_1 - p_2) \leq f(t-s)$, because $r_1 < p_2$.

Case 3 : $s < ip_2 + p_1$ and $t \geq jp_2 + p_1$. We have that: $B(s, ip_2 + p_1] \leq f(p_1 - r_1)$, $B(ip_2 + p_1, (i+2)p_2] \leq f(p_1)$, ...,$B((i+k-1)p_2, (i+k)p_2] \leq f(p_1)$ and $B((i+k)p_2, t] \leq f(p_1)$, where $i + k = j$. If $k = 1$, we would have the inequalities $B(s, ip_2 + p_1] \leq f(p_1 - r_1)$ and $B(ip_2 + p_1, t] \leq f(p_1)$. In both cases it follows that: $B(s, t] \leq kf(p_1) + f(p_1 - r_1) = f(kp_2 + p_1 - r_1) = f(t - s + p_1 - r_2) \leq f(t-s)$, because $p_1 < r_2$.

Case 4 : $s < ip_2 + p_1$and $t < jp_2 + p_1$. We have that: $B(s, ip_2 + p_1] \leq f(p_1 - r_1)$, $B(ip_2 + p_1, (i+2)p_2] \leq f(p_1)$, $B((i+2)p_2, (i+3)p_2] \leq f(p_1)$, ..., $B((i+k-1)p_2, (i+k)p_2] \leq f(p_1)$ and $B((i+k)p2, t] \leq f(t - (i+k)p_2)$, where $i + k = j$. If $k = 1$, we would have just the inequality: $B(s, t] \leq f(t-s)$ which is satisfied due to the algorithm.

Otherwise $(k > 1)$, it follows that: $B(s, t] \leq (k-1)f(p_1) + f(p_1 - r_1) + f(r_2)$.

Case 4.1 : $r_2 \leq r_1$. We have that $(k-1)f(p_1) + f(p_1 - r_1) + f(r_2) = f((k-1)p_2 + p_1 + r_2 - r_1) = f((j-i-1)p_2 + p_1 + t - jp_2 - s + ip_2) = f(t - s + p_1 - p_2) \leq f(t-s)$.

Case 4.2 : $r_1 < r_2$. We have that $(k-1)f(p_1) + f(p_1 - r_1) + f(r_2) = f(kp_2 + r_2 - r_1) = f(t-s)$.

We now give a proof sketch of the maximal regularity statement. Let us suppose, by contradiction, that there is another output B' such that $B < B'$ and let be $t = \inf\{s > 0 \mid B(s) < B'(s)\}$. Without loss of generality, we suppose that $B'(t) = B(t) + 1$. We will prove here just the case when $t \in (ip_2 + p_1, (i+1)p_2]$, where $i > 0$. The first observation is that $T(t) = 0$, because otherwise we would have a different output B.

Case 1 : $|P(t)| = 0$ (there is no more token available in P at time t).

In such a case, we will have a contradiction in the definition of the regulator (for B') over the interval $(ip_2, t]$, in the sense that $B'(t) - B'(ip_2) > \sigma + \rho(t - ip_2)$.

Case 2: $|P(t)| > 0$

Let $t_1 = \sup\{s < t \mid T(s) \geq \sigma\}$.

Case 2.1 : $t_1 > ip_2 + p_1$

In such a case we have that $B(t) - B(t_1) = \sigma + \rho(t - t_1) \Rightarrow B^{'}(t) - B^{'}(t_1) = \sigma + \rho(t - t_1) + 1$ which contradicts the regularity condition for B'.

Case 2.2 : $t_1 \in (ip_2, ip_2 + p_1]$

Using similar reasoning as above, we will obtain a contradiction on the regularity of B' over an interval $(s, t]$, where $s > (i-1)p_2$. ■