# Traffic Descriptors for the Configuration of Leaky Bucket Policing Mechanisms with General Two-State On/Off Arrival Processes

D. S. Holtsinger
H. G. Perros

Center for Communications and Signal Processing
Department of Computer Science
North Carolina State University

TR-93/16
August 1993

# TRAFFIC DESCRIPTORS FOR THE CONFIGURATION OF LEAKY BUCKET POLICING MECHANISMS WITH GENERAL TWO-STATE ON/OFF ARRIVAL PROCESSES[1]

Douglas S. Holtsinger , Harry G. Perros

Department of Computer Science, and

Center for Communications and Signal Processing

North Carolina State University

Raleigh, NC 27695

## Abstract

The leaky bucket has been widely studied as a policing mechanism for ATM networks, but very little effort has been focused on developing simple traffic descriptors which can be used for configuring the leaky bucket in practical applications. In this paper, we analyze the cell loss probability of the leaky bucket policing mechanism, assuming that the traffic source behaves as a two-state on/off type arrival process with arbitrary and independent distributions for the burst and silence periods. We demonstrate that the leaky bucket policing mechanism can be configured with substantially smaller token pool sizes as limits are placed on the maximum length of the burst period and the minimum length of the silence period. While these limits will lower the amount of traffic "burstiness", we show that traditional indicators of burstiness, such as the squared coefficient of variation of interarrival times ($C^2$) and the mean burst length, fail to capture changes in burstiness due to changes in the maximum length of the burst period. This suggests that the

[1]Supported in part by BellSouth, GTE Corporation, and NSF and DARPA under cooperative agreement NCR-8919038 with the Corporation for National Research Initiative.

maximum burst length of the source should be incorporated into the traffic descriptor as a measure of burstiness. We also propose the *dual leaky bucket* policing mechanism as a suitable device for policing the behavior of bursty traffic sources, and we approximately analyze the cell loss probability.

# 1 Introduction

Policing mechanisms have been proposed as part of a preventive congestion control technique for controlling the behavior of the traffic source at the user-network interface of an ATM network [2, 11, 23]. A policing mechanism (also known as the Usage Parameter Control (UPC) in CCITT terminology [24]) ensures that the traffic source conforms to the traffic contract which has been negotiated between the user and the network at the time of a connection setup.

The traffic characteristics of the source and the quality-of-service requirements of the source form the basis of a traffic contract between the source and the network, which is enforced by the policing mechanism. The parameters of the traffic contract will typically include limits on the mean traffic rate, peak rate, and some notion of the "burstiness" of the traffic source [2]. A burstiness parameter would allow the traffic source to temporarily exceed the negotiated mean rate for a short period of time, while staying within the peak rate and the long-term mean rate required by the traffic contract.

As long as the source complies with the limits of the traffic contract, the policing mechanism should remain transparent to the source. Transparency is usually expressed in terms of a cell loss probability of less than $10^{-9}$ [7, 19]. If the source exceeds the limits of the traffic contract, then the policing mechanism should prevent the non-compliant traffic source from inducing congestion conditions in the network. The policing mechanism may either drop, delay, or "mark" cells for later removal from the system during congestion conditions.

Many policing mechanisms have been proposed and analyzed in the literature (for example, see [3, 17, 19]), but leaky bucket-type policing mechanisms appear to have received the most attention. Many authors have examined the ability of the leaky bucket to police non-compliant traffic sources assuming that the source can be modeled with a stochastic process [1, 4, 6, 7, 10, 12, 16, 18, 20, 21].

In this paper, we examine the assumptions regarding the characteristics of the on/off traffic source which have been used in previous numerical analyses of the leaky bucket [5, 7, 19]. We are mostly interested in examining the effect of the arrival process characteristics on the configuration of the leaky bucket, rather than on examining the effectiveness of the leaky bucket for policing non-compliant traffic sources, since that has already been widely considered. Ultimately, our goal is to develop simple traffic descriptors which can be used for configuring the leaky bucket in practical applications. This is a desirable goal since the network user may not be able to specify the detailed stochastic behavior of the traffic source, such as the distribution of the burst and silence periods.

We demonstrate that the leaky bucket policing mechanism can be configured with substantially smaller token pool sizes as limits are placed on the maximum length of the burst period and the minimum length of the silence period. Traditional indicators of burstiness, such as the squared coefficient of variation of interarrival times ($C^2$) and the mean burst length, fail to capture changes in burstiness due to changes in the maximum length of the burst period. This suggests that the maximum burst length of the source should be incorporated into the traffic descriptor as a measure of burstiness. We also propose the *dual leaky bucket* policing mechanism as a suitable device for policing the behavior of bursty traffic sources, and we approximately analyze the cell loss probability.

The outline of this paper is as follows. In section 2 of this paper, we analyze the leaky bucket policing mechanism, and in section 3, we propose an improved policing mechanism known as the *dual leaky bucket.* We provide an approximate analysis of the cell loss probability, and compare our numerical results against simulation. Finally in section 4, we present our conclusions.

# 2 The Leaky Bucket Policing Mechanism

## 2.1 Description of the Leaky Bucket Policing Mechanism

In the leaky bucket policing mechanism shown in figure 1, one token is added to the token pool at the beginning of each slot, up to the capacity of the token pool, which is of size $K$. Tokens which arrive to a full token pool are dropped. Cells arrive at the end of the slot, at a maximum of one cell per slot. A cell is required to consume $N$ tokens before it may pass through the leaky bucket. If an arriving cell finds less than $N$ tokens

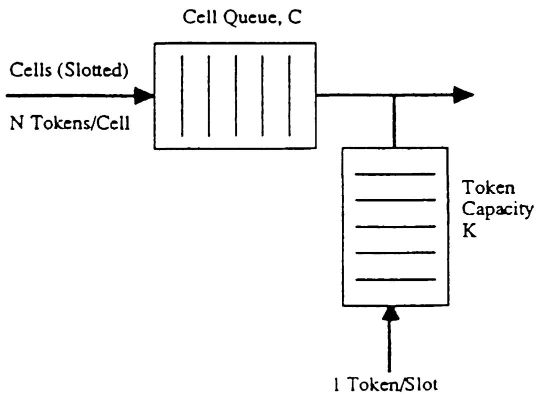


Figure 1: Leaky Bucket Policing Mechanism

in the token pool, then the cell waits in the cell queue of size $C$. When at least $N$ tokens accumulate in the token pool, a waiting cell will consume $N$ tokens and depart from the leaky bucket. If an arriving cell finds the cell queue full, then the cell is dropped.

In our analysis, we assume that a dropped cell consumes all of the remaining tokens in the token pool. Once a cell is dropped, all cells which arrive consecutively (in the same burst of cells) after the dropped cell will be lost (assuming $N > 1$). This particular feature has been shown to improve the leaky bucket's ability to restrict non-compliant traffic sources [16].

## 2.2 Analysis of the Leaky Bucket Policing Mechanism

Our analysis assumes that the source behaves as a two-state on/off arrival process with arbitrary distributions for the time spent in each state (see figure 2). When the process is in the burst state, one cell is generated per slot. During the silence state, no cells are generated. The times spent in the burst and silence states are assumed to be independent and arbitrarily distributed.

We define the following random variables:

$$\begin{aligned} B &= \text{Length of the burst period in slots} \\ S &= \text{Length of the silence period in slots} \end{aligned}$$

and the associated probability distributions are

$$\begin{aligned} b_i &= P(B = i) \\ s_i &= P(S = i) \end{aligned}$$

Our analysis of the leaky bucket is based upon a Markov chain embedded immediately before the arrival process moves from the silence state to the burst state. The state of the system at the beginning of the $n$th burst is represented by the random variable $X_n$, which is the number of slots remaining until the token pool is completely filled with tokens, assuming that no cells arrive. Thus the state of the system can be described by the equation

$$X_n = (\text{Number of Cells}) \cdot N + K - (\text{Number of Tokens})$$

where $X_n$ takes on values from 0 to $C\,N + K$.

The above system has a state space of size $(C\,N + K + 1)$, which may be too large to solve efficiently for large values of $N$ and $C$. Therefore we have developed an approximate solution method which relies on aggregating the states, and results in reducing the size of the state space to $\lfloor \frac{CN+K}{N-1} \rfloor + 1$.

The reduction in the size of the state space is not large for many of the numerical cases which we shall consider in this paper, since not every configuration of the leaky bucket which we shall consider requires a large value of $N$. Therefore, we shall rely upon the approximate solution method in the numerical results only when appropriate. Due to space considerations, we leave off the details of the exact analysis, which can be found in [15].

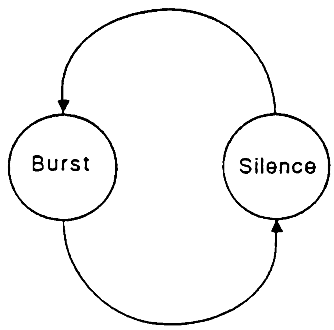


Figure 2: Two-state on/off Arrival Process

## 2.3 Approximate Solution Method

The state of the system is represented by the number of slots remaining until the token pool is filled with tokens. During each slot in which an arriving cell passes through the leaky bucket, the cell consumes $N$ tokens and one token is added to the token pool, except during a slot when the token pool is filled with tokens, in which case no token is added to the token pool. Thus during each slot in which a cell passes through the leaky bucket, the number of tokens decreases by either $N-1$ or $N$. As $N$ increases, which is of primary interest in the approximation, the difference in the decrease of the number of tokens becomes smaller, and the decrease is approximated as $N-1$ for all cases.

During the burst period, the system moves from state $i$ to state $(i+N-1)$ at each slot. The approximate solution method consists of lumping the states together to construct an aggregate Markov Chain. The lumping of states is done such that the aggregated system moves from aggregate state $i$ to aggregate state $(i+1)$ at each slot during the burst period. Thus the increase in the aggregate state during the burst period can be written directly in terms of the burst period distribution, $b_i$.

The aggregate states and their corresponding original states are shown in table 1. The largest aggregate state $Z$ is defined as the largest integer $Z$ such that $(N-1)Z < C\,N+K$, where $Z$ can be interpreted as the largest burst of cells which can pass through the leaky bucket. The aggregate states do not lend themselves to any interpretation, as they are only intended as a means for calculating the cell loss probability.

| Original States | Aggregate State |
|---|---|
| $0,\ldots,N-2$ | 0 |
| $N-1,\ldots,2N-3$ | 1 |
| $\ldots$ | $\ldots$ |
| $(N-1)i,\ldots,(i+1)(N-1)-1$ | $i$ |
| $\ldots$ | $\ldots$ |
| $(N-1)Z,\ldots,C\,N+K$ | $Z$ |

Table 1: Aggregation of States

During the silence period, a token arrives during every slot, and the state of the aggregate system does not decrease during every slot. Therefore we need an approximation to determine the decrease in the aggregrate state during the silence period. We define the following random variable:

$$\tilde{S} = \text{Potential decrease in the aggregate state during the silence period,}$$

and the associated probability distribution is

$$\tilde{s}_i = P(\tilde{S} = i).$$

Given that the system is in aggregate state $j$, the probability that the system is actually in one of the original states $(N-1)j$ through $(j+1)(N-1)-1$ is assumed to be uniformly distributed among these $(N-1)$ original states. This approximation should become better as the number of original states increases, since the relative difference between state probabilities of adjacent states has been observed to become smaller. Using this approximation, we can write

$$\tilde{s}_i = \sum_{j=0}^{N-2}\left[\frac{j}{N-1}s_{(i-1)(N-1)+j} + \frac{N-1-j}{N-1}s_{i(N-1)+j}\right] \quad i = 0,1,2,\ldots \tag{1}$$

where we recall that $s_i$ is the probability that the silence period lasts $i$ slots. While the probabilities $\tilde{s}_i$ cannot be obtained in a closed-form expression, it is only necessary to compute a finite number of these probabilities.

Let $\widetilde{X}$ denote the random variable representing the state in the aggregated system, and let $\tilde{P}_k$ denote the probability that $(\widetilde{X}=k)$ at the embedded point. We define the matrix of transition probabilities $\tilde{Q} = [\tilde{q}_{l,i}]$, where

$$\tilde{q}_{l,i} = P(\widetilde{X}_{n+1} = i \mid \widetilde{X}_n = l).$$

We can solve a linear system of equations $\tilde{P} = \tilde{Q}\tilde{P}$ by replacing one of the equations with the condition that $\sum_k \tilde{P}_k = 1$.

The state transition probabilities $\tilde{q}_{l,i}$ can be written as

$$\tilde{q}_{l,i} = \bar{b}_{Z-l}\tilde{s}_{Z-i} + \sum_{j=1}^{Z-l} b_j\tilde{s}_{l+j-i} \quad l = 0,\ldots,Z \quad i = 1,\ldots,Z \tag{2}$$

$$\tilde{q}_{l,0} = \bar{b}_{Z-l}\tilde{s}_{Z} + \sum_{j=1}^{Z-l} b_j\breve{s}_{l+j} \quad l = 0,\ldots,Z \tag{3}$$

where we define

$$\bar{b}_i = P(B > i) = 1 - \sum_{j=1}^{i} b_j$$

$$\breve{s}_i = P(\tilde{S} > i) = 1 - \sum_{j=1}^{i}\tilde{s}_j$$

## 2.4 Cell Loss Probability

Let $P_{\text{loss}}$ denote the cell loss probability, defined as the fraction of cells which are dropped by the leaky bucket. The cell loss probability is determined assuming that the state probabilities are computed using the approximate solution method. A similar method can be used in the exact solution of the state probabilities.

Suppose that the state of the system at the beginning of the $n$th burst is $\widetilde{X}_n = k$. Given that a randomly selected cell arrives in a burst of length $i$ in which at least one cell is dropped in the burst (i.e. $i + k > Z$), the probability that the randomly selected cell is dropped can be written as

$$\frac{i + k - Z}{i}. \tag{4}$$

The probability that a randomly selected arriving cell arrives during a burst of length $i$, is written as

$$\frac{i\,P(B = i)}{\sum_i iP(B = i)} = \frac{ib_i}{E[B]}, \tag{5}$$

where $E[B]$ is the expected length of a burst period in slots.

Let $P_{\text{loss}}(k)$ represent the probability that a randomly selected cell is lost in a burst, given that the system state is $k$ at the beginning of a burst period. Using equations (4) and (5), we write,

$$\begin{aligned} P_{\text{loss}}(k) &= \sum_{i=Z-k+1}^{\infty} \frac{ib_i}{E[B]} \frac{i+k-Z}{i} \\ &= \sum_{i=Z-k+1}^{\infty} \frac{(i+k-Z)b_i}{E[B]}. \end{aligned}$$

The cell loss probability can now be written as

$$P_{\text{loss}} = \sum_{k=0}^{Z} \tilde{P}_k P_{\text{loss}}(k),$$

which can be expressed as

$$P_{\text{loss}} = 1 - \frac{Z - E[\widetilde{X}] - \sum_{k=0}^{Z} \tilde{P}_k \sum_{i=1}^{Z-k} (Z-k-i)b_i}{E[B]}, \tag{6}$$

where $E[\widetilde{X}] = \sum_k k \cdot \tilde{P}_k$, and $E[B]$ is the expected length of a burst period.

## 2.5 The Truncated Geometric Arrival Process

For our numerical results, we use a discrete-time alternating renewal arrival process with parameters $p$, $q$, $T_{\max}$ and $T_{\min}$, called the truncated geometric arrival process (TGAP). The arrival process alternates between two states, the burst and silence states. While the process is in the burst state, one cell is generated per slot, and while in the silence state, no cells are generated. The distribution of the time that is spent in the burst state is defined as

$$P(B = i) = \begin{cases} (1-p)p^{i-1} & i = 1, \ldots, T_{\max} - 1 \\ p^{T_{\max}-1} & i = T_{\max} \\ 0 & \text{otherwise} \end{cases}$$

with a mean burst period duration of

$$E[B] = \frac{1 - p^{T_{\max}}}{1-p}. \tag{7}$$

The parameter $T_{\max}$ represents the maximum number of consecutive cells that may arrive during a burst period. The silence period distribution for the arrival process is defined as

$$P(S = i) = \begin{cases} 0 & i = 0, \ldots, T_{\min} - 1 \\ (1-q)q^{i-T_{\min}} & i = T_{\min}, \ldots, \infty \end{cases}$$

with a mean silence period duration of

$$E[S] = \frac{q}{1-q} + T_{\min}. \tag{8}$$

The parameter $T_{\min}$ represents the minimum amount of time that the arrival process must spend in the silence period. If we let $T_{\max} \to \infty$ and set $T_{\min} = 1$ in the truncated geometric arrival process, then we have a two-state Markov modulated arrival process defined by:

$$\begin{aligned} P(B=i) &= (1-p)p^{i-1} \qquad i=1,\ldots,\infty \\ P(S=i) &= (1-q)q^{i-1} \qquad i=1,\ldots,\infty \end{aligned}$$

which has been used in previous analyses of the leaky bucket [5, 7, 19].

The probability that a randomly selected cell arrives during a burst period of length $i$ can be written as (see equation (5))

$$\frac{iP(B=i)}{E[B]} \qquad i=1,\ldots,T_{\max},$$

where $E[B]$ is the mean burst period duration, given by equation (7). Since the randomly selected cell can arrive in any position during the burst period with equal probability, the next slot will contain an arrival with probability

$$\frac{i-1}{i}.$$

Let the random variable $X$ denote the interarrival time in slots. Now we can write

$$\begin{aligned} P(X=1) &= \sum_{i=1}^{T_{\max}} \frac{iP(B=i)}{E[B]}\frac{i-1}{i} \\ &= 1-\frac{1}{E[B]}. \end{aligned}$$

The probability that $X>1$ is then given by

$$\frac{1}{E[B]}.$$

This can be interpreted as the probability that a randomly selected cell is the last cell of a burst. We can now write the probability that $X=j$ as

$$P(X=j) = \begin{cases} 1-\frac{1}{E[B]} & ,j=1 \\ \frac{P(S=j-1)}{E[B]} & ,j=T_{\min}+1,\ldots,\infty, \end{cases} \tag{9}$$

where the random variable $S$ is the silence period duration. The second moment of the interarrival time is given by

$$
\begin{aligned}
E[X^2] &= \sum_{i=1}^{\infty} i^2\, P(X = i) \\
&= 1 - \frac{1}{E[B]} + \frac{1}{E[B]} \sum_{i=T_{\min}}^{\infty} (i+1)^2\, P(S = i) \\
&= 1 - \frac{1}{E[B]} + \frac{1}{E[B]} (E[S^2] + 2E[S] + 1) \\
&= \frac{E[B] + E[S^2] + 2E[S]}{E[B]},
\end{aligned}
$$

where $E[S^2]$ is the second moment of the silence period duration. The squared coefficient of variation of the interarrival time is

$$C^2 = \frac{E[B]E[S^2] - E^2[S]}{(E[B] + E[S])^2}.$$

From equation (9), it can be seen that the interarrival time distribution only depends upon the mean burst period and the distribution of the silence period. Thus, the squared coefficient of variation is unaffected by the distribution of the burst period duration, except through the mean. In fact, it can be seen from equation (9) that all higher order moments of the interarrival time distribution are independent of the higher order moments of the burst period distribution for any two-state on/off type arrival process with independent on/off periods.

## 2.6 Numerical Results

In previous numerical analyses of the leaky bucket, the traffic source is assumed to behave as a two-state on/off process that spends a geometrically distributed amount of time in each state [5, 7, 19]. Below we examine the effect of this assumption on the performance of the leaky bucket.

We vary the values of $T_{\max}$ and $T_{\min}$ in the truncated geometric arrival process while keeping the mean burst and silence period durations fixed, and examine the effect on the cell loss probability of the leaky bucket. Figures 3 and 4 show the cell loss probability as a function of the token pool size $K$ for values of $N = 2$ and $N = 8$ (tokens/cell), respectively, assuming that the mean burst period length is $E[B] = 10$ and the mean silence period length is $E[S] = 90$. We assume that the cell queue size $C = 0$.

| Source | Maximum Burst Length | Minimum Silence Length | $C^2$ |
|---|---|---|---|
| (a) | $\infty$ | 1 | 15.3 |
| (b) | 65 | 1 | 15.3 |
| (c) | 87 | 30 | 7.47 |
| (d) | 65 | 45 | 5.4 |

Table 2: Arrival Process Characteristics

We compare the cell loss probability that results when the source behaves as a Markov modulated arrival process, $(T_{\max}, T_{\min}) = (\infty, 1)$, and as the truncated geometric arrival process (TGAP). For a given cell loss probability, the TGAP requires a substantially smaller token pool size than the Markov modulated arrival process. For a given value of $N$, a smaller token pool size allows the leaky bucket to police non-compliant sources more effectively, since less traffic "burstiness" is able to pass through the leaky bucket [7, 19]. For example, if the cell loss probability requirement is less than $10^{-5}$ and the value of $N$ is 8, then imposing a maximum burst limit of $T_{\max} = 65$ cells reduces the token pool size from 3040 to 2080.

Table 2 shows the squared coefficient of variation ($C^2$) of the interarrival times for the four arrival processes used in figures 3 and 4. Using the previous results, we note that all moments of the interarrival time distribution for the first and second arrival processes in table 2 are identical, and the mean burst and silence periods are identical, and yet the cell loss probability is substantially lower for the second arrival process. Some authors have relied upon the $C^2$ of the interarrival times [20, 22], or the mean burst and silence lengths [7, 13, 16, 18, 19, 21] to describe traffic "burstiness" of a two-state on/off type traffic source. Our results indicate that these measures by themselves are not sufficient for capturing the effects of traffic burstiness, since they do not account for the variation of the burst period. To solve this problem, we suggest the inclusion of the maximum burst length into the traffic descriptor. The maximum burst length is a simple measure of the burst period variation which the network user can specify rather easily, unlike other measures such as the second moment of the burst period duration [13], or the generalized peakedness [9].

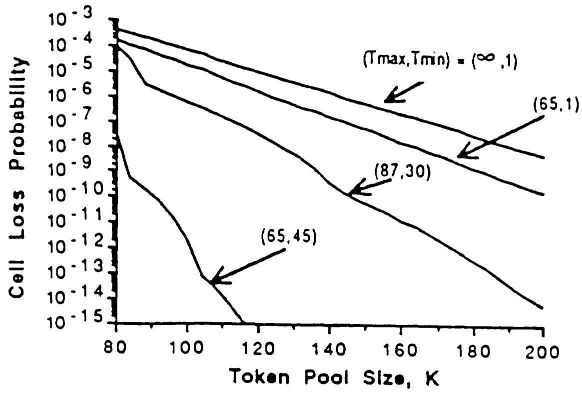


Figure 3: Cell Loss Probability, E[B]=10, E[S]=90, N=2

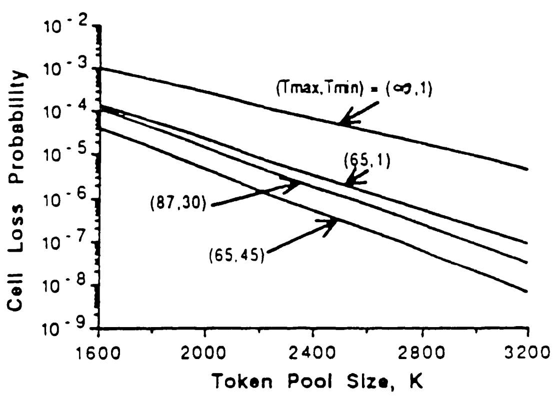


Figure 4: Cell Loss Probability, E[B]=10, E[S]=90, N=8

# 3 The Dual Leaky Bucket Policing Mechanism

A frequent criticism of leaky bucket-type policing mechanisms is that they are not effective for policing both the mean rate and the burstiness of a traffic source at the same time [7, 19]. The peak rate, mean rate, and burstiness of a source have been shown to have a significant effect on the cell loss probability and cell delay when several sources are statistically multiplexed together [8, 14].

In this section, we propose a *dual leaky bucket* policing mechanism, and using the methodology developed in the previous section, approximately analyze the cell loss probability of the dual leaky bucket. The dual leaky bucket is an attempt to overcome the disadvantages of existing leaky bucket-type mechanisms by combining two token pools and operating them as shown in figure 5. One token pool is configured to police the mean rate of the traffic source, and the other token pool is configured to police the burstiness of the source.

The dual leaky bucket consists of two token pools, TP1 and TP2, which have a capacity of $K_1$ and $K_2$ tokens, respectively. At the beginning of each slot, one token is added to token pool TP1 and token pool TP2, up to the capacity of each token pool. Cells arrive at the end of the slot, at a maximum of one cell per slot. A cell is required to consume $N_1$ tokens from token pool TP1 and $N_2$ tokens from token pool TP2 before it may pass through the dual leaky bucket. If one or both token pools has insufficient tokens, then the arriving cell is dropped, otherwise the cell passes through the dual leaky bucket. If insufficient tokens are found in a token pool, then all of the tokens in that pool are consumed. This policy has been demonstrated to improve the ability of the leaky bucket to restrict non-compliant traffic sources [16].

In the following section, we approximately analyze the loss probability of cells for the dual leaky bucket policing mechanism, assuming that the source behaves as a two-state on/off arrival process with arbitrary distributions for the time spent in each state.

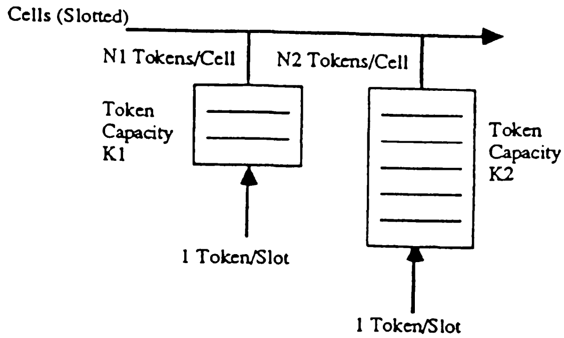


Figure 5: Dual Leaky Bucket Policing Mechanism

## 3.1 Analysis of the Dual Leaky Bucket Policing Mechanism

The approximate analysis of the dual leaky bucket begins by analyzing the two token pools in isolation. We assume that the two token pools operate independently from each other, and we analyze the state probability distribution of each token pool. For the purposes of analyzing the cell loss probability, these state probability distributions are assumed to be independent. We can obtain the state probability distributions of the token pools through the analysis of the previous section.

Through a comparison to simulation, we have found that this is a good approximation if the two token pool sizes are disparate. If the token pool sizes (and $N_1, N_2$) are identical, then the token pools will operate identically, and the approximation will not be accurate. But as the token pool sizes (and $N_1, N_2$) begin to differ, the states of the two token pools will change over different time scales, and the approximation becomes more accurate.

Due to space considerations, we leave off the details of the analysis, which can be found in [15]. The cell loss probability is analyzed in a similar manner as before in section 2.4. Each token pool is configured separately to achieve a cell loss probability which is less than half the required cell loss probability in the traffic contract. By adding the cell loss probabilities contributed by each token pool, this method will give an upper bound to the cell loss probability of the dual leaky bucket under compliant traffic source conditions [15].

## 3.2 Comparison with Simulation

We use simulation to validate the approximate model developed for the cell loss probability of the dual leaky bucket in section 3.1. Due to the limitations of simulation, it becomes unreasonable to simulate the leaky bucket using a cell loss probability requirement of less than $10^{-9}$. Therefore we have selected $10^{-5}$ as the cell loss probability requirement. Our simulations were carried out using the method of independent replications, with 10 replications at 16 million cells per replication. Table 3 describes the configuration of the dual leaky leaky bucket used in the experiments.

| | |
|---|---|
| Tokens/Cell $N_1$ | 8 |
| Token Pool Size $K_1$ | 1008 |
| Tokens/Cell $N_2$ | 2 |
| Token Pool Size $K_2$ | 52 |

Table 3: Configuration of the Dual Leaky Bucket

Figures 6, and 7 show a comparison between the simulation (dotted line) and the cell loss probability predicted by the model (solid line), using a TGAP arrival process. The 95% confidence intervals are small enough that they are obscured by the graph marks. In each plot, the maximum burst length is held constant at $T_{max} = 90$ and the minimum silence length is held at $T_{min} = 20$.

In figure 6, the mean burst length of the compliant source is increased from 5 to 15 while the mean silence length is held fixed at $E[S] = 45$. This also increases the mean arrival rate $\rho$, or the expected number of arrivals per slot. In figure 7, the mean burst length and the mean silence length are increased, while the mean arrival rate is held fixed at $\rho = 0.1$. Several comparisons of the approximation with simulation show that the worst case relative error does not exceed 0.15.

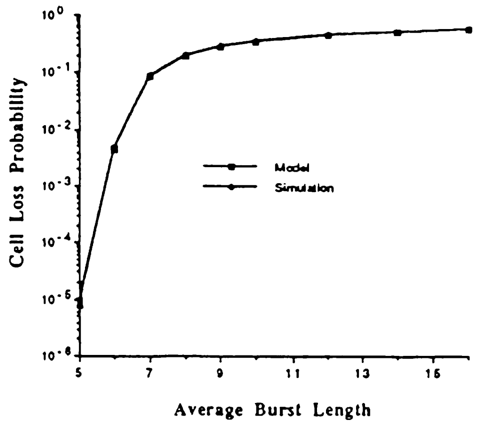


Figure 6: Cell Loss Probability, Average Silence Length $E[S] = 45$

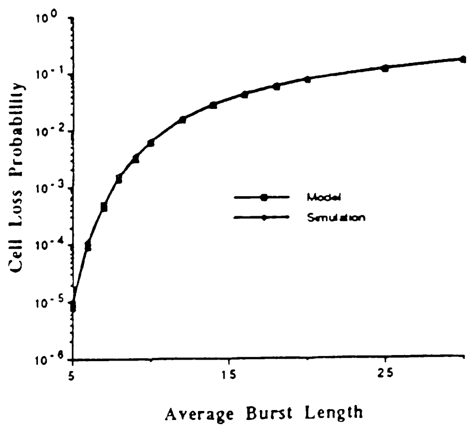


Figure 7: Cell Loss Probability, Mean Arrival Rate $\rho = 0.1$

## 3.3 Numerical Results

In this section, we compare the ability of the dual leaky bucket policing mechanism to restrict non-compliant traffic sources against a leaky bucket configured with only a single token pool, as described in section 2.

First we configure the dual leaky bucket and the single leaky bucket to achieve a cell loss probability requirement of less than $10^{-5}$ when the compliant traffic source is applied to each leaky bucket. We use a compliant traffic source given in table 4. We use three different configurations of the single leaky bucket which meet a cell loss probability requirement of less than $10^{-5}$, corresponding to two extreme values of $N$, and one value of $N$ between the two extremes. For each value of $N$, we use the smallest possible value of the token pool size $K$ that will result in a cell loss probability of less than $10^{-5}$. The configurations of the dual leaky bucket and the single leaky buckets are given in table 5.

| Average Burst Length | Average Silence Length | Mean Arrival Rate | $C^2$ |
|---|---|---|---|
| 10 | 90 | 0.1 | 9.6 |

Table 4: Compliant traffic source

Next we compare these three single leaky buckets against one configuration of the dual leaky bucket for restricting the source when it becomes non-compliant. The dual leaky bucket is configured according to the method outlined in section 3.1. All traffic sources are truncated geometric arrival processes, and have a maximum burst length of $T_{max} = 90$ and a minimum silence length of $T_{min} = 20$.

| $K_1$ | $N_1$ | $K_2$ | $N_2$ |
|---|---|---|---|
| 2400 | 8 | – | – |
| 600 | 5 | – | – |
| 92 | 2 | – | – |
| 2600 | 8 | 96 | 2 |

Table 5: Leaky bucket configurations

Figures 8 and 9 show the cell loss probability for the three different configurations of the single leaky bucket and the dual leaky bucket when a non-compliant traffic source is used.

In figure 8, the mean silence length is held fixed at $E[S] = 90$ and the mean burst length is varied from 10 to 40. The dual leaky bucket and the single leaky bucket with a token pool size of $K = 2400$ have nearly identical performance for restricting the non-compliant traffic source. The dual leaky bucket performs much better than the other two configurations of the single leaky bucket. In figure 9, the mean arrival rate is decreased from $\rho = 0.1$ to $\rho = 0.02$, and the mean burst length and the mean silence length are increased. Here the dual leaky bucket performs much better than the single leaky bucket with a token pool size of $K = 2400$, and it performs about as well as the other two configurations of the single leaky bucket. Additional numerical results can be found in [15].

Since the behavior of the non-compliant traffic source cannot be predicted, the network designer cannot select a leaky bucket configuration with a single token pool that will perform best for restricting the source under all non-compliant traffic conditions. While a particular single leaky bucket configuration may perform slightly better than the dual leaky bucket for a given non-compliant traffic source (for example, figure 9, $K = 92$, $N = 2$), the dual leaky bucket can perform far better than the same single leaky bucket under different non-compliant traffic conditions (figure 8). The numerical results here indicate that the dual leaky bucket performs about as well as the *best performing* configuration of the single leaky bucket for restricting a given non-compliant traffic source.

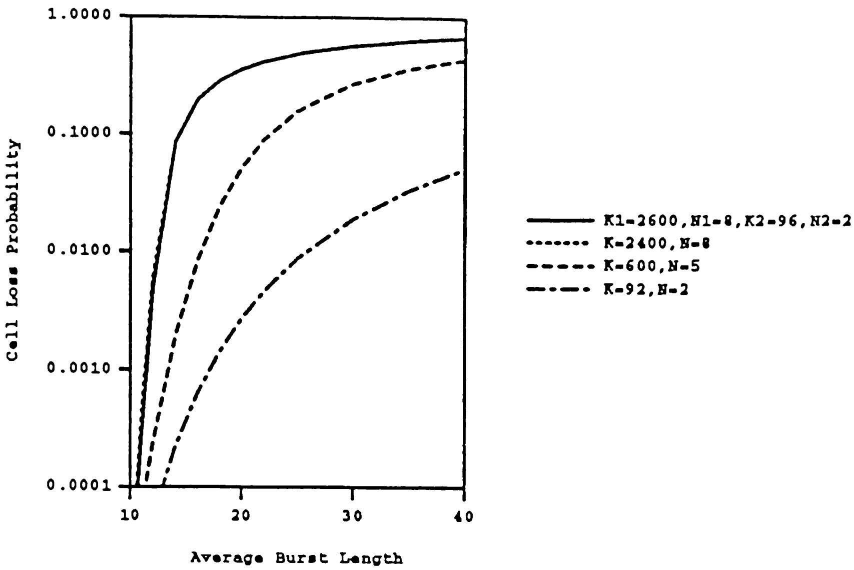


Figure 8: Cell Loss Probability, $E[S] = 90$

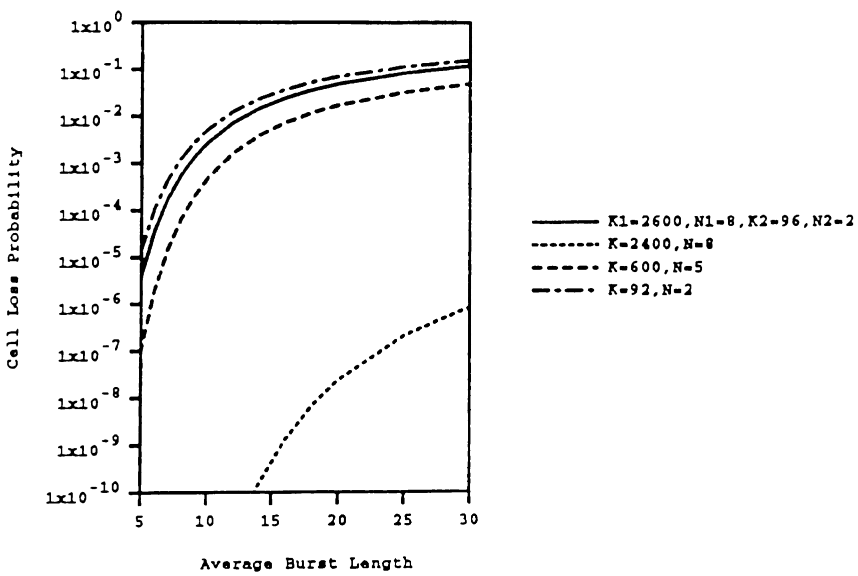


Figure 9: Cell Loss Probability, $\rho = 0.02$

# 4 Summary and Conclusions

We have provided an analysis of the leaky bucket policing mechanism, assuming that the source behaves as a two-state on/off type arrival process. We focused on examining the effect of the arrival process characteristics on the configuration of the leaky bucket. Our goal is to develop simple traffic descriptors for configuring the leaky bucket in practical applications. We found that traditional descriptors of traffic burstiness, such as the squared coefficient of variation of the interarrival times and the mean burst length, are not sufficient for capturing the effect of traffic burstiness. Our results suggest that the maximum burst length of the source should be incorporated into the traffic contract for the purposes of configuring the leaky bucket policing mechanism.

The call admission mechanism must allocate network resources based upon the worst-case traffic which can pass through the policing mechanism, and as the token pool size decreases, less traffic burstiness can pass through the leaky bucket. If a network user can specify limits on the maximum burst length and the minimum silence length of the source, then the token pool size required to achieve a given cell loss probability will be lower than if these characteristics are left unspecified. As a consequence of the smaller token pool size, the network does not have to over-allocate as much resources to the connection to safeguard against congestion conditions caused by a non-compliant source.

We have also shown that the dual leaky bucket provides superior policing capabilities over a range of non-compliant traffic sources. Our numerical results here indicate that the dual leaky bucket performs about as well as the *best performing* configuration of the single leaky bucket for restricting a given non-compliant traffic source.

# References

[1] H. Ahmadi, R. Guérin, and K. Sohraby. Analysis of leaky bucket access control mechanism with batch arrival process. In *IEEE Globecom*, pages 344–349, 1990.

[2] J.J. Bae and T. Suda. Survey of traffic control schemes and protocols in ATM networks. *Proceedings of the IEEE*, 79(2):170–189, February 1991.

[3] K. Bala, I. Cidon, and K. Sohraby. Congestion control for high speed packet switched networks. In *IEEE Infocom*, pages 520–526, 1990.

[4] A.W. Berger. Performance analysis of a rate-control throttle where tokens and jobs queue. *IEEE Journal on Selected Areas in Communications*, 9(2):165–170, February 1991.

[5] A.W. Berger, A.E. Eckberg, T. Hou, and D.M. Lucantoni. Performance characterizations of traffic monitoring, and associated control, mechanisms for broadband "packet" networks. In *IEEE Globecom*, pages 350–354, 1990.

[6] A.D. Bovopolous. Performance evaluation of a traffic control mechanism for ATM networks. In *IEEE Infocom*, pages 469–478, 1992.

[7] M. Butto', E. Cavallero, and A. Tonietti. Effectiveness of the leaky bucket policing mechanism in ATM networks. *IEEE Journal on Selected Areas in Communications*, 9(3):335–342, April 1991.

[8] L. Dittmann and S.B. Jacobsen. Statistical multiplexing of identical bursty sources in an ATM network. In *IEEE Globecom*, pages 1293–1297, 1988.

[9] A.E. Eckberg. Generalized peakedness of teletraffic processes. In *Proc. 10th International Teletraffic Conference*, Montreal, 1983.

[10] 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. In *Int. Teletraffic Cong. Specialist Sem.*, Adelaide, Australia, September 1989.

[11] A.E. Eckberg, D.T. Luan, and D.M. Lucantoni. Meeting the challenge: congestion and flow control strategies for broadband information transport. In *IEEE Globecom*, pages 1769–1773, 1989.

[12] A.I. Elwalid and D. Mitra. Analysis and design of rate-based congestion control of high speed networks, I: Stochastic fluid models, access regulation. *Queueing Systems*, 9:29–64, 1991.

[13] R. Guérin, 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, Sept. 1991.

[14] M. Hirano and N. Watanabe. Characteristics of a cell multiplexer for bursty ATM traffic. In *Proc. of the ICC*, pages 399–403, 1989.

[15] D.S. Holtsinger. *Performance analysis of leaky bucket policing mechanisms for high-speed networks.* PhD thesis, Department of Electrical and Computer Engineering, North Carolina State University, 1992.

[16] B. Laguë, C. Rosenberg, and F. Guillemin. A generalization of some policing mechanisms. In *IEEE Infocom*, pages 767–775, 1992.

[17] J.A.S. Monteiro, M. Gerla, and L. Fratta. Input rate control for ATM networks. In *Proc. 13th International Teletraffic Conference*, pages 117–122, 1991.

[18] M. Murata, Y. Ohba, and H. Miyahara. Analysis of flow enforcement algorithm for bursty traffic in ATM networks. In *IEEE Infocom*, pages 2453–2462, 1992.

[19] E.P. Rathgeb. Modeling and performance comparison of policing mechanisms for ATM networks. *IEEE Journal on Selected Areas in Communications*, 9(3):325–334, April 1991.

[20] M. Sidi, W. Liu, I. Cidon, and I. Gopal. Congestion control through input rate regulation. In *IEEE Globecom*, pages 49.2.1–49.2.5, 1989.

[21] K. Sohraby and M. Sidi. On the performance of bursty and correlated sources subject to leaky bucket rate-based access control schemes. In *IEEE Infocom*, pages 426–434, 1991.

[22] 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, Sept. 1986.

[23] G.M. Woodruff and Rungroj Kositpaiboon. Multimedia traffic management principles for guaranteed ATM network performance. *IEEE Journal on Selected Areas in Communications*, 8(3):437–446, April 1990.

[24] CCITT Study Group XVIII. Draft recommendation I.311, B-ISDN general network aspects. May 1990.

[9] Y. Ohba, M. Murata, and H. Miyahara, "Analysis of interdeparture process for bursty traffic in ATM networks," *IEEE J. Select. Areas Commun.*, no. 3, pp. 468–476, Apr. 1991.

[10] O. Hashida, Y. Takahashi, and S. Shimogawa, "Switched Batch Bernoulli Process(SBBP) and the discrete-time SBBP/G/1 queue with application to statistical multiplexer performance," *IEEE J. Select. Areas Commun.*, no. 3, Apr. 1991.

[11] T. D. Morris and H. G. Perros, "Performance analysis of a multi-buffered Banyan ATM switch under bursty traffic," in *INFOCOM '92*, pp. 436–445, Florence, Italy, May 1992.