

FAST SIMULATION OF THE LEAKY BUCKET ALGORITHM

Victor F. Nicola *
Gertjan A. Hagesteijn
Byung G. Kim †

Tele-Informatics and Open Systems
University of Twente
P.O. Box 217
7500 AE Enschede, The Netherlands

ABSTRACT

In this paper we use fast simulation methods, based on importance sampling, to efficiently estimate cell loss probability in queueing models of the Leaky Bucket algorithm. One of these models was introduced by Berger (1991), in which the rare event of a cell loss is related to the rare event of an empty finite buffer in an "overloaded" queue. In particular, we propose a heuristic change of measure for importance sampling to efficiently estimate the probability of the rare empty-buffer event in an asymptotically unstable $GI/GI/1/k$ queue. This change of measure is, in a way, "dual" to that proposed by Parekh and Walrand (1989) to estimate the probability of a rare buffer overflow event. We present empirical results to demonstrate the effectiveness of our fast simulation method. Since we have not yet obtained a mathematical proof, we can only conjecture that our heuristic is asymptotically optimal, as $k \rightarrow \infty$.

1 INTRODUCTION

In an Asynchronous Transfer Mode (ATM) network, data is transported in fixed-size cells. An ATM connection is established with an admission contract which specifies the traffic characteristics of the source and the quality of service (QOS) requirement to be guaranteed by the network. In order for the network to ensure that the admission contract is not violated, the usage parameter control (UPC) procedure is invoked to monitor the actual traffic and to police the excess traffic violating the contract.

The Leaky Bucket (LB) algorithm is a popular UPC procedure and can easily be implemented with counters (see Turner (1986).) Each time a cell arrives, the counter is incremented by one. As long as the counter has a positive value, it is decremented

at fixed intervals, d . When the cell arrival "rate" exceeds the periodic decrement "rate," the counter value will increase. If the counter reaches a pre-specified limit, say, k_T , then the source is considered to have exceeded its admission contract, and subsequent cells are discarded (or marked for policing) until the counter value falls below the limit again. The operation of this LB algorithm can be modeled as a $GI/D/1/k_T$ queue, in which the service time is set identical to the decrement interval, d . For a two-phase burst/silence source model (see Section 3.1), the stationary cell loss probability can be obtained by a numerical method, whose complexity grows in proportion to the value of k_T (Ratgeb 1991.)

The concept of the LB algorithm is sometimes described by means of control tokens and a cell-delay buffer. Tokens are generated at fixed intervals (corresponding to the counter decrement interval introduced above) and are stored in a token bank (buffer) with a finite capacity k_T . A token which arrives to find the token bank full is lost. If the token bank is not empty, then an arriving cell is transmitted instantly and removes one token from the bank. If the token bank is empty, then an arriving cell is either queued in a cell-delay buffer or lost, if this buffer has reached its capacity k_C . For the LB algorithm with two buffers and a certain class of arrival processes, it is shown in Berger (1991) that the steady-state throughput and cell loss depend only on the sum of the token and cell-delay buffer capacities. In other words, with respect to throughput and cell loss, the LB algorithm with two buffers is equivalent to another, with only a token buffer of size of $k_S = k_T + k_C$, and no cell-delay buffer. Therefore, the LB algorithm can be modeled as a $D/GI/1/k_S$ queue, in which the server operates in a slightly non-standard manner (see Section 3.3). In this model, the arrival process corresponds to the token generation process and the service process corresponds to the cell arrival process. For a Markovian arrival process (MAP) of cells, Berger (1991) considers the analysis of an embedded Markov renewal process to determine the probability that a token is

*Victor F. Nicola is currently on leave from IBM T.J. Watson Research Center, Yorktown Heights, NY.

†Byung G. Kim is currently on leave from Department of Computer Science, University of Massachusetts, Lowell, MA.

blocked, which is directly related to the cell loss probability. The main disadvantage of this approach is that a relatively small numerical error in the evaluation of the (high) token blocking probability leads to a significant relative error in the evaluation of the (low) cell loss probability. We note that the rare cell loss event in the operation of the LB algorithm is directly related to the rare empty-buffer event in Berger's model. This is a key observation, since the problem of estimating the cell loss probability in the LB algorithm can be solved if we have a method to estimate the probability of an empty buffer in an "overloaded" $D/GI/1/k_S$ queue. Informally, by "overloaded" we mean that the arrival "rate" is higher than the service "rate," in which case the queue is asymptotically (as the buffer size tends to infinity) unstable.

In order to avoid restrictions necessary for analytic tractability and/or numerical feasibility, simulation is often preferred for the evaluation of realistic models of the LB algorithm. However, accurate estimation of the cell loss probability requires observing numerous cell loss events. But, if the cell loss probability is 10^{-9} , then each cell loss event takes place approximately once in 10^9 cells. Observing a sufficiently large number of cell loss events will take extremely long simulation time.

Importance sampling (Hammersley and Handscorn 1964) has been used effectively to achieve significant speed ups in simulations involving rare events, such as failure in a reliable computer system or cell loss in an ATM communication network. See Nicola et al. (1993) for a review of techniques for fast simulation of highly dependable systems, and Heidelberger (1993) for a survey of efficient simulation methods to estimate buffer overflow probabilities in communication systems. The basic idea of importance sampling is to simulate the system under a different probability measure (i.e., with different underlying probability distributions), so as to increase the probability of typical sample paths involving the rare event of interest. For each sample path (observation) during the simulation, the measure being estimated is multiplied by a correction factor, called the *likelihood ratio*, to obtain an unbiased estimate of the measure in the original system. Asymptotically optimal change of measures (to use in importance sampling) have been found to estimate small probabilities of buffer overflow in relatively simple queueing models (see, Parikh and Walrand (1989), Sadowsky (1991), Chang et al. (1993) and others.) In this paper, we use these change of measures to estimate very small full-buffer probabilities in $GI/D/1/k$ queueing models of the LB algorithm; this is motivated by the model in Rathgeb (1991). We also propose a change of measure to efficiently estimate very small empty-buffer probabilities in asymptotically unstable $D/GI/1/k$ queues; this is motivated by the model in Berger (1991), as

described above. This change of measure is, in a way, "dual" to that proposed by Parikh and Walrand (1989). We experiment with importance sampling methods in simulations of queueing models of the LB algorithm. Two cell arrival processes are considered, namely, a Poisson process (mainly for validation purposes) and a more realistic two-phase burst/silence process (see Section 3.1). Empirical results demonstrate the effectiveness of our method to estimate very small empty-buffer (and cell loss) probabilities. Since we have not yet formally established its asymptotic optimality, this can only be conjectured.

The rest of this paper is organized as follows. In Section 2, we consider change of measures used in importance sampling to speed up simulations of simple queues; both, rare full-buffer and rare empty-buffer events are considered. Experiments with these change of measures to simulate queueing models of the LB algorithm are presented in Section 3. Conclusions are given in Section 4.

2 FAST SIMULATION OF SIMPLE QUEUES

In this section we consider two types of rare events in a simple queue with a finite buffer. Our interest in these rare events is motivated by models of the LB algorithm described in Section 1. The first is the *full-buffer* event, when the cell arrival "rate" is much smaller than the service "rate". Efficient simulation involving a rare full-buffer event has been considered by many (see, for example, Parikh and Walrand (1989) and Sadowsky (1991).) Another rare event of interest is the *empty-buffer* event, when the arrival "rate" is much higher than the service "rate." (Note that the queue is always stable because of its finite buffer.) To the best of our knowledge, efficient simulation involving the latter rare event using importance sampling has not been considered previously.

2.1 Rare Full-Buffer Event

Consider an $GI/GI/1/k$ queue, in which the arrival "rate" is much smaller than the service "rate." Therefore, buffer overflow is a rare event. The probability density function (pdf) of the inter-arrival (resp., service) time is given by $f_A(t)$ (resp., $f_S(t)$.) Let $N(t)$ be the number of jobs in the queue (including that in service) at time t , and denote by $t_j, j = 0, 1, 2, \dots$, the consecutive instants in time at which $N(t)$ jumps from 0 to 1, i.e., for all $j = 0, 1, 2, \dots$, $N(t_j^-) = 0$ and $N(t_j^+) > 0$. Define a *busy cycle* to be the evolution of the process $N(t)$ between two such consecutive instants, say, t_j and t_{j+1} . Note that $t_j, j = 0, 1, 2, \dots$, constitute renewal points, and, therefore, busy cycles are i.i.d. (independent and identically distributed.)

The length of a busy cycle is a r.v. T ; for the j -th busy cycle $T_j = t_j - t_{j-1}$, $j = 1, 2, \dots$. The number of arrivals during a busy cycle is a r.v. N which, because of buffer overflow, is not necessarily equal to the number of departures in the same busy cycle; for the j -th busy cycle it is denoted by N_j . Furthermore, denote by B_j the total time in the j -th busy cycle during which the buffer is full, i.e., $B_j = \int_{t=t_{j-1}}^{t_j} I(N(t) = k) dt$, where $I(\cdot)$ is the indicator function. Also, denote by O_j the total number of cells lost because of a buffer overflow during the j -th cycle. B_j and O_j are realizations of the random variables B and O , respectively. It follows that the long-run (steady-state) fraction of time the buffer is full, \mathcal{F} , and cell loss probability, \mathcal{O} , are given by

$$\mathcal{F} = \frac{E(B)}{E(T)}, \quad (1)$$

$$\mathcal{O} = \frac{E(O)}{E(N)}. \quad (2)$$

Let us consider the estimation of the steady-state full-buffer probability, \mathcal{F} , using the above ratio representation. It is more efficient to estimate $E(T)$ using standard simulation and $E(B)$ using importance sampling. This approach is known as “measure specific dynamic importance sampling” or MS-DIS (see Goyal et al. (1992)). Using importance sampling, we have $E(B) = E_f(B) = E_g(BL)$, where f and g are the original and the new probability measures, respectively, and L is the likelihood ratio. Denote by $dg(\omega)$ the probability of a sample path ω according to the new probability measure g . (Similarly, $df(\omega)$ is the probability of a sample path ω according to the original probability measure f .) Then $L(\omega) = df(\omega)/dg(\omega)$ is the likelihood ratio associated with a sample path ω ; it can be computed easily during the simulation. For example, let $t_{A,j}^i$ (resp., $t_{S,j}^i$), $i = 1, 2, \dots, N_j$, be the cell arrival (resp., departure) instants in the j -th busy cycle. Furthermore, let $g_{A,j}^i(t)$ (resp., $g_{S,j}^i(t)$) be the new i -th inter-arrival (resp., service) time density used to simulate the system with importance sampling. The likelihood ratio, L_j , associated with the j -th busy cycle, takes the form

$$L_j = \prod_{i=1}^{N_j} \frac{f_A(t_{A,j}^{i+1} - t_{A,j}^i)}{g_{A,j}^i(t_{A,j}^{i+1} - t_{A,j}^i)} \times \frac{f_S(t_{S,j}^i - t_{A,j}^i)}{g_{S,j}^i(t_{S,j}^i - t_{A,j}^i)}. \quad (3)$$

Note that $t_{A,j}^{N_j+1} = t_{A,j+1}^1$ is the instant at which the j -th busy cycle ends and the $j+1$ -th busy cycle begins. Thus, L_j can be computed recursively at arrival and departure events during the simulation.

Now, let n_B be the number of independent “biased” (using importance sampling) busy cycles used

to obtain estimates for the mean and the variance of the r.v. BL . These estimates are given by

$$\hat{\mu}_B = \sum_{j=1}^{n_B} B_j L_j / n_B, \quad \hat{\sigma}_B^2 = \sum_{j=1}^{n_B} (B_j L_j - \hat{\mu}_B)^2 / (n_B - 1).$$

Let n_T be the number of independent “normal” (using standard simulation) busy cycles used to obtain estimates for the mean and the variance of the r.v. T . These estimates are given by

$$\hat{\mu}_T = \sum_{j=1}^{n_T} T_j / n_T, \quad \hat{\sigma}_T^2 = \sum_{j=1}^{n_T} (T_j - \hat{\mu}_T)^2 / (n_T - 1).$$

Asymptotically unbiased estimate of \mathcal{F} is given by

$$\hat{\mathcal{F}} = \frac{\hat{\mu}_B}{\hat{\mu}_T}.$$

The relative half-width (in percentage) of the 99% confidence interval for the above estimator is given by $2.56(\hat{\sigma}_{\mathcal{F}}/\hat{\mathcal{F}}) \times 100$, where

$$\hat{\sigma}_{\mathcal{F}}^2 = \left(\frac{\hat{\sigma}_B^2}{n_B} + \hat{\mathcal{F}}^2 \frac{\hat{\sigma}_T^2}{n_T} \right) / \hat{\mu}_T^2.$$

In the following we consider the optimal change of measure (importance sampling distribution) to efficiently estimate $E(B)$. Let $F_A(\theta) = \int_{t=0}^{\infty} e^{\theta t} f_A(t) dt$ be the moment generating function of the inter-arrival times. Define $f_A^\theta(t) = e^{\theta t} f_A(t) / F_A(\theta)$; this is another pdf obtained by exponentially tilting (twisting) the pdf $f_A(t)$ at a parameter θ . Similarly, $F_S(\theta) = \int_{t=0}^{\infty} e^{\theta t} f_S(t) dt$ is the moment generating function of the service times, and $f_S^\theta(t) = e^{\theta t} f_S(t) / F_S(\theta)$ is the corresponding exponentially tilted pdf.

Using heuristic arguments based on the theory of large deviations (Bucklew 1990), Parekh and Walrand (1989) proposed an importance sampling distribution to efficiently estimate the probability of buffer overflow in a $GI/GI/1/k$ queue. In Sadowsky (1991), this distribution was proved to be the unique asymptotically (as $k \rightarrow \infty$) optimal change of measure. Let θ^* be the solution of the equation

$$F_A(-\theta^*) F_S(\theta^*) = 1. \quad (4)$$

Then the optimal change of measure is obtained by simulating the $GI/GI/1/k$ queue with the exponentially tilted densities $g_A(t) = f_A^{-\theta^*}(t)$ and $g_S(t) = f_S^{\theta^*}(t)$. Importance sampling is “turned on” at the start of each busy cycle, and is “turned off” at the occurrence of the rare event. The moment generating functions for the new (optimal) inter-arrival and service times are given by

$$G_A(\theta) = \frac{F_A(\theta - \theta^*)}{F_A(-\theta^*)}, \quad G_S(\theta) = \frac{F_S(\theta + \theta^*)}{F_S(\theta^*)}. \quad (5)$$

Consider the $M/M/1/k$ queue with its arrival rate λ much smaller than its service rate μ (i.e., $\lambda \ll \mu$), so that a full buffer is a rare event. $F_A(-\theta) = \lambda/(\lambda + \theta)$ and $F_S(\theta) = \mu/(\mu - \theta)$, for $\theta < \mu$. Solving the equation $F_A(-\theta^*) F_S(\theta^*) = 1$ for θ^* , we get $\theta^* = \mu - \lambda$. It follows that $G_A(\theta) = \mu/(\mu - \theta)$ and $G_S(\theta) = \lambda/(\lambda - \theta)$, i.e., optimally, the $M/M/1/k$ queue is simulated with arrival rate μ and service rate λ . This change of measure accelerates the arrival process relative to the service process, thus increasing the probability of a full buffer in the simulated system.

In section 3.2, we determine the optimal importance sampling distribution (as outlined above) and use it to estimate very small full-buffer probabilities in $GI/D/1/k$ queueing models of the LB algorithm.

2.2 Rare Empty-Buffer Event

Consider the $GI/GI/1/k$ queue, with its arrival "rate" much higher than its service "rate." Note that the queue is always stable because the buffer is finite, however, reaching the empty-buffer state is a rare event. It follows that busy cycles (as defined in Section 2.1) are extremely long, and it is not practical to use them as individual samples in a regenerative simulation. Therefore, we need to define other, more frequent, cycles to use in our importance sampling simulation. Consider, for example, the consecutive points in time $t_j, j = 0, 1, 2, \dots$, at which an arrival causes the buffer to be full or finds it already full. That is, for all $j = 0, 1, 2, \dots$, t_j is an arrival instant such that $N(t_j^+) = k$. Define a *full-buffer cycle* to be the evolution of the process $N(t)$ between two such consecutive instants, say, t_j and t_{j+1} . Note that, in general, $t_j, j = 0, 1, 2, \dots$, are not renewal points (because of the age of the current service), and, therefore, these full-buffer cycles are neither regenerative nor independent. A methodology to overcome this complication has been described in Nicola et al. (1993). However, in this paper we restrict our discussion to situations where regenerative full-buffer cycles can be identified at some (but not necessarily all) t_j instants. For example, when the service time distribution is of Phase type, a regenerative full-buffer cycle is identified upon an arrival causing a re-entry to the full-buffer state, while the current service is in a particular memoryless state. More formally, let $A(t)$ and $S(t)$ denote the state of the current inter-arrival and the current service (if any) at time t , respectively. Denote by $t_j, j = 0, 1, 2, \dots$, the consecutive instants at which the process $\{N(t), A(t), S(t)\}$ enters the same (regenerative) full-buffer state, i.e., begins a regenerative full-buffer cycle. These regenerative cycles are i.i.d., which we, simply, refer to as *full-buffer cycles*. The length of the j -th full-buffer cycle is $X_j = t_j - t_{j-1}, j = 1, 2, \dots$; a realization of a r.v. X . Furthermore, denote by Y_j the total time in

the j -th busy cycle during which the buffer is empty, i.e., $Y_j = \int_{t=t_{j-1}}^{t_j} I(N(t) = 0) dt$, where $I(\cdot)$ is the indicator function. Y_j is a realization of a r.v. Y . It follows that the long-run (steady-state) fraction of time the buffer is empty, \mathcal{E} , is given by

$$\mathcal{E} = \frac{E(Y)}{E(X)}. \quad (6)$$

In the above equation, the quantity $E(X)$ is easy to estimate using standard simulation. However, estimating $E(Y)$ is difficult, since an empty buffer before the end of a full-buffer cycle is a rare event. Here, too, we use standard simulation to estimate $E(X)$ and importance sampling to estimate $E(Y) = E_f(Y) = E_g(YL)$. An estimate $\hat{\mathcal{E}}$ of \mathcal{E} and its 99% confidence interval can be obtained as described in Section 2.1.

To the best of our knowledge, the problem of estimating the probability of a rare empty-buffer event using importance sampling has not been considered before. In the following we give a heuristic change of measure to efficiently estimate $E(Y)$. This change of measure is, in a way, "dual" to that given in Section 2.1 to estimate the probability of a rare full-buffer event. Following the same notation as in Section 2.1, let θ^* be the solution of the equation

$$F_A(\theta^*) F_S(-\theta^*) = 1. \quad (7)$$

Our "heuristic" change of measure is to simulate the $GI/GI/1/k$ queue with the exponentially tilted densities $g_A(t) = f_A^{\theta^*}(t)$ and $g_S(t) = f_S^{-\theta^*}(t)$. Importance sampling is "turned on" at the beginning of each full-buffer cycle, and is "turned off" at either the occurrence of the rare event or the start of the next cycle. The moment generating functions for the new inter-arrival and service times are given by

$$G_A(\theta) = \frac{F_A(\theta + \theta^*)}{F_A(\theta^*)}, \quad G_S(\theta) = \frac{F_S(\theta - \theta^*)}{F_S(-\theta^*)}. \quad (8)$$

Consider the $M/M/1/k$ queue with its arrival rate λ much higher than its service rate μ (i.e., $\lambda \gg \mu$), so that an empty buffer is a rare event. According to the above heuristic, importance sampling involves simulating the $M/M/1/k$ queue with arrival rate μ and service rate λ . Here too, the arrival and service rates are interchanged. However, this is done to accelerate the service process relative to the arrival process, thus increasing the probability of an empty buffer in the simulated system.

In section 3.3, we follow the above heuristic to determine an importance sampling distribution which we use to estimate very small empty-buffer probabilities in $D/GI/1/k$ queueing models of the LB algorithm. Empirical results demonstrate the effectiveness of our heuristic and supports our conjecture of its asymptotic optimality (see Sadowsky (1991).)

3 EXPERIMENTAL RESULTS

In this section we use fast simulation methods discussed in Sections 2.1 and 2.2 to evaluate two models of the LB algorithm. The empirical results displayed here are limited to the estimation of the steady-state probability of a full cell-delay buffer. The probability of buffer overflow, as seen by arriving cells (i.e., cell loss), can be estimated simultaneously from the same simulation, but it is not included in our results.

The operation of the LB algorithm is described in Section 1, first by means of a counter, then by means of a token bank together with a cell-delay buffer. It should be noted that in the first (“counter”) implementation, cells are either lost or transmitted instantly upon arrival, i.e., they do not experience any delays. Furthermore, counter decrement (corresponding to token generation) stops at level 0, and resumes only when a cell arrives, i.e., it is not totally independent of cell arrivals. On the other hand, in the second (“token bank”) implementation, cells may experience delays if, upon arrival, there are no tokens and the cell-delay buffer is not full. Also, the token generation process is totally independent of cell arrivals. Therefore, the “counter” and “token bank” implementations of the LB algorithm are close but not identical, neither with respect to cell delays nor with respect to throughput and cell loss.

3.1 The Cell Arrival Process

Without loss of generality, we shall consider two cell arrival processes. The first is a Poisson process, for which numerical (non-simulation) results can be obtained; thus making it possible to validate our statistical output from simulation. The second is a two-phase burst/silence process (see Rathgeb (1991)), which we will refer to as TPBS process. This arrival process has been used to model different types of sources, such as packetized voice (see Heffes and Lucantoni (1986)) and interactive data services, and, therefore, it is often used to compare various policing mechanisms. The number of cells per burst is geometrically distributed with a parameter α , and the inter-cell time during a burst is deterministic given by τ . Therefore, transitions from burst to silence occur with a probability α , only at multiples of τ . The duration of the silence phase is exponentially distributed with a mean β^{-1} . The peak cell arrival “rate” is $1/\tau$, and the average cell arrival “rate” $\lambda = (\tau + \alpha/\beta)^{-1}$. It follows that the burstiness, b , of the TPBS source is given by $b = 1/\lambda\tau = 1 + \alpha/\beta\tau$. The pdf of the TPBS inter-arrival time is given by

$$f_A(t) = \begin{cases} 0, & \text{if } t < \tau, \\ 1 - \alpha, & \text{if } t = \tau, \\ \alpha\beta e^{-\beta(t-\tau)}, & \text{if } t > \tau, \end{cases} \quad (9)$$

and its moment generating function is given by

$$F_A(\theta) = e^{\theta\tau} \left[(1 - \alpha) + \alpha \frac{\beta}{\beta - \theta} \right]. \quad (10)$$

$g_A(t) = f_A^{\theta^*}(t)$ is the corresponding exponentially tilted pdf (with a tilting parameter θ^*); its moment generating function is given by $G_A(\theta) = F_A(\theta + \theta^*)/F_A(\theta^*)$. It can be shown that the tilted pdf, $g_A(t)$, is also a TPBS process with the same deterministic burst inter-cell time τ , and with its parameters, $\beta^* = \beta + \theta^*$, and $\alpha^* = \alpha\beta/(\beta + (1 - \alpha)\theta^*)$. In Sections 3.2 and 3.3, the appropriate tilting parameter θ^* is determined, and the tilted pdf, $g_A(t)$, is used as the (new) inter-arrival time density for simulation with importance sampling.

3.2 GI/D/1/k Queueing Model

First we consider the “counter” implementation of the LB algorithm (see Section 1.) The operation of this algorithm corresponds to a single buffer of size k_T , whose content is incremented with each (accepted) cell arrival and is decremented periodically at the token generation rate, $1/d$. An arriving cell is lost if it finds a full buffer. Therefore, in this case, the LB algorithm can be modeled as a $GI/D/1/k_T$ queue (see Rathgeb (1991)) in which a buffer overflow corresponds to a cell loss.

In this section we use importance sampling to efficiently estimate the steady-state probability of a full buffer in the above $GI/D/1/k_T$ queueing model. As outlined in Section 2.1, the optimal change of measure can be obtained by solving Equation (4) for θ^* . The corresponding (exponentially tilted) inter-arrival and service time densities can now be determined from their generating functions as given in Equation (5). It follows that the optimal service times are also deterministic and identical to the original (i.e., no change in the service process.) However, the arrival process does change, so as to increase the probability of the rare full-buffer event. Frater (1990) considered deterministic service and batch-Poisson arrival process. For Poisson and TPBS cell arrival processes, Tables 1 and 2 display point estimates of the full-buffer probability (see Equation (1)) and the relative half-width (in percentage) of their 99% confidence intervals.

First, we consider a Poisson arrival process with rate $\lambda = 0.75$ cells per unit time. The new (optimal) arrival process is also Poisson, however, at an increased rate $\lambda^* = \lambda + \theta^*$, where (from Equation (4)) $\theta^* > 0$ is the non-trivial solution of $\lambda + \theta^* = \lambda e^{d\theta^*}$. The (deterministic) service time is set to $d = 1$ time unit, and we vary the buffer size k_T . In this case, we are simulating an $M/D/1/k_T$ queue, and, therefore, we can validate our statistical output from simulation with numerical results obtained from analysis

(see Cooper (1981).) In Table 1, we list fast simulation estimates of the full-buffer probability, as well as numerical results from analysis. 25600 “biased” (with importance sampling) busy cycles were used to get an estimate of the numerator $E(B)$. Another 25600 “normal” (with standard simulation) busy cycles were used to get an estimate of the denominator $E(T)$. Typically, “normal” busy cycles are much shorter than “biased” busy cycles, and, therefore, the effort to estimate $E(T)$ is much less than that to estimate $E(B)$. Using the same total simulation effort (in real time), for only two entries (corresponding to relatively high full-buffer probabilities) it was possible to obtain meaningful results from standard simulation. Note that fast simulation results are in good agreement with the numerical results from analysis. Also, the relative accuracy is practically the same, regardless of how small is the full-buffer probability being estimated. This verifies the asymptotic optimality of the change of measure used in importance sampling.

For a TPBS arrival process with parameters β , α and τ , the new (optimal) arrival process is also a “TPBS,” with the same τ . However, the parameters β^* and α^* are changed (as given in Section 3.1), so as to appropriately accelerate the cell arrival process. From Equation (4), the optimal tilting parameter $\theta^* > 0$ is the non-trivial solution of $\beta + \theta^* = (\beta + (1 - \alpha)\theta^*)e^{(d-\tau)\theta^*}$. For the original arrival process, we fix α at 0.1, and vary β and τ to experiment with burstiness (b) and the average cell arrival “rate” (λ). In Table 2, λ is decreased in two ways; either by fixing β and increasing τ , or by fixing τ and decreasing β . (Note that, by decreasing either β or τ , we are increasing the burstiness, since $b = 1 + \alpha/(\beta\tau)$.) The (deterministic) service time, d , is set to 10 time units, and the buffer size, k_T , is set to 150. Exact analysis of the $TPBS/D/1/k_T$ queueing model under consideration is difficult because of state dimensionality (see Rathgeb (1991).) For each fast simulation entry in Table 2, 25600 “biased” busy cycles were used to get an estimate of the numerator $E(B)$. Another 25600 “normal” busy cycles were used to get an estimate of the denominator $E(T)$. Using the same total simulation effort (in real time), for only two entries (corresponding to relatively high full-buffer probabilities) we are able to compare with estimates from standard simulation. For all entries in Table 2, fast simulation gives stable estimates and tight confidence intervals; which is an indication of the effectiveness of the used importance sampling method. As expected, the full-buffer probability decreases as we decrease the average cell arrival “rate,” however, the effect of increasing the burstiness is not quite apparent in this experiment.

3.3 $D/GI/1/k$ Queueing Model

In this section we consider the “token bank” implementation of the LB algorithm (as described in Section 1), in which there is a token bank of capacity k_T , as well as a cell-delay buffer of size k_C . For some cell arrival processes (including TPBS), it can be shown that the operation of this algorithm (with respect to the steady-state throughput and cell loss) corresponds to another with only a token bank of capacity $k_S = k_T + k_C$, but no cell-delay buffer. (For Markovian Arrival Processes (MAP), this has been established formally by Berger (1991).) In this case, the LB algorithm can be modeled by a $D/GI/1/k_S$ queue in which the server (independently from the arrival process) continues to service virtual customers, also when the system is empty. Only at virtual service epochs, the service of an actual customer, if any, is completed. In other words, it operates as a standard $D/GI/1/k_S$ queue, except that an arrival to an empty system has a service time which is identical to the time until the next virtual service epoch. In this model, the arrival process corresponds to the (deterministic) token generation process and the service time corresponds to the cell inter-arrival time. Note that, an empty buffer in this model corresponds to a full cell-delay buffer in the LB algorithm. As outlined in Section 2.2, the proposed change of measure is obtained by solving Equation (7) for θ^* . The corresponding (exponentially tilted) inter-arrival and service time densities can now be determined from their generating functions as given in Equation (8). The new arrival (token generation) process is also periodic (deterministic) at the same “rate” $1/d$. However, the service (cell arrivals) process does change, so as to accelerate the occurrence of the rare empty-buffer event. For Poisson and TPBS cell arrival processes, Tables 3 and 4 display point estimates of the empty-buffer probability (see Equation (6)) and the relative half-width (in percentage) of their 99% confidence interval.

First, we consider an exponential service time (corresponding to a Poisson cell arrival process) with rate $\mu = 0.75$ cells per unit time. The new service times are also exponential, however, at an increased rate $\mu^* = \mu + \theta^*$, where (from Equation (7)) $\theta^* > 0$ is the non-trivial solution of $\mu + \theta^* = \mu e^{d\theta^*}$. The inter-arrival (or token generation) period is set to $d = 1$ time unit, and we vary the aggregate buffer size k_S . In this case, we are simulating an $D/M/1/k_S$ queue, and, therefore, we can validate our statistical output from simulation with numerical results obtained from analysis (see Hokstad (1975).) In Table 3, we list fast simulation estimates of the empty-buffer probability, as well as numerical results from analysis. 102400 “biased” full-buffer cycles (as defined in Section 2.2) were used to get an estimate of the numerator $E(Y)$. Another 25600 “normal” full-buffer cycles were used

to get an estimate of the denominator $E(X)$. Typically, "normal" full-buffer cycles are much shorter than "biased" full-buffer cycles, and, therefore, the effort to estimate $E(X)$ is much less than that to estimate $E(Y)$. For each table entry, the same total simulation effort (in real time) is used to obtain a standard simulation estimate. For only two entries (corresponding to relatively high empty-buffer probabilities) it was possible to obtain meaningful results from standard simulation. Note that fast simulation results are in good agreement with the numerical results from analysis. Also, the relative accuracy is practically the same, regardless of how small is the empty-buffer probability being estimated. This supports our claim of asymptotic optimality for the proposed change of measure.

For a TPBS service (corresponding to cell arrival) process with parameters β , α and τ , the new service process is also a "TPBS," with the same τ . However, the parameters β^* and α^* are changed (as given in Section 3.1), so as to appropriately accelerate the service process. From Equation (7), the tilting parameter $\theta^* > 0$ is the non-trivial solution of $\beta + \theta^* = (\beta + (1 - \alpha)\theta^*)e^{(d-\tau)\theta^*}$. For the original service process, we fix α at 0.1, and, alternatively, vary the average cell arrival "rate" (λ) and the burstiness (b). For each entry in Table 4, given λ and b , other parameters, namely, τ and β can be determined uniquely (note that τ can be assigned only integer values.) The inter-arrival time (corresponding to the token generation period), d , is set to 20 time units, and the buffer size, k_S , is set to 150. For each fast simulation entry in Table 4, 204800 "biased" full-buffer cycles were used to get an estimate of the numerator $E(Y)$. Another 204800 "normal" full-buffer cycles were used to get an estimate of the denominator $E(X)$. Using the same total simulation effort (in real time), for a few entries (corresponding to relatively high full-buffer probabilities) we are able to compare with estimates from standard simulation. For all entries in Table 4, fast simulation results give stable estimates and tight confidence intervals, which again demonstrates the effectiveness of our importance sampling method. As expected, the probability of a full cell-delay buffer increases with λ . Generally, it also increases with b , however, it appears much less sensitive to burstiness at higher average cell arrival "rates." The exact analysis for the standard $D/GI/1/k_S$ queue is perhaps feasible for some service processes. However, including the non-standard service feature of our particular model of the LB algorithm makes it considerably more difficult. Here, importance sampling is very useful, as it makes simulation a practical evaluation alternative.

4 CONCLUSIONS

In this paper we have proposed an importance sampling change of measure to efficiently estimate the probability of a rare empty-buffer event in an asymptotically unstable $GI/GI/1/k$ queue. This change of measure is, in a way, "dual" to that proposed by Parekh and Walrand (1989) to estimate the probability of a rare buffer overflow event in an asymptotically stable queue. The problem is motivated by a queueing model (Berger 1991) in which the rare empty-buffer event corresponds to a full cell-delay buffer (which is closely related to cell loss) in the LB algorithm. Experimental results demonstrate the asymptotic efficiency of our heuristic, however, this has not yet been established formally.

REFERENCES

- Berger, A.W. 1991. Performance analysis of a rate-control throttle where tokens and jobs queue. *IEEE J. Select. Areas Commun.* 9, 2: 165-170.
- Bucklew, J.A. 1990. *Large Deviation Techniques in Decision, Simulation and Estimation*. New York, NY: J. Wiley & Sons, Inc.
- Chang, C.S., P. Heidelberger, S. Juneja and P. Shahabuddin. 1993. Effective bandwidth and fast simulation of ATM intree networks. In *Proc. of the Performance '93 conference*.
- Cooper, R.B. 1981. *Introduction to Queueing Theory*. London: Arnold.
- Goyal, A., P. Shahabuddin, P. Heidelberger, V.F. Nicola and P.W. Glynn. 1992. A unified framework for simulating Markovian models of highly reliable systems. *IEEE Trans. Comp.* 41: 36-51.
- Frater, M.R., J. Walrand and B.D.O. Anderson. 1990. Optimally efficient simulation of buffer overflows in queues with deterministic service times via importance sampling. *Australian Telecommunications Research* 24: 1-8.
- Hammersley, J.M. and D.C. Handscomb. 1964. *Monte Carlo Methods*. London: Methuen.
- Heffes, H. and D.M. Lucantoni. 1986. A Markov modulated characterization of packetized voice and data traffic and related statistical multiplexer performance. *IEEE J. Select. Areas Commun.* 4, 6: 856-868.
- Heidelberger, P. 1993. Fast simulation of rare events in queueing and reliability models. In *Models and Techniques for Performance Evaluation of Computer and Communications Systems*, Springer-Verlag, Lecture Notes in Comp. Sc., 729: 165-202.
- Hokstad, P. 1975. The $G/M/m$ queue with finite waiting room. *J. Appl. prob.* 12: 779-792.
- Nicola, V.F., P. Shahabuddin and P. Heidelberger. 1993. Techniques for fast simulation of highly dependable systems. In *Proc. of the Second Inter-*

national Workshop on Performability Modelling of Computer and Communication Systems.

- Nicola, V.F., P. Shahabuddin, P. Heidelberger and P.W. Glynn. 1993. Fast simulation of steady-state availability in non-Markovian highly dependable systems. In *Proc. of the Twenty-Third International Symposium on Fault-Tolerant Computing*, 38-47, IEEE Computer Society Press.
- Parekh, S. and J. Walrand. 1989. A quick simulation method for excessive backlogs in networks of queues. *IEEE Trans. Autom. Contr.* 34, 1: 54-66.
- Rathgeb, E.P. 1991. Modeling and performance comparison of policing mechanisms for ATM networks. *IEEE J. Select. Areas Commun.* 9, 3: 325-334.
- Sadowsky, J.S. 1991. Large deviations theory and efficient simulation of excessive backlogs in a $GI/GI/m$ queue. *IEEE Trans. Autom. Contr.* 36, 12: 1383-1394.
- Turner, J.S. 1986. New directions in communications (or which way to the information age?). *IEEE Commun. Mag.* 25, 10: 8-15.

AUTHOR BIOGRAPHIES

VICTOR F. NICOLA holds the Ph.D. degree in computer science from Duke University, North Carolina. From 1979 he held scientific and research staff positions at Eindhoven University and Duke University. In 1987, he joined IBM T.J. Watson Research Center as a Research Staff Member. He is currently on leave from IBM and holds the position of a Visiting Associate Professor in the group of Tele-Informatics and Open Systems at the University of Twente. His research interests include performance and reliability modeling of computer and communication systems, queueing theory, fault-tolerance and simulation.

GERTJAN A. HAGESTEIJN is a graduate student working towards his M.Sc. degree at the University of Twente. He is currently involved in the implementation of fast simulation techniques for the evaluation of ATM systems.

BYUNG G. KIM obtained the Ph.D. degree in 1980 from the University of Massachusetts at Amherst. He was at Wright State University in Dayton, Ohio, and joined the Computer Science Department of the University of Massachusetts at Lowell in 1984. In 1993-1994, he spent a year at the University of Twente, The Netherlands. His research interests are in the areas of computer networking, telecommunications, and performance evaluations. He is a coauthor of a recent book, *Discrete-time Models for Communication Systems Including ATM*, published by Kluwer Publishing.

Table 1: Estimates of the Full-Buffer Probability in an $M/D/1/k_T$ Queueing Model of the LB Algorithm

k_T	Std. Sim.	Fast Sim.	Anal.
10	1.54×10^{-3} $\pm 13.55\%$	1.47×10^{-3} $\pm 3.72\%$	1.48×10^{-3}
15	8.79×10^{-5} $\pm 47.62\%$	9.40×10^{-5} $\pm 3.72\%$	9.40×10^{-5}
30	---	2.46×10^{-8} $\pm 3.77\%$	2.45×10^{-8}
45	---	6.39×10^{-12} $\pm 3.74\%$	6.38×10^{-12}
60	---	1.66×10^{-15} $\pm 3.80\%$	1.66×10^{-15}

Table 2: Estimates of the Full-Buffer Probability in a $TPBS/D/1/k_T$ Model of the LB Algorithm

τ		$\beta = 0.01$	$\beta = 0.005$	$\beta = 0.001$
1	Std. Sim.	2.74×10^{-3} $\pm 6.55\%$	---	---
1	Fast Sim.	2.80×10^{-3} $\pm 7.43\%$	2.60×10^{-6} $\pm 2.73\%$	1.63×10^{-9} $\pm 1.97\%$
2	Std. Sim.	7.41×10^{-4} $\pm 13.64\%$	---	---
2	Fast Sim.	7.86×10^{-4} $\pm 5.81\%$	4.83×10^{-7} $\pm 2.65\%$	2.90×10^{-10} $\pm 2.09\%$
4	Std. Sim.	---	---	---
4	Fast Sim.	3.08×10^{-6} $\pm 3.14\%$	1.20×10^{-9} $\pm 2.48\%$	6.94×10^{-13} $\pm 2.14\%$

Table 3: Estimates of the Empty-Buffer Probability in a $D/M/1/k_S$ Queueing Model of the LB Algorithm

k_T	Std. Sim.	Fast Sim.	Anal.
10	1.11×10^{-3} $\pm 8.92\%$	1.14×10^{-3} $\pm 3.19\%$	1.14×10^{-3}
15	7.13×10^{-5} $\pm 28.95\%$	7.24×10^{-5} $\pm 3.24\%$	7.23×10^{-5}
30	---	1.85×10^{-8} $\pm 3.27\%$	1.88×10^{-8}
45	---	4.90×10^{-12} $\pm 3.18\%$	4.91×10^{-12}
60	---	1.29×10^{-15} $\pm 3.27\%$	1.28×10^{-15}

Table 4: Estimates of the Empty-Buffer Probability in a $D/TPBS/1/k_S$ Model of the LB Algorithm

b		$\lambda = .0139$	$\lambda = .0208$	$\lambda = .0417$
6	Std. Sim.	---	---	7.76×10^{-4} $\pm 15.34\%$
6	Fast Sim.	1.86×10^{-16} $\pm 2.08\%$	1.16×10^{-9} $\pm 2.83\%$	7.90×10^{-4} $\pm 6.67\%$
12	Std. Sim.	---	---	9.68×10^{-4} $\pm 11.97\%$
12	Fast Sim.	1.32×10^{-9} $\pm 2.62\%$	2.23×10^{-7} $\pm 2.83\%$	9.90×10^{-4} $\pm 7.14\%$
24	Std. Sim.	---	---	8.71×10^{-4} $\pm 10.89\%$
24	Fast Sim.	3.36×10^{-8} $\pm 2.84\%$	8.73×10^{-7} $\pm 3.42\%$	8.83×10^{-4} $\pm 7.36\%$