Queueing Systems 20 (1995) 117–138

A comparison of the sliding window and the leaky bucket

Arthur W. Berger
AT&T Bell Laboratories, Holmdel, NJ 07733-3030, USA

Ward Whitt
AT&T Bell Laboratories, Murray Hill, NJ 07974-0636, USA

Received 15 January 1994; revised 13 May 1994

In this paper we compare the sliding-window (SW) and leaky-bucket (LB) input regulators. These regulators reject, or treat as lower priority, certain arrivals to a queueing system, so as to reduce congestion in the queueing system. Such regulators are currently of interest for access control in emerging high-speed communication networks. The SW admits no more than a specified number W of arrivals in any interval of specified length L. The LB is a counter that increases by one up to a maximum capacity C for each arrival and decreases continuously at a given drain rate to as low as zero; an arrival is admitted if the counter is less than or equal to $C - 1$. To indirectly represent the impact of the regulator on the performance of the queueing system, we focus on the maximum bursts admissible at the peak rate. We show that the SW admits larger bursts than the LB at any given peak rate and admissible average rate. To make the comparison, we use a special construction: We start with a sample path of an arrival process with a given peak rate. We choose a window length L for the SW and find the minimum window content W that is just conforming (so there are no rejections). We then set the LB drain rate equal to W/L, so that the two admissible average rates are identical. Finally, we choose the LB capacity C so that the given arrival process is also just conforming for the LB. With this construction, we show that the SW will admit larger bursts at the peak rate than the LB. We also develop approximations for these maximum burst sizes and their ratio over long time intervals based on extreme-value asymptotics. We use simulations to confirm that these approximations do indeed enable us to predict the burst ratios for typical stochastic arrival processes.

Keywords: Communication networks, ATM, broadband ISDN, leaky-bucket algorithm, sliding-window algorithm, partial characterization of point processes, extreme values, traffic descriptors, queues.

1. Introduction

An area of active research where queueing theory is playing an important role is the design of new communication networks. The general goal is to be able to support a wide variety of services, exploiting high bandwidth; see Roberts [14]. Our paper is concerned with developing methods for describing and regulating the traffic in these emerging *broadband integrated services digital networks*

(B-ISDNs), assuming that they will exploit the *asynchronous transfer mode* (ATM) technology, which is (partly) characterized by small fixed packets called *cells*.

Work on traffic description for B-ISDNs can roughly be divided into two categories: (1) *stochastic models* and (2) *deterministic algorithms*; in this paper we focus on deterministic algorithms. Stochastic models are primarily being developed to describe important classes of traffic or service types. A good example is the multi-state Markov chain model for video teleconference service traffic developed by Heyman, Tabatabai and Lakshman [8]. These stochastic models have many uses, including guiding the design of equipment and the operation of networks.

However, stochastic models are awkward for *standardized traffic descriptors* to be employed by sources and network operators to control the flow on a connection in real time. It is envisioned that users will make *traffic contracts* with network service providers. Such a contract may include a description of the traffic to go over the connection. Given that the user has chosen to provide a description of the traffic (or that a particular service requires a description of the traffic), the "typically" foreseen scenario is that the user equipment will *shape* the traffic to be conforming to the traffic descriptor, and network equipment will *police* the traffic to confirm that it is indeed conforming.

The problem, then, is to develop (standardized) traffic descriptors. In order that at all times during a connection both the user and the network operator can determine whether the flow is conforming to the contract, it is convenient if the traffic descriptor is specified by a deterministic algorithm. We call a traffic descriptor defined in terms of a deterministic algorithm an *operational traffic descriptor* (OTD).

In this paper, we compare the performance of two OTD's: the *sliding window* (SW) and the *leaky bucket* (LB). The SW admits no more than a specified number W of arrivals in any interval of specified length L. (We consider the intervals to be half open, i.e., $[t, t+L)$.) The LB is a counter that increases by one up to a maximum capacity C for each arrival and decreases continuously at a given drain rate D to as low as zero; an arrival is admitted if the counter is less than or equal to $C-1$ (so that after the arrival it will be less than or equal to C).

We would like to determine which OTD most effectively reduces congestion in a following queueing system. In order to make a fair comparison, we stipulate that the SW and the LB should have the same peak rate and admissible average rate. We define these two rates operationally as well. We define *peak rate* as the reciprocal of the minimum distance between successive arrivals. We define the *admissible average rate* as the reciprocal of the smallest distance between successive arrivals in a hypothetical deterministic equally spaced arrival process that is just conforming with the algorithm. We make a distinction between the admissible average rate and the actual long-run average rate of the source. Note that the admissible average rate is an upper bound on the long-run average rate of sources that conform to the OTD. For the SW, the admissible average rate is W/L; for the LB the admissible average rate is D.

It remains to specify how the OTD reduces congestion in a following queueing system. Instead of focusing on this directly, *we consider the largest burst at the peak rate that is conforming with the OTD*. We are thus relying on the intuitive idea that congestion depends on burstiness, and that a regulator will be less effective in controlling congestion if it admits larger bursts at the peak rate.

We show that the SW can admit larger bursts at the peak rate than the LB. Thus the SW is inferior to the LB with respect to this performance measure. Thus, we contend that the LB should be more attractive to the network provider. On the other hand, users would typically prefer the additional flexibility of larger admissible bursts, everything else being equal. However, with larger bursts, the network may need to allocate more resources per connection, and thus need to have a higher price per connection; i.e., everything else may *not* be equal. (We remark that some of the results in this paper contributed to decisions by an industry forum to adopt the LB for specifying the admissible average rate, there referred to as the "sustainable" rate.)

Our performance measure, the largest burst at the peak rate that is conforming to the OTD, is frequently used to construct the stressful ON/OFF traffic pattern of a source that bursts at the peak rate for as long as allowed and then is idle just long enough for the overall rate to equal the average rate. These hypothetical sources are sometimes called the *worst case traffic* that is conforming to the OTD, where worst is in the sense of arrivals loss. However, for the case of multiple sources and a finite buffer, Doshi [6] has shown circumstances where the ON/OFF traffic is in fact not worst, but rather that higher cell losses can occur when the source, after the ON period, continues at the average rate for a time period before becoming idle. Also, Lee [12] has shown that this latter traffic pattern can lead to higher mean delay in an infinite buffer. Since the worst case non-ON/OFF sources of Doshi and Lee still make use of the largest burst at the peak rate that is conforming to the OTD, our performance measure is still of major interest from that perspective.

Our work here continues work by Reibman and Berger [13]. They applied the SW and LB algorithms to actual data from variable bit rate (VBR) video teleconferencing sequences. They chose SW and LB parameters with specified common admissible average rate so that the given teleconferencing data is just conforming. Under these conditions, they found that the largest burst size that could be admitted is several times larger for SW than for LB. Now we want to see what happens more generally.

To compare SW and LB, we consider a (finite segment of a) sample path of an arrival process with peak rate p and perform a special construction: For any window length L, we let W be the minimum value for which SW is just conforming (all arrivals are admitted). This makes W/L the admissible average rate of the SW. We then let the drain rate of the LB be $D = W/L$, so that LB has the same admissible average rate as SW. We then choose the LB capacity C to be the minimum value that is conforming for the same arrival process sample path. Finally, we determine the maximum burst sizes at the peak rate that are conforming with the SW and the LB, with the determined parameters. We call these maximum bursts B_{SW} and B_{LB}.

It is not difficult to see that the maximum burst sizes are

$$B_{SW} = \begin{cases} W & \text{if } p > W/L, \\ \infty & \text{if } P \leq W/L, \end{cases} \tag{1.1}$$

and

$$B_{LB} = \begin{cases} \lfloor (C-1)/(1-D/p) \rfloor + 1 & \text{if } p > D, \\ \infty & \text{if } p \leq D, \end{cases} \tag{1.2}$$

where $\lfloor x \rfloor$ is the greatest integer less than or equal to x. (To determine (1.2), let X_i be the content of an initially empty, infinite capacity LB just after the ith arrival of the burst. Then $X_1 = 1$, and for $p > D, X_i = X_{i-1} - D/p + 1$, for $i \geq 2$, and thus,

$$X_i = (i-1)(1-D/p) + 1.$$

Noting that B_{LB} is the maximum i such that $X_i \leq C$ yields (1.2).)

Here is a summary of our results and an indication of how the rest of this paper is organized. In section 2 we establish the general inequality

$$2 \leq B_{LB} \leq B_{SW} \tag{1.3}$$

under minor regularity conditions, which implies that the *burst ratio* B_{SW}/B_{LB} *is* always greater than or equal to 1. In section 3 we develop approximations for $W = B_{SW}$, D, C and B_{LB} as a function of the time interval T (when T is large) and arrival process characteristics for quite general stochastic arrival processes based on extreme-value asymptotics. It is interesting that for the SW the relevant extreme-value behavior is for the *arrival counting process*, while for the LB the relevant extreme value behavior is for the workload in a $G/D/1$ *queueing model*, because the LB can be modelled as a $G/D/1$ queue.

This analysis leads us to propose a simple rough approximation to predict the burst ratio, namely,

$$B_{SW}/B_{LB} \approx \sqrt{8L/c^2 \log T}, \tag{1.4}$$

where L is the window length, T is the time interval and c^2 is the *asymptotic variance* or limiting value of the *index of dispersion for counts* (IDC), defined in (3.2) below. For a renewal arrival process, c^2 is the *squared coefficient of variation* (SCV, variance divided by the square of the mean) of an interarrival time. Approximation (1.4) is not extraordinarily accurate, but it gives a good idea how the burst ratio depends on the arrival process and the variables L and T. We also develop more accurate approximations, which are more complicated. The extreme-value asymptotics supporting these approximations are of independent interest, and are being presented separately in Berger and Whitt [3, 4].

In section 4 we provide additional insight by showing what the burst ratio is for stylized deterministic ON/OFF arrival processes. In section 5 we describe simulation experiments conducted to see what the burst ratio is for typical stochastic arrival processes and to investigate the accuracy of the approximations proposed in section 3.

2. Bounds on conforming burst sizes

We restrict the set of sample paths of arrival times to those where a peak rate and admissible average rate are well defined. In particular, we make the following two *assumptions*:

(A1) The minimum interarrival time, p^{-1}, exists and is positive.

(A2) There exists a choice for L such that $p > W/L$.

Note that assumption (A2) eliminates the uninteresting case of *all* arrivals having equal spacing, i.e., continuous bit rate (CBR) connections. (These assumptions are satisfied in the communication network application of interest.) As the typical case, we have in mind

$$p^{-1} \ll L \ll T, \tag{2.1}$$

i.e., p^{-1} is much smaller than L, which in turn is much smaller than the duration of the connection, which we take to be our time interval T. However, this is not assumed in the following.

For a sample path of arrival times that satisfies (A1), let N denote the maximum number of consecutive interarrival times that are equal to the minimum interarrival time. Note that the number of arrivals in this "burst" is $N+1$. The following lemma provides a *necessary condition* for (A2). It says that in order for the admissible average rate to be less than the peak rate, the window length needs to be chosen greater than the longest burst of arrivals at the peak rate.

LEMMA 1

If $W/L < p$, then $L > (N+1)p^{-1}$.

Proof

Suppose the contrary: $L \leq (N+1)p^{-1}$. Then there exists an integer $k \leq N$ such that $kp^{-1} < L \leq (k+1)p^{-1}$. From the definition of N, the sample path contains at least one portion with N consecutive interarrival times each with length p^{-1}. Thus, as a window with the above length L passes over the sample path, the window will enclose $k+1$ arrivals. Thus, $W = k+1$, and $W/L = (k+1)/L \geq (k+1)/[(k+1)p^{-1}] = p$. But then the admissible average rate, W/L, is greater than or equal to p, which is a contradiction. Thus, $L > (n+1)p^{-1}$. □

Note that *the converse of Lemma 1 is false*. That is, L can be greater than $(N+1)p^{-1}$ and W/L is still greater than p. The idea is that the definition of N says nothing about a burst of arrivals with spacing greater than p^{-1}, and there can be a burst of equally spaced arrivals but with a spacing that is just a bit bigger than p^{-1}. As a simple example, suppose $p = 1$, and $N = 1$ and $L = 2.5$. Suppose the sample path includes three consecutive arrivals whose spacing is 1.1. A window of length 2.5 would enclose these three arrivals, and thus W is, at least, 3. Thus, W/L is at least 1.2, which is greater than p.

The previous paragraph and Lemma 1 reveal an awkwardness of the SW algorithm that is not present with the LB algorithm. With the LB we can a priori choose the drain rate to be less than the peak rate. However, with SW we need to guess a value of L that is big enough in order that the resulting admissible average rate, W/L, is indeed less than the peak rate.

We now establish our main comparison result.

THEOREM 1

If a sample path satisfies assumptions (A1) and (A2) and the SW and LB are just conforming with $p > W/L = D$, then (1.3) holds or, equivalently,

$$1 \leq \frac{B_{SW}}{B_{LB}} \leq \frac{W}{2}. \tag{2.2}$$

Proof

Consider any sample path that satisfies assumptions (A1) and (A2). Then p is determined and we can choose L such that $p > W/L$, and we do. Given p, L, and W, the SW maximum allowable burst B_{SW} is simply W by (1.1). As for B_{LB}, consider the set of $\mathcal{S}$ of all sample paths that satisfy (A1) and (A2) and that are just conforming with the above values of p, L, and W. We imagine that the sample paths in $\mathcal{S}$ are used to increment an (infinite capacity) leaky bucket with drain rate $D = W/L$. We are interested in those sequences in $\mathcal{S}$ that yield the smallest and largest values for C. By (1.2), sample paths that yield the smallest value for C yield the smallest value for B_{LB}, and likewise for the largest value for C.

Consider first the smallest possible value for C. We know that there are at least two arrivals a distance p^{-1} apart. Suppose the bucket were empty prior to the first of these arrivals. Thus, C must be at least $2 - D/p$ to be big enough to accommodate these two arrivals. Moreover, it is possible that C be no bigger. For example, consider a sample path where the first interarrival time is p^{-1}, followed by an interarrival time at least long enough for the bucket to empty (at least

$(2 - D/p)D)$ and then all subsequent interarrival times are of length L/W. Thus, the smallest possible value for C is $2 - D/p$. With C equal to this value, B_{LB} is obviously equal to 2. (Substituting $2 - D/p$ for C in (1.2) yields $B_{LB} = 2$.) Thus, the minimum value for B_{LB} is 2.

To determine the largest value of C for the arrival processes in $\mathscr{S}$, consider a busy period of the LB that begins at some time t. To allow as much flexibility as possible, suppose that the idle period before this busy period is long enough so that the possible arrival times during the busy period are not constrained by the SW algorithm for windows that begin before time t. The key observation is that the LB will accumulate the greatest content during the busy period if as many arrivals come as possible and if they come as quickly as possible, i.e., if W arrivals occur with a spacing of p^{-1}. After the Wth arrival the content of the LB will be

$$X_W = (W - 1)(1 - D/p) + 1.$$

Thus, for the arrival processes in $\mathscr{S}$, the maximum possible content in the LB is $(W - 1)(1 - D/p) + 1$. Substituting $(W - 1)(1 - D/p) + 1$ for C in (1.2) shows that the maximum burst at the peak rate that is conforming to this LB is W. Thus, for the arrival processes in $\mathscr{S}$, the largest possible value for B_{LB} is W, which equals B_{SW}. Taking the reciprocal of the terms in (1.3), multiplying by B_{SW}, and using (1.2) yields (2.2). □

The inequalities in (2.2) highlight the fact that the burst ratio B_{SW}/B_{LB} can be as small as one but can be very large, since W increases with L. For the example video teleconferencing sequences in Reibman and Berger [13], typical values for W were in the range of 100 to 1,000.

3. Approximations from extreme-value asymptotics

In this section we develop approximations based on extreme-value asymptotics for the burst ratio B_{SW}/B_{LB}. From (1.1), we have that $B_{SW} = W$, provided p is not too small, which we assume is the case. From (1.2) we see that $B_{LB} \approx C$ when $D \ll p$, and we assume that this is the case as well. Hence, we want to estimate W and C for a given time horizon T, window length L and given stochastic arrival process $A \equiv \{A(t) : t \geq 0\}$ with arrival rate λ and peak rate p. (The random variable $A(t)$ counts the number of arrivals in the interval $[0, t]$.) Without loss of generality, we choose the measuring units so that the arrival rate is $\lambda = 1$. We first estimate W and then estimate C assuming a drain rate of $D = W/L$.

We develop approximations for W by exploiting extreme-value asymptotics for the specified arrival process A. In general, the extreme-value asymptotics depends on the detailed properties of the arrival process A; see Leadbetter et al. [11]. However, if L is suitably large, then we might hope that the window content for any fixed interval is approximately normally distributed, so that it might suffice to consider extreme values for normally distributed random variables depending only on appropriate mean and variance parameters. Of course, we must consider L in relation to T. To obtain useful new extreme-value results, we consider the asymptotic behavior as T and L get large, with T much larger than L, but L much larger than $\log T$. Motivated by the present problem, in Berger and Whitt [3] we show that in the regime $T \gg L \gg \log T$ we can use extreme value theory for Brownian motion, which means that the arrival process can be characterized by only a single parameter beyond its rate. In particular, in [3] we justify the approximation

$$W \approx EW \approx L + \sqrt{c^2 L}\Big(\sqrt{2\log(T/L) + 2\log\log T} - 1.28/\sqrt{2\log(T/L)}\Big), \quad (3.1)$$

where the arrival rate equals one and where c^2 is the *asymptotic variance* (limiting value of the index of dispersion for counts), i.e.,

$$c^2 = \lim_{t\to\infty} \frac{\operatorname{Var} A(t)}{EA(t)}; \quad (3.2)$$

see (2) and Theorem 1 of [3]. For a renewal arrival process, c^2 is just the *squared coefficient of variation* (SCV, variance divided by the square of the mean) of an interarrival time. Approximation (3.1) is appealing because it is relatively simple. Moreover, having $T \gg L \gg \log T$ seems realistic for communication network applications.

As a simplification of (3.1), we also propose the following rough approximation:

$$W \approx EW \approx L + \sqrt{2c^2 L \log(T/L)}. \quad (3.3)$$

Approximation (3.3) is obtained from (3.1) by dropping asymptotically negligible terms. (For greater accuracy we suggest (3.1), but we will show that (3.3) is quite close to (3.1) in our numerical examples in section 5.)

Before proceeding, note that (3.3) states that W is approximately equal to a single mean window content L plus a standard deviation $\sqrt{c^2 L}$ (obtained from a central limit theorem argument) times a multiplicative factor $\sqrt{2\log(T/L)}$, which reflects the growth due to the maximization. The simplification in (3.3) applies only to this multiplicative factor.

Given the approximation for W, we obtain the associated approximation for the LB drain rate $D = W/L$. The LB capacity C is then the maximum buffer content in $[0, T]$ with the given arrival process A and drain rate D. If we use (3.3), then we obtain the approximation

$$D = W/L \approx 1 + \xi, \quad (3.4)$$

where

$$\xi = \sqrt{2c^2 \log(T/L)/L}. \tag{3.5}$$

We use (3.5) to develop simple approximations (3.17) and (3.18) below. We used D obtained from (3.1) for more accurate approximations.

If we change the measuring units for time so that the drain rate is 1 and the arrival rate is $\rho \equiv D^{-1}$ (and keeping unit jumps), then the content of the leaky bucket is equivalent to the workload process in a $G/D/1$ queue with arrival rate ρ and service time 1. Again motivated by this problem, in Berger and Whitt [4] we develop approximations for the maximum values of queueing processes over long time intervals in general single-server queues. In [4] the service rate is 1, so the results there apply after making the change of the time scale by ρ^{-1}. In particular, in [4] we develop the approximation (for our time scale here)

$$C \approx EC \approx \gamma(\log(\rho^{-1}T) + \log\beta + 0.577), \tag{3.6}$$

where γ and β are parameters that can be approximated given the arrival process with arrival rate ρ and service time 1.

The approximation (3.6) is natural in view of extreme-value limit theorems; see Leadbetter et al. [11] and Iglehart [9]. In [4] we extend the class of systems previously considered and develop approximations for the parameters γ and β. In general, γ and β depend on the arrival process beyond c^2 though, so we will consider refinements. Approximation (3.6) is not limited to deterministic service times, but here we will only consider this special case.

The critical parameter γ in (3.6) is the reciprocal of the asymptotic decay rate of the steady-state workload, say Z; i.e., $\gamma \equiv \eta^{-1}$ where

$$P(Z > x) \sim \alpha e^{-\eta x} \quad \text{as} \quad x \to \infty, \tag{3.7}$$

with $f(x) \sim g(x)$ as $x \to \infty$ meaning that $f(x)/g(x) \to 1$ as $x \to \infty$.

The asymptotic decay rate η in (3.7) can be calculated exactly for many arrival processes; see Glynn and Whitt [7], Whitt [15] and references cited there. For example, for renewal purposes, η is the unique root of the equation $Ee^{\theta(V-U)} = 1$, where V is a service time and U is an interarrival time. We will use this formula in our numerical examples in Section 5.

Methods for approximating the asymptotic decay rate η in (3.7) in quite general single-server queues have also been developed by Abate et al. [1] and Choudhury and Whitt [5]. The simple heavy-traffic approximation in the case of a general arrival process is

$$\eta \approx \eta_{HT} = 2(1-\rho)/c^2, \tag{3.8}$$

from which we obtain

$$\gamma \approx \gamma_{HT} = c^2/2(1-\rho). \tag{3.9}$$

The HT approximation in (3.9) is appealing because it only depends on the available parameters c^2 and ρ. However, it is known to be not very accurate. We can do better if we know more about the arrival process A. A refinement for the special case of renewal arrival processes with known first three moments is

$$\eta \approx \eta_{HT}(1-\eta^*(1-\rho)), \tag{3.10}$$

where

$$\eta^* = \frac{2-2m_3+3c^2(c^2+2)}{3(c^2)^2}, \tag{3.11}$$

with m_3 being the third moment of the rate-1 arrival process; see Theorem 3 of [1] and (6.14) of [4].

The parameter β in (3.6) can be approximated by $\beta = \alpha\theta$, where α is the asymptotic constant in (3.7) and

$$\theta \approx 2(1-\rho)^2/c^2; \tag{3.12}$$

see (1.10) and (5.9) of [4]. Other approximations for θ are also given in [4], but we will not consider them here. The exact formula for θ for renewal arrival processes is given in Iglehart [9], but it is somewhat complicated. The asymptotic parameters α and η in (3.7) can often be computed for fully specified models. A simple approximation for the workload asymptotic constant from [1] and Theorem 2 of [2] is

$$\alpha \approx (\eta E(W))\frac{\rho(e^\eta-1)}{\eta} = E(W)\rho(e^\eta-1), \tag{3.13}$$

where W is the steady-state waiting time (seen by arrivals and thus not the workload) in the $G/D/1$ model.

For renewal arrival processes we can use the Kraemer and Langenbach-Belz [10] approximation for EW, specialized to D service, to obtain the approximation

$$EW \approx \frac{\rho c^2}{2(1-\rho)}h(\rho,c^2), \tag{3.14}$$

where

$$h(\rho,c^2) = \begin{cases} \exp(-[2(1-\rho)/3\rho](1-c^2)^2/c^2), & c^2 \le 1, \\ \exp(-(1-\rho)(c^2-1)/c^2), & c^2 \ge 1. \end{cases} \tag{3.15}$$

In summary, a *full approximation* for C in the case of a rate-1 renewal arrival process partially characterized by the second and third moments of the interarrival time is (3.6) with $\gamma = \eta^{-1}$ for η in (3.10) and (3.11), and with $\beta = \alpha\theta$ for θ (3.12) and α in (3.13), (3.14) and (3.15). A simple rough *heavy-traffic* (HT) approximation based on (3.6) and (3.9) only is

$$C \approx c^2 \log(T)/2(1-\rho). \tag{3.16}$$

We now consider simple approximations based on (3.4). Combining (3.4) and (3.16), and replacing $1+\xi$ by 1 and $\log(T/L)$ by $\log T$, we obtain another *simple approximation* for C, namely,

$$C \approx \frac{(1+\xi)}{2\xi} c^2 \log(T) \approx \sqrt{\frac{c^2 L \log T}{8}}. \tag{3.17}$$

Finally, we can combine approximations for W and C to obtain an approximation for the burst ratio B_{SW}/B_{LB}. In particular, using (3.4) and the middle term in (3.17), we obtain

$$\frac{B_{SW}}{B_{LB}} \approx \frac{W}{C} \approx \frac{(1+\xi)L}{C} \approx \frac{2\xi L}{c^2 \log T} \approx \sqrt{\frac{8L\log(T/L)}{c^2}}\,\frac{1}{\log T}. \tag{3.18}$$

Approximating $\log(T/L)$ by $\log T$, we obtain (1.4) from (3.18).

The relatively simple approximations for W, C and B_{SW}/B_{LB} in (3.3), (3.17) and (1.4) are obviously appealing because they clearly reveal the dependence upon the basic model parameters T, L and c^2. We evaluate these approximations by making comparisons with simulations in section 5.

We conclude this section by pointing out that the approximations for W in (3.1) and (3.3) and the approximations for C in (3.6), (3.16) and (3.17) are really for the expected values. For single sample paths there will be fluctuations about these expected values.

First, from (4) of [3], we see that the standard deviation of W is approximately

$$STD(W) \approx 1.28\sqrt{c^2 L/2\log(T/L)}. \tag{3.19}$$

To see the impact of (3.19) on the drain rate D and the G/D/1 traffic intensity ρ, note that, by (3.3) and (3.19),

$$\frac{STD(W)}{E(W-L)} \approx \frac{0.64}{\log(T/L)}. \tag{3.20}$$

For $T/L = 1000$, the ratio in (3.20) is 0.093.

Second, given the drain rate $D = \rho^{-1}$, from (1.6) of [4], we see that the standard deviation of the LB capacity C is approximately

$$STD(C) \approx 1.28\gamma. \tag{3.21}$$

To see the impact on our estimate of C, note that, by (3.6) and (3.21),

$$\frac{STD(C)}{EC} \approx \frac{1.28}{\log(\rho^{-1}T) + \log\beta + 0.577}. \tag{3.22}$$

For the examples in section 5 with $T = 10^6$, the ratio in (3.22) is also approximately 0.10. Since these ratios are only about 10%, we conclude that the main effect is in the means.

4. Burst ratios for stylized deterministic arrival processes

To provide additional insight, we now consider a stylized deterministic sample path of arrivals, for which we can easily determine the bursts allowed by the OTD's. Suppose the source is ON/OFF with constant length ON and OFF periods. In particular, the source repeats the following, constant length, cycle: it emits a fixed number of arrivals, say B, at the peak rate, p, and then the source is idle for a fixed period that is longer than p^{-1} (so that the average rate of source is less than the peak rate).

We consider in turn three cases for the length of the window, L, with the third case being of greatest interest:

1. L is less than or equal to the length of the ON period, the latter being $(B-1)/p$.
2. L is greater than the length of the ON period and less than or equal to the length of one ON + OFF cycle.
3. L is greater than the length of one ON + OFF cycle.

The first case is degenerate since the SW and LB have an admissible average rate that is equal to (or even a bit greater than) the peak rate, and thus the maximum conforming burst at the peak rate is infinite for both the SW and the LB. In particular, suppose for some integer n, $(n-1)p^{-1} < L \leq np^{-1}$. Then the maximum count in the window is n, and the admissible average rate is n/L. Thus, if L happens to exactly equal np^{-1}, then the admissible average rate is the peak rate, otherwise the admissible average rate is actually a bit bigger than the peak rate. In this degenerate case, the SW and LB impose no further constraint on the class of conforming processes than that imposed by the peak rate alone.

In the second case, W equals the number of arrivals in an ON period, and thus $B_{SW} = B$. Note that the admissible average rate, W/L, is greater than or equal to the actual long-run average rate. As for the leaky bucket, since the drain rate equals W/L and thus is greater than or equal to the long-run average rate, the leaky bucket can work off the cells of one ON period before the next ON period begins. Thus, for the given sample path, the capacity of the leaky bucket is determined by a burst of B arrivals at the peak rate. Thus, the leaky bucket, by construction, will allow a maximum, conforming burst of B arrivals at the peak rate. (That is, with the drain rate set to the admissible average rate, the capacity, C, is given by $(B-1)(1-D/p)+1$, which when substituted into $\lfloor (C-1)/(1-D/p) \rfloor + 1$ yields B again.) Thus in the second case, $B_{SW} = B_{LB} = B$ and the ratio B_{SW}/B_{LB} is one.

In the third case, the sliding window will enclose arrivals from more than one burst, and W will be greater than B. Now here is the key point: Although for this sample path the burst size is always B, when one asks what is the maximum conforming burst at the peak rate, the SW algorithm would allow all of the arrivals in the window to be clumped as one burst. In contrast, the leaky bucket would not allow arrivals beyond one ON period, because the drain rate, which is equal to W/L, will again be greater than or equal to the long-run average rate and the reasoning of the second case pertains. (D will equal the long-run average rate when L equals an integral multiple of the length of the ON + OFF cycle, otherwise D is greater than the long-run average rate.) Hence, for the leaky bucket, the maximum conforming burst size is still B. In summary, for L greater than the length of one ON + OFF cycle,

$$B_{SW} = W > B = B_{LB}. \tag{4.1}$$

If we think of L as being "large" such that $W = nB$ for some integer n, then

$$B_{SW}/B_{LB} = W/B = n, \tag{4.2}$$

and thus the ratio can be arbitrarily large, given that n can be arbitrarily large. If we consider alternative stylized ON/OFF sample paths that differ in the value of the burst size, B, then the ratio B_{SW}/B_{LB} is greatest when B is its minimum value of 2. Thus, we attain the upper bound $W/2$ in (2.2).

5. Simulations of stochastic arrival processes

Theorem 1 implies that the burst ratio B_{SW}/B_{LB} must be in the interval $[1, W/2]$. The stylized deterministic example in section 4 shows that the burst ratio can be anywhere in this interval, depending on the arrival process. It also gives an idea about when the ratio will be near 1 and when it will be substantially greater than 1. Now we are interested in stochastic arrival processes. The approximations in section 3 yield predictions of the burst ratio, but it remains to determine how accurate these approximations are. Hence, we simulated some stochastic arrival processes to get further insight.

In our simulations we consider sample paths of 1,000,000 arrivals (which, given the 48 byte information field of an ATM cell and ignoring the overhead of higher layers, corresponds to 48 Mbytes of information being sent during the course of an ATM connection), and we consider a normalized arrival rate of 1 arrival per time unit.

As a first example, we consider the arrival process to be Poisson, and we obtained the results shown in figs. 1(a) through 1(g) for a simulated sample path. For the given sample path, we consider a range of values for L from 10 to 1,000.

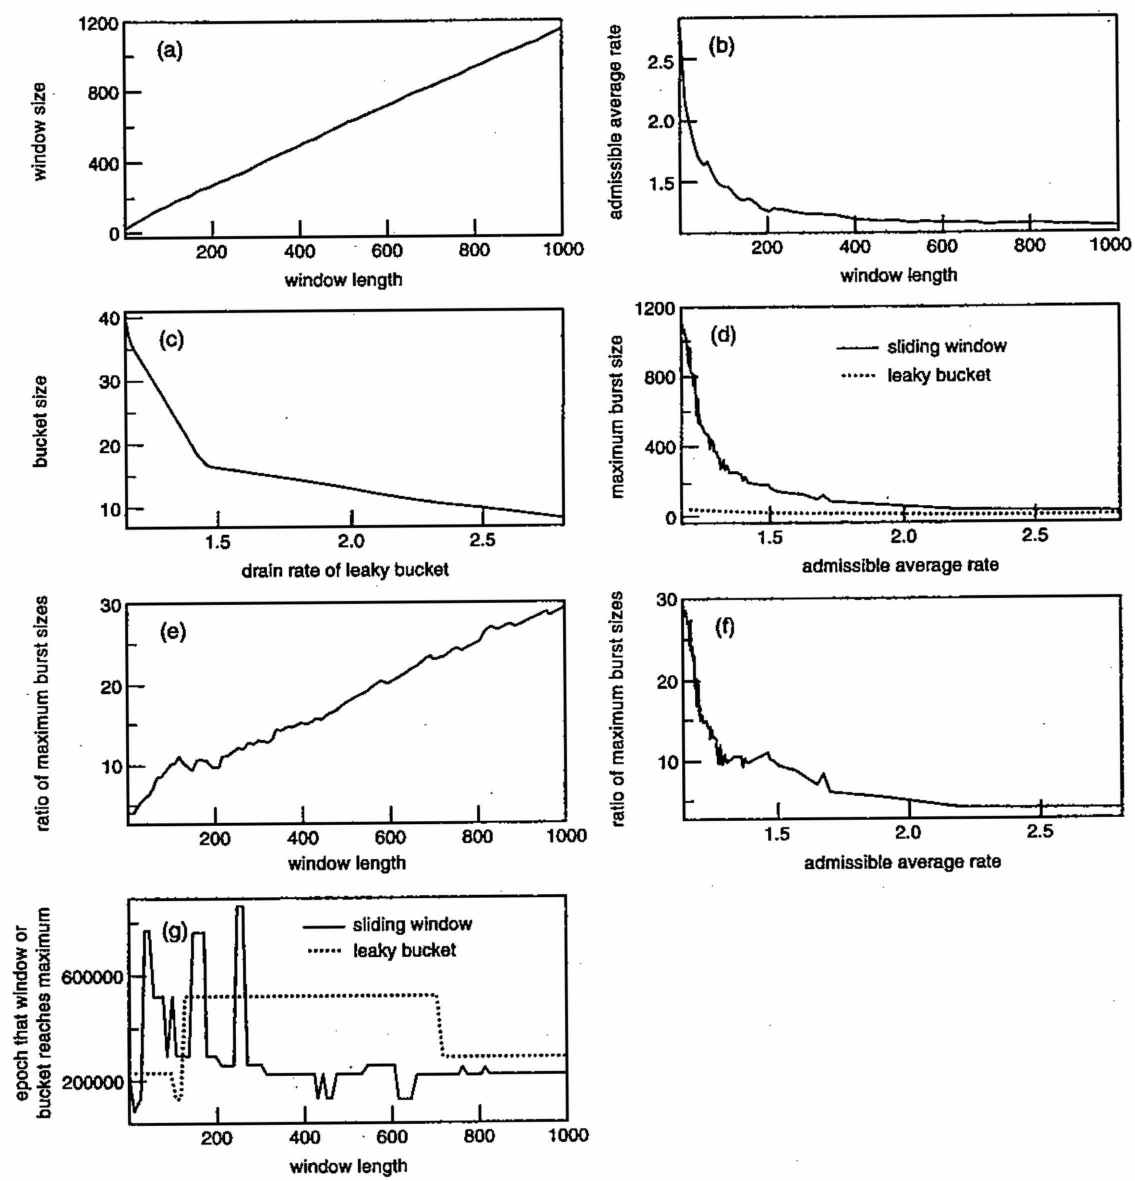

Fig. 1. Simulation results for a Poisson arrival process.

For each value of L, we obtain the value for W, and the admissible average rate, W/L, as displayed in figs. 1(a) and 1(b), respectively. Figure 1(b) shows that, overall, the admissible average rate is decreasing in L. However, there can be intervals where an increase in L yields an "atypically" large increase in W, sufficient to cause the admissible average rate to increase, not decrease.

Given L and W, we set D, the drain rate of the LB, equal to W/L so that both algorithms have the same admissible average rate. For each value of D, we obtain the value for C, as shown in fig. 1(c). For the sample path, the minimum inter-epoch time is roughly 3.4E-07, and its reciprocal, 2,900,000, is the peak rate ρ. (The peak rate is "large" because in simulating the Poisson process we did not a priori enforce a minimum inter-epoch time.) Given p, D, and C, then B_{LB} is determined from (1.2), and B_{SW} is W. The maximum burst sizes B_{SW} and B_{LB} versus the admissible average rate are shown in fig. 1(d), and the burst ratio B_{SW}/B_{LB} versus L and versus the admissible average rate is shown in figs. 1(e) and 1(f), respectively. Lastly, fig. 1(g) plots the epoch at which the window or bucket content reaches a maximum for each value of the window length. Figure 1(g) shows that the epochs need not coincide for the two algorithms and that for each algorithm, as L increases, the epoch at which the maximum occurs can shift significantly. *The key observation from fig. 1 is that B_{SW} is significantly bigger than B_{LB}*, roughly an order of magnitude bigger for admissible average rates up to twice the long-run average rate.

We also considered multiple sample paths of the Poisson process. Table 1 reports the results from twenty runs, each of 1,000,000 arrivals and differing in the choice of the initial seed. The table includes the 90% confidence intervals using the t-distribution under the standard assumptions of independence and normality. (The drain rate of the leaky bucket is obtained from table 1 by simply dividing the reported average value of W by L.) The results in table 1 are consistent with those in fig. 1.

For a second example, we consider a renewal arrival process with $D + H_2$ interarrival times, i.e., with an interarrival time of a constant plus a hyperexponentially distributed random variable. The constant is 0.01 and the parameters of the hyperexponential distribution were chosen so that it has balanced means and so that the interarrival times have mean 1 and SCV 4. Balanced means holds when

Table 1
Data from twenty sample paths of a Poisson arrival process with rate 1.

Window length L	Maximum window count W	Leaky bucket capacity C	Burst ratio B_{SW}/B_{LB}
10.	29.3 ± 0.5	7.9 ± 0.2	4.0 ± 0.2
100.	148.2 ± 1.1	17.6 ± 0.7	8.8 ± 0.3
1,000.	$1,126.8 \pm 4.3$	45.8 ± 1.4	25.1 ± 0.9

$\pi\lambda_1^{-1} = (1-\pi)\lambda_2^{-1}$ for the H_2 density $\pi\lambda_1 e^{-\lambda_1 x} + (1-\pi)\lambda_2 e^{-\lambda_2 x}$. The specific parameters are $\pi = 0.1106437$, $\lambda_1 = 0.223523$ and $\lambda_2 = 1.796680$. This H_2 distribution has SCV 4.0812.

For a sample path of 1,000,000 arrivals, the realized minimum inter-epoch time, to four significant figures, was 0.01000 and hence the peak rate was 100.00. Repeating the same construction as with the first example, fig. 2 shows the results from a single sample path and table 2 shows the sample averages from twenty sample paths. Qualitatively, the results are similar to those of the first example,

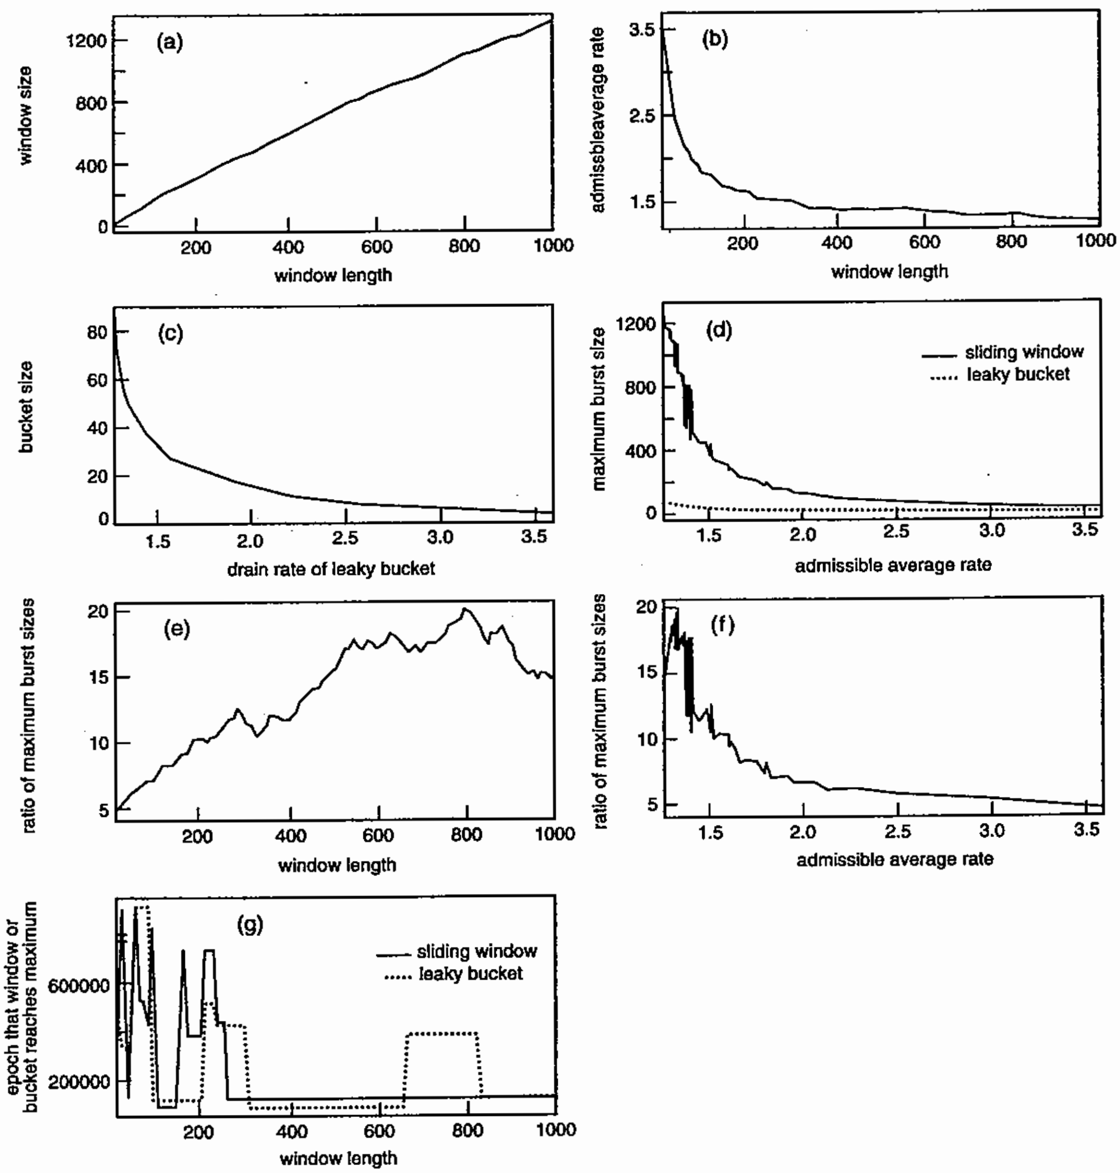

Fig. 2. Simulation results for a (constant + H_2) renewal arrival process.

Table 2
Data from twenty sample paths of a renewal process with a constant plus hyperexponentially distributed inter-epoch time.

Window length L	Maximum window count W	Leaky bucket capacity C	Burst ratio B_{SW}/B_{LB}
10.	37.5 ± 0.6	9.0 ± 0.3	4.3 ± 0.2
100.	185.6 ± 2.4	24.2 ± 0.9	7.8 ± 0.4
1,000.	$1,245.5 \pm 8.4$	85.2 ± 3.4	14.7 ± 0.7

and again B_{SW} is roughly an order of magnitude larger than B_{LB} for admissible average rates up to twice the long-run average rate.

For our last example, we simulated the popular ON/OFF process with a geometrically distributed number of arrivals during an ON time with a mean of 100 arrivals and with a constant spacing of 0.1 during the ON time, and with an OFF time equal to 0.1 plus an exponentially distributed random variable. The parameter of the exponential distribution was chosen so that the overall mean interarrival time is still 1. The resulting exponential distribution has mean 90. Thus, the arrival process is again a renewal process and for these parameters the SCV is 161.2, and the minimum inter-epoch time is 0.1, or equivalently, the peak rate is 10.

Note that from Lemma 1 in section 2, in order for W/L to be less than p, L needs to be chosen longer than the longest ON period of the sample path. The expected length of an ON period is 10, and for the given sample path, we found that L equal to 100 happened to be longer than the longest realized ON period. For $L = 100$, W was 995, or equivalently, admissible average rate was 9.95, which is just a bit less than the peak rate of 10. Using the same construction procedure as in the prior two examples, with L varying from 100 to 5,000, fig. 3 shows the results for a single sample path of the ON/OFF arrival process, and table 3 shows the averages for twenty sample paths. The ratio B_{SW}/B_{LB} is smaller in this example than in the prior two, though it is still greater than or equal to one, of course. The ratio is greater than 2 for admissible average rates up to 3 times the long-run average rate. Though, for admissible average rates above 6, or equivalently for L less

Table 3
Data from twenty sample paths of a renewal, ON/OFF process with a geometrically distributed number of arrivals during an ON period with a mean of 100 arrivals.

Window length L	Maximum window count W	Leaky bucket capacity C	Burst ratio B_{SW}/B_{LB}
200.	1359.5 ± 33.1	337.4 ± 23.9	1.31 ± 0.07
1000.	2978.4 ± 102.7	894.4 ± 45.2	2.37 ± 0.13
5000.	8481.0 ± 220.6	1575.6 ± 95.5	4.55 ± 0.22

than 250, the ratio of burst sizes is relatively small, between 1.1 and 1.3, see figs. 3(f) and 3(e), respectively.

To provide some insight into the relatively small burst ratios in this example, we can heuristically compare the present, stochastic ON/OFF process with the stylized, deterministic ON/OFF process in section 4. Note that in the stochastic example, the burst ratio is close to 1 for small values of L, values on the order of 100, which is the mean length of an ON+OFF cycle. These small values of L roughly correspond to the second case of the deterministic ON/OFF process, where

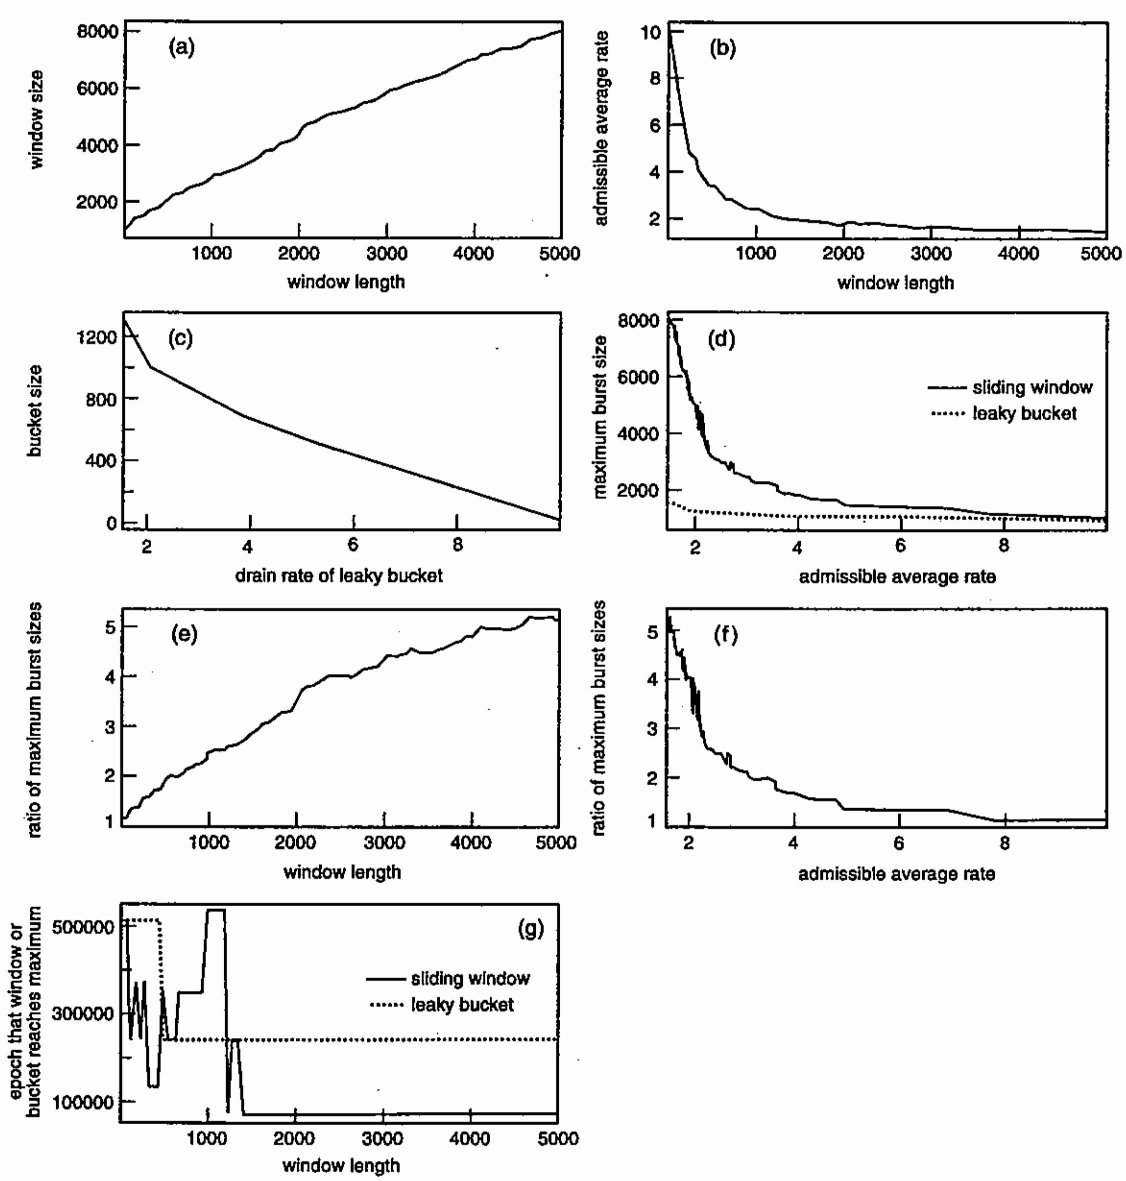

Fig. 3. Simulation results for an ON/OFF arrival process.

L is greater than the length of the ON period but not greater than the length of an ON + OFF cycle. This second case of the stylized example yields the lower bound of 1 for the burst ratio.

We compare the approximations in section 3 to the simulation estimates for these three examples in table 4. Since the approximations for the maximum window

Table 4
A comparison of approximations with simulation estimates for three rate-1 renewal arrival processes over a time interval $T = 10^6$.

Characteristics	Exponential, $c^2 = 1$		$D + H_2, c^2 = 4$		ON/OFF, $c^2 = 161.2$	
L	100	1000	100	1000	1000	5000
W						
simulation est.	148	1127	186	1246	2978	8481
approx. (3.1)	146	1127	191	1254	2615	8221
approx. (3.3)	143	1118	186	1235	2492	7922
$\rho = D^{-1} = L/W$						
simulation est.	0.675	0.887	0.539	0.803	0.336	0.589
approx. using (3.1)	0.686	0.887	0.523	0.797	0.382	0.608
arrival process parameters						
third moment, m_3	6.0		60.51		43,789	
η^* in (3.11)	−0.33		−0.980		−0.111	
LB capacity approx.						
γ_{HT} in (3.9)	1.59	4.43	4.19	9.86	130.5	205.7
refined γ in (3.10)	1.44	4.27	2.85	8.23	122.1	197.1
exact γ	1.47	4.59	1.90	8.29	112.2	204.2
approx. EW in (3.14)	1.09	3.93	1.53	6.75	27.0	84.8
approx. α in (3.13)	0.732	0.849	0.554	0.690	0.092	0.253
θ in (3.12)	0.197	0.0255	0.114	0.0206	0.00473	0.00190
$\log \beta = \log(\alpha\theta)$	−1.94	−3.83	−2.76	−4.26	−7.73	−7.64
LB capacity C						
simulation est.	17.6	45.8	24.2	85.2	894	1576
approx. (3.6) with exact γ	18.9	49.0	23.3	85.9	855	1481
approx. (3.6) with (3.10)	18.5	45.6	35.1	85.1	931	1430
approx. (3.16)	22.0	61.2	57.9	136.2	1803	2842
approx. (3.17)	13.1	41.6	26.3	83.1	528	1180
burst ratio B_{SW}/B_{LB}						
simulation est.	8.8	25.1	7.8	14.7	2.37	4.55
approx. (3.6) and (3.1)	7.7	23.0	8.2	14.6	3.06	5.55
approx. (3.6), (3.1), (1.2)	8.1	23.0	8.3	14.6	2.26	4.64
approx. (1.4)	7.6	24.1	3.8	12.0	1.90	4.24

content W are intended for the regime $T \gg L \gg \log T$, we consider only the two largest values of L in tables 1–3. Thus we consider two window lengths L for each arrival process, making six cases overall.

Given $T = 10^6$ and L, we consider the approximations for W in (3.1) and (3.3). We see that both approximations agree quite closely with the simulation estimates in all cases, although both approximations underestimate the simulation estimates by about 5–15% for the highly bursty ON/OFF example. From table 4 we see that the simple approximation (3.3) performs about as well as the more detailed approximation (3.1). However, we use (3.1) for the remainder of the study, e.g., in specifying the drain rate D.

Given the (approximate) drain rate D, there are several possible approximations for the LB capacity C. Our most detailed approximation is (3.6), for which we consider three different values for γ. First, we consider the exact value obtained from solving the equation

$$Ee^{\theta(V-U)} \equiv e^{\theta} Ee^{-\theta U} = 1, \tag{5.1}$$

where V is a deterministic service time equal to one, and $U = \rho^{-1} Z$ where Z is an interarrival time with mean one, as in the three examples herein for which $Ee^{-\theta\rho^{-1}Z}$ is easily constructed. We also consider the heavy traffic (HT) approximations γ_{HT} in (3.9) and the refined approximation in (3.10). Consistent with the supporting theory, the HT approximations for γ perform better at higher ρ, which corresponds to large L. The refined approximation in (3.10) has under 10% error in all cases but one. For the $D + H_2$ example with $L = 100$, it has 50% error. From the formulas and the numerical results, it is evident that it is desirable to obtain the exact value of γ. Fortunately, it is often possible to obtain it, even for non-renewal processes.

The refined γ in (3.10) depends on η^* in (3.11) and thus the third moment of the interarrival time for each renewal arrival process. The third moment m_3, the parameter η^* in (3.11) and the approximations for γ in (3.9) and (3.10) are all given in table 4.

Our approximation for C in (3.6) employs $\log\beta$ for $\beta = \alpha\theta$ for θ in (3.12) and α in (3.13). The approximation for α in (3.13) depends on the approximation for EW in (3.14) and depends on γ; we use the exact γ from (5.1). All these are displayed in table 4 as well. Due to the logarithm in $\log\beta$, it seems less critical to have high precision in the approximations for α and θ.

In table 4 we compare four approximations for the LB capacity C and three approximations for the burst ratio B_{SW}/B_{LB} to the simulation estimates. We display the approximation for C in (3.6) with the leading term γ being either the exact γ or the refined γ in (3.10). (In both cases, $\log\beta$ is computed with the exact γ.) From table 4 we see that, overall, all of the approximations have a ball-park accuracy. Approximation (3.6) for the LB capacity C with the exact γ is consistently good, with less than 10% error in all cases. The approximations for the burst ratio using

(3.6) all use the exact value of γ. For the third example, in which the peak rate is only 10, the approximation is significantly improved by using (1.2) for B_{LB} instead of $B_{LB} \approx C$. However, the simplest approximations (3.17) for C and (1.4) for B_{SW}/B_{LB} seem remarkably good, given their simplicity.

Finally, in order to further evaluate the simple approximation (1.4), we performed 20 replications for each example for 16 different values of L and plotted the burst ratio B_{SW}/B_{LB} as a function of $\sqrt{L}$. The results in fig. 4 strongly support the simple approximation in (1.4). In fig. 4 we also display the 90% confidence intervals for each estimate of the burst ratio, the regression line fit to the sample means and the line provided by formula (1.4). The close fit of the regression line to the sample means indicates that the burst ratio is well approximated by a linear function in $\sqrt{L}$. Also, a surprising good approximation for the slope is given by (1.4).

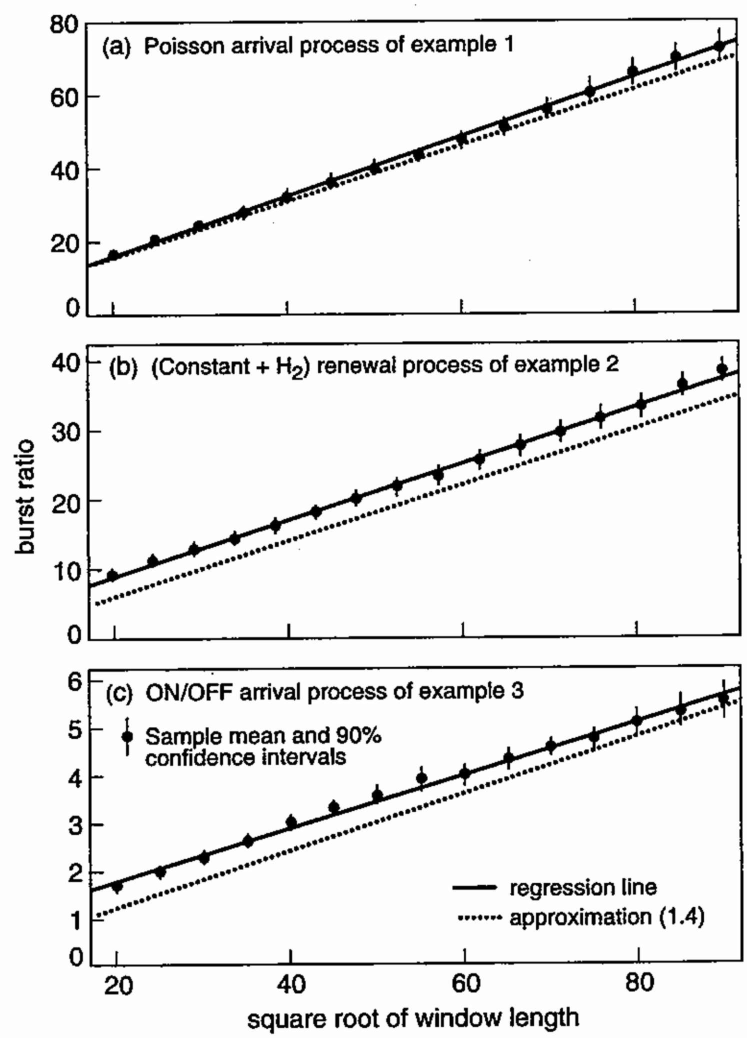

Fig. 4. Mean of burst ratios realized over 20 sample paths.

Acknowledgement

We thank Bharat Doshi for helpful comments on the presentation of this paper.

References

[1] J. Abate, G.L. Choudhury and W. Whitt, Exponential approximation for tail probabilities in queues, I: waiting times, Oper. Res., to appear.

[2] J. Abate, G.L. Choudhury and W. Whitt, Exponential approximations for tail probabilities in queues, II: sojourn time and workload, Oper. Res., to appear.

[3] A.W. Berger and W. Whitt, Asymptotics for open-loop window flow control, J. Appl. Math. Stoch. Anal. 7 (1994) 337–356.

[4] A.W. Berger and W. Whitt, Maximum values in queueing processes, Prob. Eng. Inf. Sci., to appear.

[5] G.L. Choudhury and W. Whitt, Heavy traffic asymptotic expansions for the asymptotic decay rates in the BMAP/G/1 queue, Stoch. Models 10 (1994) 453–498.

[6] B. Doshi, Deterministic rule based traffic descriptors for Broadband ISDN: worst case behavior and connection admission control, *Globecom*, Houston, Texas (1993) p. 1759–1764.

[7] P.W. Glynn and W. Whitt, Logarithmic asymptotics for steady-state tail probabilities in a single-server queue, in: *Studies in Applied Probability*, eds. J. Galambos and J. Gani (Appl. Prob. Trust, Sheffield, UK, 1994) pp. 131–156.

[8] D.P. Heyman, A. Tabatabai and T.V. Lakshman, Statistical analysis and simulation study of video teleconference traffic in ATM networks, IEEE Trans. Circuits and Syst. for Video Tech. 2 (1992) 49–59.

[9] D.L. Iglehart, Extreme values in the GI/G/1 queue, Ann. Math. Statist. 43 (1972) 627–635.

[10] W. Kraemer and M. Langenbach-Belz, Approximate formulae for the delay in the queueing system GI/G/1, *8th Int. Teletraffic Congress*, Melbourne (1976) pp. 235/1–8.

[11] M.R. Leadbetter, G. Lindgren and H. Rootzén, *Extremes and Related Properties of Random Sequences and Processes* (Springer, New York, 1983).

[12] D. Lee, Tight performance bounds for ·/D/1 queues with leaky-bucket-regulated arrivals, Working Paper, MIT (1993).

[13] A.R. Reibman and A.W. Berger, Traffic descriptors for VBR video teleconferencing over ATM networks, IEEE/ACM Trans. Networking 3 (1995) to appear.

[14] J.W. Roberts, *Performance Evaluation and Design of Multiservice Networks*, COST 224, Final Report (Commission of the European Communities, Luxembourg, 1992).

[15] W. Whitt, Tail probabilities with statistical multiplexing and effective bandwidths for multi-class queues, Telecom. Syst. 2 (1993) 71–107.