Proceedings of the 1994 Winter Simulation Conference
ed. J. D. Tew, S. Manivannan, D. A. Sadowski, and A. F. Seila

FAST SIMULATION OF PACKET LOSS RATES IN COMMUNICATION NETWORKS WITH PRIORITIES

Perwez Shahabuddin

IBM T.J. Watson Research Center
P.O. Box 704
Yorktown Heights, New York 10598

ABSTRACT

We review some recently developed fast simulation concepts and techniques used for estimating the probability of rare buffer overflows in a class of queueing models that arise in the analysis of ATM (asynchronous transfer mode) communication networks. We then discuss some new applications of these concepts to estimate buffer loss probabilities in ATM networks where packets have different types of priorities.

1 INTRODUCTION

This paper deals with the estimation of buffer overflow probabilities in a class of queueing models that arise in the context of ATM (asynchronous transfer mode) communication networks. Broadly speaking, an ATM network consists of several nodes connected to each other through a network of switches. Packets of information are sent from one node to another via a sequence of switches called its route. Each packet contains information in its header, about its final destination and/or the route to be used to get to that destination. The function of a switch is to read the header information on each packet and direct it to the next switch on its route. Hence, the rate at which packets depart from a switch is governed by the time it takes for the switch to do this. All switches have finite buffers and packets which arrive at a switch when its buffer is full, are lost. These communication networks are also characterized by the fact that the packet arrival stream from each node may be significantly autocorrelated.

It is customary to model an ATM network as a queueing network in discrete time, i.e., the time axis is discretized. We will assume that each queue in the network has the capacity to serve a fixed number of packets in each discrete time unit. Several external sources feed packets into different switches in the network. To model the autocorrelation in the packet arrival streams from external sources, different models are used. The most commonly used seems to be the Markov Modulated Arrival Process (MMAP), in which each source can be in one of several states. The source may change states after each discrete time unit with transitions being governed by a discrete time Markov chain. The distribution of the number of arrivals from a source in each discrete time unit depends on the state of the source over that time unit. Other autocorrelated arrival stream models, like autoregressive arrival processes are also used to model the arrival streams in such networks.

The queueing network model with the MMAPs described above, may be solved for the buffer overflow probabilities, using Markov chain methods. However, for networks with many switches and external arrival sources, the state space of the Markov chain may be prohibitively large. The problem gets worse in networks with multiple packet classes and priorities. One alternative may be to use simulation.

The main problem with simulation is the rarity of the buffer overflow event. A large number of events have to be simulated in the model before any samples of buffer overflow may be obtained. Thus we have what is termed in the literature as a rare event simulation problem. An approach that is used for such problems is called importance sampling (see, e.g., Hammersley and Handscomb 1964, Glynn and Iglehart 1989). In importance sampling, the stochastic model is simulated with a new probability dynamics, that makes the events of interest occur more frequently. The sample value is then adjusted to make the final estimate unbiased. However, choosing any change of measure that makes the event of interest occur frequently is not enough; *how* it is made to happen more frequently is also very important. For example, an arbitrary change of measure that makes the rare event happen more frequently may give an estimator with an infinite variance. Hence the main problem in importance sampling is to come up with an appropriate change of measure for the rare event simulation problem in hand.

Provably efficient changes of measures for queueing models have been proposed and studied in the past (e.g., Siegmund 1976, Cottrell et. al. 1983, Parekh and Walrand 1989, Sadowsky 1991, etc.). Until recently, most of the study dealt with one queue systems where the arrival stream constitutes an i.i.d. renewal process. Extension of these provably efficient changes of measures to networks have been few and apply mainly to Markovian tandem networks (e.g., Tsoucas 1992, Glasserman and Kou 1993). Heuristical approaches for fast simulation of more general networks may be found in Parekh and Walrand (1989), Frater et. al. (1991) and Devetsikiotis and Townsend (1993). Provably efficient changes of measures for queues with autocorrelated arrival processes were studied in Chang et. al. (1993a) and Juneja (1993) (see, e.g., Asmussen 1985, Lehtonen and Nehriyen 1992 for analogous concepts in the context of risk analysis), fast simulation of Markov fluid models of such queues have been studied in Kesidis and Walrand (1993) and Ridder (1994). Chang et. al. (1993a) also linked fast simulation techniques for ATM switches to the concept of effective bandwidth (see, e.g., Guerin et. al. 1981) of the arrival sources, thus generalizing the class of source models that can be handled and allowing the study to be extended to the class of intree networks. Some critical concepts in Chang et. al. (1993a), dealing with effective bandwidths in intree networks, were also developed independently in de Veciana et. al. (1993). The reader is referred to Heidelberger (1993) for a comprehensive survey covering most of the above results, as well as a detailed list of references.

In this paper we review some of the basic concepts and fast simulation techniques developed for ATM networks in Chang et. al. (1993a). We then discuss some new applications of this theory to ATM networks with multiple classes of packets and different priority rules. In particular, we consider networks with two classes of packets (high priority and low priority) under two different priority rules. In the first priority rule, the higher priority packet gets what we call service priority. This means that in any discrete time instance first all the high priority packets are transmitted. Then if there is any remaining capacity, the lower priority packets get transmitted. In the second priority rule, the higher priority packet gets service priority as well as what we call buffer priority. In buffer priority, as long as the buffer is not full, both high and low priority packets are admitted. When the buffer is full and a high priority packet arrives, then a low priority packet is displaced if it is present in the system. Otherwise the high priority packet is lost. If a low priority packet arrives and the buffer is full, then the low priority packet is always lost.

Another type of priority rule has been considered in Chang et. al. (1993b). Whenever the queue length (buffer contents) is less than a fixed $B_s < B$, then both high and low priority traffic streams are accepted by the switch. However, when the queue length crosses B_s, then low priority arrivals are rejected and only high priority arrivals are admitted to the buffer. When the buffer gets full, both low and high priority packets are rejected. For details of the fast simulation technique, the reader is referred to Chang et. al. (1993b).

The rest of the paper is organized as follows. In Section 2, the general technique of importance sampling for rare event simulation is described. Some of the basic importance sampling concepts and techniques for single queues and intree networks are discussed in Section 3. In Section 4, we study the two different priority rules applied to intree networks with multiple packet classes. Experimental results are given in Section 5 and the conclusion is presented in Section 6.

2 IMPORTANCE SAMPLING

First we illustrate the basic problem of rare event simulation. Let X be a random variable with density $p(\cdot)$ and let $\mathcal{A}_B$ be a rare set in which X can take values. The rare set $\mathcal{A}_B$ is parameterized by B, so that as $B \to \infty$, the set becomes rarer, i.e., $P(\mathcal{A}_B) \to 0$. The problem is to estimate $\gamma_B = P(\mathcal{A}_B) = E_p(1_{\{X\in\mathcal{A}_B\}})$ where the subscript in the expectation denotes the density of the random variable X. In standard simulation, we generate n samples of the random variable X, say $X_1, X_2, \ldots X_n$ and estimate γ_B by using $\hat{\gamma}_B \equiv \sum_{i=1}^{n} 1_{\{X_i\in\mathcal{A}_B\}}/n$. For fixed n, the half width (HW_B) of the confidence interval is (approximately) directly proportional to $\sqrt{Var(1_{\{X\in\mathcal{A}_B\}})} = \sqrt{\gamma_B - \gamma_B^2}$. Consequently, the relative error $RE_B \equiv HW_B/\gamma_B$, is directly proportional to $\sqrt{Var(1_{\{X\in\mathcal{A}_B\}})}/\gamma_B$. It is then easy to see that $RE_B \to \infty$ as $B \to \infty$.

Let $p'(\cdot)$ be another density so that $p'(x) > 0$ whenever $p(x) > 0$ for all x in $\mathcal{A}_B$. Then

$$\gamma_B = \int 1_{\{X\in\mathcal{A}_B\}} p(x)dx = E_{p'}(1_{\{X\in\mathcal{A}_B\}} L_{p'}(X)) \quad (1)$$

where the subscript in the expectation denotes the new density with respect to which the expectation is taken and $L_{p'}(\cdot)$ is the likelihood ratio, i.e., $L_{p'}(x) \equiv p(x)/p'(x)$ whenever $p'(x) > 0$ and 0 otherwise. Equation 1 suggests that we can use $p'(\cdot)$ instead of $p(\cdot)$ to generate n samples of X and then use $\hat{\hat{\gamma}} \equiv \sum_{i=1}^{n} 1_{\{X_i\in\mathcal{A}_B\}} L_{p'}(X_i)$ as an unbiased estimator of γ_B. This is called importance sampling. The problem is to choose a $p'(\cdot)$ so that

$$Var_{p'}(1_{\{X\in\mathcal{A}_B\}} L_{p'}(X)) \ll Var_p(1_{\{X\in\mathcal{A}_B\}}).$$

Suppose the $p'(\cdot)$ is such that for all $X \in \mathcal{A}_B$,

$$d_1 f(B) \leq L_{p'}(X) \leq d_2 f(B) \quad (2)$$

where d_1 and d_2 are constants and $f(\cdot)$ is a function taking positive values with the property that that $f(B) \to 0$ as $B \to \infty$. Then using Equation 1, it can be shown that

$$d_1 f(B) E_{p'}(1_{\{X \in \mathcal{A}_B\}}) \leq \gamma_B \leq d_2 f(B) \tag{3}$$

and $E_{p'}(1_{\{X \in \mathcal{A}_B\}} L^2_{p'}) \leq d_2^2 f^2(B)$. Hence

$$Var_{p'}(1_{\{X \in \mathcal{A}_B\}} L^2_{p'}) \leq d_2^2 f^2(B). \tag{4}$$

Suppose we are able to show that $\liminf_{B \to \infty} E_{p'}(1_{\{X \in \mathcal{A}_B\}}) > 0$. Then using Equation 3 and Equation 4 we can show that the relative error of the importance sampling estimate remains bounded as $\epsilon \to 0$. This is referred to as the *bounded relative error* property.

3 IMPORTANCE SAMPLING FOR SINGLE QUEUE AND INTREE NETWORKS

First consider the problem of estimating the steady state loss probability in a single queue system. Assume that the queue is fed by K external sources, each of which is a MMAP. For simplicity, we consider the simplest form of such an arrival process. Let the kth source be in any of the M_k states $\{0, 1, \ldots M_{k-1}\}$. Let $Y_k(t)$ be the state of the kth source after time t, and let $p_k(i,j) = P(Y_k(t+1) = j | Y_k(t) = i)$. Let the number of packets a source transmits per unit of discrete time, $a_k(t)$, be equal to the current state of the source. Let $A_k(t) = \sum_{i=1}^{t} a_k(t)$ be the total number of arrivals from the kth source until time t, $a(t) = \sum_{k=1}^{K} a_k(t)$ and $A(t) = \sum_{k=1}^{K} A_k(t)$. We assume that the switch has the capacity to dispatch c packets every unit of discrete time. Let B now denote the size of the buffer. Then the number of people in the system at time t is governed by the following Lindley type recursion: $Q(t+1) = (\min(Q(t) + a(t+1), B) - c)^+$. Note that the recursion assumes that, all packets arriving when the buffer is full, are lost. The problem is to estimate the steady state probability of packet loss when B is large.

Consider the Markov chain that has the $Y_k(t)$'s and the $Q(t)$ as its state at time t. We will call this the process Markov chain. Let A be the set of states of this Markov chain with $Q(t) = 0$. Define an A-cycle to be the process between any two consecutive times when the Markov chain enters a state in A. Let X now denote a random sample path of the process Markov chain over an A-cycle. Then the steady state loss probability may be expressed as (see, e.g., Cogburn 1975)

$$\alpha = \frac{E(W)}{E(D)} \tag{5}$$

where $W \equiv W(X)$ is the number of packets lost in an A-cycle (in steady state) and $D \equiv D(X)$ is the total number of packets that arrive in an A-cycle (in steady state). Let $\mathcal{A}_B$ denote the event of hitting the buffer in an A-cycle, when the buffer size is B. Note that as $B \to \infty$, the event $\mathcal{A}_B$ becomes rarer. Thus, as mentioned in Section 2, it becomes harder to simulate for $P(\mathcal{A}_B) = E_p(1_{\{X \in \mathcal{A}_B\}})$ (where now $p(\cdot)$ is the original measure on the sample path X). Note that $E(W) = E(W|\mathcal{A}_B)P(\mathcal{A}_B)$. Given that $\mathcal{A}_B$ has occurred, W is mostly positive. So any importance sampling change of measure that can efficiently simulate for $P(\mathcal{A}_B)$, should also be efficient for simulating for $E(W)$. Also, since D is rarely 0, one can estimate $E(D)$ by standard simulation. Hence we will concern ourselves only with efficient changes of measure for estimating $P(\mathcal{A}_B)$.

Let $\lambda_{k,\theta}$ be the spectral radius of the matrix with elements $\mathcal{A}_k(i,j) = e^{\theta j} p_k(i,j)$ and let $h_{k,\theta}(j)$ be the corresponding eigenvector. Define a new Markov chain transition matrix by

$$p'_{k,\theta}(i,j) = \frac{e^{\theta j} p_k(i,j) h_{k,\theta}(j)}{\lambda_{k,\theta} h_{k,\theta}(i)}.$$

Let $t = 0$ denote the time of the start of an A-cycle in steady state. Let the new measure $p'_\theta(\cdot)$ on the sample space of X be one in which we use the $p'_{k,\theta}(i,j)$'s until the first time when either the A-cycle completes or $Q(t) = B$. If $Q(t) = B$ then we use the $p_k(i,j)'s$ until the A-cycle completes. Let T_B be the time to buffer overflow in an A-cycle. Then for all $X \in \mathcal{A}_B$, the likelihood ratio $L_\theta(X) \equiv L_{p_\theta}(X)$ accumulated over a cycle is given by

$$\begin{aligned} L_\theta(X) &= \prod_{k=1}^{K} \prod_{t=1}^{T_B} e^{-\theta Y_k(t)} \lambda_{k,\theta} \frac{h_{k,\theta}(Y_k(t-1))}{h_{k,\theta}(Y_k(t))} \\ &= e^{-\theta A(T_B)} \prod_{k=1}^{K} (\lambda_{k,\theta})^{T_B} \frac{h_{k,\theta}(Y_k(0))}{h_{k,\theta}(Y_k(T_B))} \end{aligned}$$

Let $\underline{H}(\theta)$ (resp. $\overline{H}(\theta)$) be the minimum (resp. maximum) of $\prod_{k=1}^{K} (h_{k,\theta}(i_k)/h_{k,\theta}(i'_k)$ over all possible $(i_1, \ldots, i_K, i'_1, \ldots, i'_K)$. Because all the arrival processes have a finite number of states, therefore such a minimum (resp. maximum) exists. Then

$$\underline{H}_\theta \leq \frac{e^{\theta A(T_B)} L_\theta(X)}{(\prod_{k=1}^{K} \lambda_{k,\theta})^{T_B}} \leq \overline{H}_\theta \tag{6}$$

Let $B(t)$ be the number of departures from the queue until time t. Since for sample paths in $\mathcal{A}_B$, $Q(t)$ does not hit 0 before T_B, $B(T_B) = cT_B$. Then $A(T_B) = cT_B + O(T_B) + B$ where $O(T_B)$ is the overshoot at time T_B. Suppose we use a $p'_\theta(\cdot)$ with $\theta = \theta^*$, where θ^* is the solution to

$$\frac{\sum_{k=1}^{K} \log(\lambda_{k,\theta})}{\theta} = c \quad \text{and} \quad \theta > 0. \tag{7}$$

(it can be shown that this set of equations has a unique solution). Substituting $\theta = \theta^*$ in Equation 6 and using the relationship between $A(T_B)$ and $B(T_B)$ described above, we get that

$$\underline{H}(\theta^*)e^{-\theta^* B}e^{-O(T_B)} \leq L_{\theta^*}(X) \leq e^{\theta^* B}\overline{H}(\theta^*).$$

Since all the MMAP's have finite number of states, the $O(T_B)$ is bounded from above. Hence $L_{\theta^*}(X)$ satisfies Equation 2 with $f(B) = e^{-\theta^* B}$. If in addition, we are able to prove that $\liminf_{B\to\infty} E_{p'_{\theta^*}}(1_{\{X\in\mathcal{A}_B\}}) > 0$ then we will also have *bounded relative error*. Intutitively, $p'_{\theta^*}(\cdot)$ is such that it makes the queue unstable and so this lower bound is very likely to be true.

Changes of measure with the bounded relative error property may also be found for arrival processes other than MMAP's. The key is the connection between effective bandwidth and fast simulation (Chang et. al. 1993a). For any general arrival process, let $A(t)$ be the number of arrivals in time $(0, t]$. Then the effective bandwidth is defined to be

$$a^*(\theta) = \lim_{t\to\infty} \frac{\log[E(e^{\theta A(t)})]}{\theta t}$$

(if the limit exists). For example, consider the MMAP described above. Let $l_\theta(t)$ denote the likelihood ratio accumulated till time t if we use the change of measure given by $p'_\theta(\cdot)$. Then Equation 6 holds with T_B replaced by t and $L_\theta(X)$ replaced by $l_\theta(t)$. Using that equation we see that

$$a^*(\theta) = \lim_{t\to\infty} \frac{\log[E_{p'_\theta}(e^{\theta A(t)}l_\theta(t))]}{\theta t} = \frac{\log(\prod_{i=1}^{K} \lambda_{k,\theta})}{\theta}.$$

Hence the θ^* used in the simulation for the queue with the MMAP's is the solution of what is termed in the literature as the effective bandwidth equation, i.e., $a^*(\theta) = c$. More generally, suppose that for any general arrival process $\{A(t) : t \geq 0\}$, there exists a family of change of measures $p_\theta(\cdot)$ such that

$$\underline{H}(\theta) \leq \frac{e^{\theta A(t)}l_\theta(t)}{e^{\theta v(\theta)t}} \leq \overline{H}(\theta)$$

for some functions $\underline{H}(\theta)$, $\overline{H}(\theta)$ and $v(\theta)$ that are independent of t. Then $a^*(\theta) = v(\theta)$. Furthermore, if we use the change of measure $p_{\theta^*}(\cdot)$, where θ^* is the solution of the equation $a^*(\theta) = c$, then Equation 2 is satisfied.

Now consider the case of intree networks. Intree networks are feed forward networks where the queues are arranged in the form of tree. The feed forward is in the direction of the root of the tree. External arrivals can also occur at the intermediate queues. The extension of the previous concepts to these networks is based on the following. Consider a queue with capacity c and arrival process $\{A(t) : t \geq 0\}$. Let $\{B(t) : t \geq 0\}$ denote the departure process (which is also the arrival process to the next queue) from the queue. Suppose there exists a family of change of measure $p_{\theta,A}(\cdot)$ on $\{A(t) : t \geq 0\}$ such that

$$\frac{e^{\theta A(t)}L_{\theta,A}(t)}{e^{\theta a^*(\theta)t}} \leq \overline{H}(\theta)$$

for some functions $\overline{H}(\theta)$ and $a^*(\theta)$ (which is the same as the effective bandwidth). Then

$$\frac{e^{\theta B(t)}L_{\theta,A}(t)}{e^{\theta b^*(\theta)t}} \leq \overline{B}(\theta)$$

where $b^*(\theta)$ is defined to be the following. If $\tilde{\theta}$ is the root of the equation $\frac{d}{d\theta}(\theta a(\theta)) = c$, then $b^*(\theta) = a^*(\theta)$ for $\theta \leq \tilde{\theta}$ and $b^*(\theta) = c-(\tilde{\theta}/\theta)[c-a^*(\theta)]$ for $\theta > \tilde{\theta}$ (see Chang et. al. 1993a for an intuitive explanation; see de Veciana et. al. 1993 for independent development of similar ideas).

Let $\{B(t) : t \geq 0\}$ feed into another queue, that also has an external arrival process $\{\breve{A}(t) : t \geq 0\}$ and service capacity $\breve{c}$. Assume that there exists $p_{\theta,\breve{A}}$ so that

$$\frac{e^{\theta \breve{A}(t)}L_{\theta,\breve{A}}(t)}{e^{\theta \breve{a}^*(\theta)t}} \leq \overline{\breve{A}}(\theta)$$

for some functions $\overline{\breve{A}}(\theta)$ and $\breve{a}^*(\theta)$ independent of t. Then

$$\frac{e^{\theta(B(t)+\breve{A}(t))}L_{\theta,\breve{A}}(t)L_{\theta,A}(t)}{e^{\theta(\breve{a}^*(\theta)+b^*(\theta))t}} \leq \overline{\breve{A}}(\theta)\overline{B}(\theta). \qquad (8)$$

Define an A-cycle to start when the second queue becomes empty and let $\mathcal{A}_B$ be the event of buffer overflow in the second queue during an A-cycle. Equation 8 implies that if we use the changes of measure $p_{\theta^*,A}(\cdot)$ and $p_{\theta^*,\breve{A}}(\cdot)$, where θ^* is given by the root of the equation $\breve{a}^*(\theta) + b^*(\theta) = \breve{c}$, then we will get the upper bound in Equation 2. This will guarantee a large variance reduction in the estimation of $P(\mathcal{A}_B)$. However, because we do not have a lower bound, it is difficult to prove the bounded relative error property.

The exact simulation procedure uses a splitting technique combined with batch means (Nicola et. al. 1993). We first run a few A-cycles so that the system (approximately) reaches steady state. After that, each time the queue empties (i.e., an A-cycle starts), we split a process from the original process. The split process uses the importance sampling change of measure as described in this section. We use this split A-cycle to get a sample of W and $L_{\theta^*}(X)$, say W_i and $L_{\theta^*,i}(X)$, and use the original A-cycle to get a sample of D, say D_i. Each original A-cycle together with the corresponding split A-cycle, is referred to as an A-cycle pair. If we simulate n A-cycle pairs (after the initial transient deletion), then Equation 5 suggests that α may be estimated by

$$\hat{\alpha} = \frac{\sum_{i=1}^{n} W_i L_{\theta^*,i}(X)}{\sum_{i=1}^{n} D_i}.$$

Note that if the A-cycles were mutually independent we could have used the regenerative simulation method (Crane and Iglehart 1975) to construct confidence intervals on this estimate. However, since this is not the case here, we have to use a batch means type of procedure where we combine consecutive A-cycle pairs into fixed size batches. Then as is usual, we assume that if the batch size is sufficiently large, then the batches may be treated as being (approximately) independent. For details of this procedure, the reader is referred to Chang et. al. (1993) or Nicola et. al. (1993).

4 APPLICATION TO FAST SIMULATION OF ATM NETWORKS WITH PRIORITIES

In this section we will show how the theory described in Section 3 can be used to simulate for the loss probabilities in networks with priorities. We consider the different priority schemes described in Section 1. First we introduce some notation. We label the high priority packets to be of Class 1 and the low priority packets to be of Class 2. Let $W^{(1)}$ (resp. $W^{(2)}$) be the number of Class 1 (resp. Class 2) packets lost in an A-cycle and $D^{(1)}$ (resp. $D^{(2)}$) be the total number of Class 1 (resp. Class 2) packets that arrive in an A-cycle. Let α_1 and α_2 be the loss probabilities of Class 1 and Class 2 packets respectively. We assume that out of the K sources, the first K_1 sources produce Class 1 packets and the next $K_2 = K - K_1$ sources produce Class 2 packets.

Priority Rule 1 - Service Priority: We can again express $\alpha^{(i)}$, $i = 1, 2$, as

$$\alpha^{(i)} = \frac{E(W^{(i)})}{E(D^{(i)})} \qquad (9)$$

As for the single packet class case, $E(D^{(i)})$'s are easy to estimate using standard simulation, but estimating the $E(W^{(i)})$'s are difficult because of the rareness of the buffer overflow event. Again define $\mathcal{A}_B$ to be the buffer overflow event (i.e., when the total population in the system exceeds B). Then we can express $E(W^{(i)}) = E(W^{(i)}|\mathcal{A}_B)P(\mathcal{A}_B)$. The probability of a Type i packet being lost, given that $\mathcal{A}_B$ has ocurred is not rare. Hence techniques used for the efficient estimation of $P(\mathcal{A}_B)$ will also work for the $E(W^{(i)})$'s.

Priority Rule 2 - Service and Buffer Priority: For estimating the loss probability of Class 2 packets, we use the same change of measure as described in Section 3. This works for estimating $E(W^{(2)})$ as the probability of a Class 2 packet being lost given that $\mathcal{A}_B$ has occurred, is very high. In fact, the probability of a Class 2 packet being lost, at the instant the buffer is crossed, is very close to 1. This is because, given that the buffer is crossed at time t, the only scenario where a Class 2 packet is not lost at time t, is when at time $t-1$ the buffer only had Class 1 packets (which itself has a very small probability of occurring because Class 1 packets get service priority) and all the arriving packets at time t are of Class 1. Consequently $E(W^{(2)}|\mathcal{A}_B)$ is very small and hence techniques that work for estimating $P(\mathcal{A}_B)$ are not likely to be efficient here.

However, note that because Class 1 packets gets both service and buffer priority, they are totally unaffected by the arrival pattern of the Class 2 packets. Hence from their point of view, they are the only packets in the system and the same change of measure as described in the last section can be used here (where now we assume that the system only has the sources that transmit Class 1 packets).

Priority Rule 1 and Rule 2 applied to intree networks: Now consider the case of a two queue intree network as described in the Section 3. Consider estimating the loss probabilities in second queue for the two priority rules. As before we now define an A-cycle to start at the points when the second queue becomes empty. Consider the case of Priority Rule 1. As for the single queue case, we want to make the total number of packets in the second queue reach its buffer level within the A-cycle. Once this happens, then we are assured of a significant probability of packet loss within the A-cycle, for each of the packet types. Hence we need to do the change of measure described in the intree section of Section 3.

For the case of Priority Rule 2, the Class 1 packets are again not affected by the arrival stream of Class 2 packets. Hence to estimate the Class 1 packet loss at the second queue, we simulate the network with only Class 1 arrival sources, using the change of measure described in Section 3. For Class 2 packets, we again have to make the event $\mathcal{A}_B$ happen more frequently. Conditional on $\mathcal{A}_B$, the probability of Class 2 packet loss at station 2 is significant. Hence we again need to do the change of measure that was suggested in Section 3.

5 EXPERIMENTAL RESULTS

Priority Rule 1 - Service Priority: We consider an example from Chang et. al. (1993a). There is a single switch with 16 MMAP sources of the type described in Section 3. Each source may be in either of two states: 0 or 1. The $p_k(0,0)$'s and $p_k(1,1)$'s of the two sources are given by

[0.9, 0.9, 0.9, 0.9, 0.8, 0.8, 0.8, 0.8
0.8, 0.7, 0.7, 0.7, 0.7, 0.6, 0.6, 0.5]

and

[0.7,0.8,0.6,0.5,0.3,0.6,0.6,0.4
0.8,0.5,0.9,0.6,0.5,0.8,0.6,0.9]

respectively. Note that $p_k(0,0) + p_k(1,1) - 1 > 0$ so that all the generated streams are positively correlated as one would expect in practice. The example in Chang et. al. (1993a) assumed that all packets have the same priority. However, we assume that the first 15 sources generate Class 1 packets and the last source generates Class 2 packets. The total arrival rate from the sources, r, is computed to be 6.517. The rates for each class are $r_1 = 4.14$ and $r_2 = 2.37$, respectively. The capacity c is taken to be 8. We estimate the steady state loss probabilities for different values of the buffer size.

Solving 7 we get $\theta^* = 0.34375$. In the actual simulation, for each value of the buffer size, we first run 300 A-cycles without importance sampling so that the simulation reaches steady state. Then we simulate for 60000 A-cycle pairs. The 60000 A-cycle pairs are divided into 2000 batches, with 30 A-cycle pairs in each batch. For each value of the buffer size, we also ran standard simulations for the same CPU time as the one with importance sampling. In the standard simulation, we again run 300 A-cycles so that the system approximately reaches steady state. The instead of simulating A-cycle pairs, we just simulate single A-cycles, without importance sampling, until the CPU time expires. We get a sample of W and $L(X)$ from each A-cycle. Then we form batches of 30 A-cycles each, and use the batch means type of procedure mentioned before to estimate confidence intervals on the estimate.

Results of the experiments are given in Table 1. The third column gives the estimate and the relative error (RE) using importance sampling and the fourth column gives the same using standard simulation. Note how the RE in the importance sampling estimate remains bounded as B increases, whereas in standard simulation the RE seems to increase without bound. In cases indicated by '?', the variance had not stabilized in the standard simulation. Note that in some of the standard simulation cases, there was no buffer loss in any of the simulated A-cycles.

Priority Rule 2 - Service and Buffer Priority: We use the same example as in the case of Priority Rule 1. As mentioned in the previous section, for each buffer size, we have to run two different simulations. The first simulation is with only Class 1 packets (the first 15 sources) and it is used to estimate α_1. The θ^* for this simulation came out to be 1.25. After the initial transient deletion, we again ran for 60000 A-cycle pairs (2000 batches of 30 A-cycle pairs each. The second simulation is with both classes of packets and is used to estimate α_2. The θ^* here is the same as the Priority Rule 1 case. The number of A-cycle pairs simulated was the same as for estimating α_1. Finally we ran standard simulations for the sum of CPU times that was utilized in estimating α_1 and α_2. Results are presented in Table 2. Again note that the RE using importance sampling remains bounded, whereas the RE using standard simulation tends to grow without bounds.

Let r_1 and r_2 be the total arrival rates of of Class 1 and Class 2 packets, respectively. Then, the total rate of packet loss, $\alpha_1 r_1 + \alpha_2 r_2$, should be the same for both Priority Rule 1 and Priority Rule 2. This is because, in the original system, the arrival processes and the buffer sizes in the two cases are the same. Only how one labels the packets that are lost, is different. Let us estimate $\alpha_1 r_1 + \alpha_2 r_2$ by $\hat{\alpha}_1 r_1 + \hat{\alpha}_2 r_2$. The value of this estimator for $B = 10$ is 7.58×10^{-2} and 7.70×10^{-2} in the two cases, respectively. As expected, these two estimators are very close.

Priority Rule 1 and Rule 2 for Intree Networks: As mentioned in the previous section, both Priority Rule 1 and Priority Rule 2 can easily be extended to intree networks. For brevity, we will only consider the case of Priority Rule 2. We modify an example of an intree network from Chang et. al. (1993a) by attaching priorities to the external arrival streams. The network is composed of two queues. Departures from the first queue feed into the second queue. There are eight MMAP sources feeding into the first queue, each with two states 0 and 1. For all these sources $p_k(0,0) = 0.8$ and $p_k(1,1) = 0.5$. The first seven of these sources generate Class 1 packets and the last one generates Class 2 packets. In addition to the departures from the first queue that feed into the second queue, there are also two external sources feeding into the second queue. These are also two state MMAP's with $p_k(0,0) = 0.6$ and $p_k(1,1) = 0.7$. The c for each of the two queues is taken to be 4.

We will only estimate the loss probabilities in the second queue. We fix the buffer size in the first queue at 25 and estimate the loss probabilities for various values of the buffer of the second queue. As in the single queue case, we have to run two different simulations. The first simulation is with only Class 1 packets in the system, and is used for estimating α_1. For this simulation, $\tilde{\theta}$ and θ^* for the importance sampling came out to be 0.4688 and 0.2734, respectively. The second simulation is with both types of packets in the system, and is used for estimating α_2. Here the $\tilde{\theta}$ and θ^* came out to be 0.613 and 0.875, respectively. As mentioned before, for both these simulations, an A-cycle is defined to start whenever the second queue is empty. The run lengths used for the importance sampling estimates were again 60000 A-cycles divided into 2000 batches. Results for experiments are presented in Table 3. Again note the order of magnitude increase in simulation efficiency by using importance sampling. Also note that in the importance sampling cases, the RE remains almost constant for a wide range of values of the loss probabilities ranging from 10^{-5} to 10^{-21} (though no theoretical bounded relative error proof is available yet).

6 CONCLUSION

We reviewed some basic concepts and techniques used in the fast simulation of ATM networks with a single class of packets. We then described some new applications of these concepts to estimating the packet loss rate in an ATM networks with two priority classes. Investigations were carried out for two different priority rules, but the same importance sampling method applies to some more general priority rules. For example, instead of first serving only high priority packets and then using the remaining capacity on the low priority ones, we can allot capacities c_1 and c_2 (with $c_1 + c_2 = c$) to the higher and lower priority packets, respectively. This is called generalized processor sharing (see, e.g., De Veciana and Kesidis 1994). Each priority class definitely gets the capacity allotted to it. It may also consume any capacity that is remaining from the other class. The Priority Rule 1 described in this paper was a special case of this priority rule with $c_1 = c$ and $c_2 = 0$. Using similar arguments as in Section 4, one can see that the importance sampling method for Priority Rule 1 will remain unchanged if we use this more general type of service priority. However some complications arise in the case of Priority Rule 2 with $c_2 > 0$, as now Class 1 packets, can no longer be treated independently of the Class 2's.

Similar to partitioning the service capacity, we can also partition the buffer. We can allot buffers B_1 and B_2 to the high and low priority packets respectively. Each priority first fills its own buffer space. If they find their buffer space full, then they can encroach on the buffer space of the other priority, provided they find some space (they migrate to their own buffer once there is some space there). But then they have to leave the system when a packet of the other priority arrives and finds its allocated buffer full. The second case we studied was a special case of this one with $B_1 = B$ and $B_2 = 0$. For $B_1 > 0$ and $B_2 > 0$, the simulation is no longer straightforward and more research is needed.

REFERENCES

Asmussen, S. 1985. Conjugate processes and the simulation of ruin problems. *Stochastic Processes and their Applications* 20: 213-229.

Chang, C.S., P. Heidelberger, S. Juneja, and P. Shahabuddin. 1993a. The application of effective bandwidth to fast simulation of communication networks. IBM Research Report RC 18877, Yorktown Heights, New York. Earlier version to appear in *Proceedings of the Performance '93 Conference.*

Chang, C.S., P. Heidelberger, and P. Shahabuddin. 1993b. Fast simulation of packet loss rates in shared buffer communications switch. IBM Research Report RC 19250, Yorktown Heights, New York. Submitted for publication.

Cogburn, R. 1975. A uniform theory for sums of Markov chain transition probabilities. *The Annals of Probability* 3: 191-214.

Cottrell, M., J.C. Fort, and G. Malgouvres. 1983. Large deviations and rare events in the study of stochastic algorithms. *IEEE Transactions on Automatic Control AC-28*: 907-920.

Crane, M.A., and D.L. Iglehart. 1975. Simulating stable stochastic systems III: regenerative processes and discrete event simulation. *Operations Research* 23: 33-45.

de Veciana, G., and G. Kesidas. 1994. Bandwidth allocation for multiple classes of service using generalized processor sharing approach. Report number SCC-94-01, Dept. of Electrical and Computer Engineering, University of Texas at Austin.

de Veciana, G., C. Courcoubetis, and J. Walrand. 1993. Decoupling bandwidth for networks: a decomposition approach for resource management for networks. Submitted. Earlier version in *IEEE INFOCOM'94 Proceedings*, 466-473, IEEE Computer Society Press.

Devetsikiotis, M., and J.K. Townsend. 1993. Statistical optimization of dynamic importance sampling parameters for efficient simulation of communication networks. *IEEE/ACM Transactions on Networking* 1: 293-305.

Frater, M.R., T.M. Lenon, and B.D.O. Anderson. 1991. Optimally efficient estimation of the statistics of rare events in queueing networks. *IEEE Transactions on Automatic Control* 36: 1395-1405.

Glasserman, P., and S.G. Kou. 1993. Overflow probabilities in Jackson networks. *Proceedings of the 32nd IEEE Conference on Decision and Control*, 3178-3182, IEEE Press, New York.

Glynn, P.W., and D.L. Iglehart. 1989. Importance sampling for stochastic simulations. *Management Science* 35: 1367-1392.

Guérin, R., H. Ahmadi and M. Naghshineh. 1991. Equivalent capacity and its application to bandwidth allocation in high-speed networks. *IEEE J. Select. Areas Commun.* 9: 968-981.

Hammersley, J.M., and D.C. Handscomb, 1964. *Monte Carlo Methods.* London: Methuen and Co., Ltd.

Heidelberger, P. 1993. Fast simulation of rare events in queueing and reliability models. IBM Research Report RC 19028, Yorktown Heights, New York. Preliminary version appears in *Performance Evaluation of Computer and Communications Systems*, Springer Verlag, Lecture Notes in Computer Science, Volume 729, 165-202.

Juneja, S. 1993. *Efficient Rare Event Simulation of Stochastic Systems.* Ph.D. Thesis, Department of Operations Research, Stanford University, California.

Kesidis, G., and J. Walrand. 1993. Quick simulation of ATM buffers with on-off multiclass Markov fluid sources. To appear in *ACM Transactions on Modeling and Computer Simulation.*

Lehtonen, T., and H. Nyrhinen. 1992. Simulating level-crossing probabilities by importance sampling. *Advances in Applied Probability* 24: 858-874.

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 *Proceedings 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 Transactions on Automatic Control* 34: 54-56.

Ridder, A. 1994. Fast Simulation of Markov Fluid Models. Technical Report, Dept. of Economics and Econometrics, Free University of Amsterdam, Netherlands.

Sadowsky, J.S. 1991. Large deviations and efficient simulation of excessive backlogs in a GI/G/m queue. *IEEE Transactions on Automatic Control* 36: 1383-1394.

Siegmund, D. 1976. Importance sampling in the Monte Carlo study of sequential tests. *The Annals of Statistics* 4: 673-684.

Tsoucas, P. 1992. Rare events in series of queues, *Journal of Applied Probability* 29: 168-175.

AUTHOR BIOGRAPHIES

PERWEZ SHAHABUDDIN is a Research Staff Member at the IBM T.J. Watson Research Center. He received his B.Tech in Mechanical Engineering from the Indian Institute of Technology, Delhi, in 1984. After working for an year in Engineers India Limited, India, he joined Stanford University from where he received a M.S. in Statistics in 1987 and a Ph.D. in Operations Research in 1990. His research interests include modeling and analysis of the performance/availability of computer and communication systems, rare event simulation and simulation variance reduction. He won the first prize in the 1990 George E. Nicholson Student Paper Competition organized by the Operations Research Society of America.

APPENDIX A: TABLES

B	Loss Prob.	Import. Samp. Estimate	Standard Sim. Estimate
10	α_1	$4.40 \times 10^{-4} \pm 9.3\%$	$4.52 \times 10^{-4} \pm 19\%$
	α_2	$3.12 \times 10^{-2} \pm 2.9\%$	$3.12 \times 10^{-2} \pm 4.8\%$
20	α_1	$1.47 \times 10^{-5} \pm 7.9\%$	$1.22 \times 10^{-5} \pm 72\%$
	α_2	$8.45 \times 10^{-4} \pm 3.2\%$	$7.82 \times 10^{-4} \pm 26.5\%$
40	α_1	$1.69 \times 10^{-8} \pm 7.9\%$	0.00±?
	α_2	$1.00 \times 10^{-6} \pm 3.4\%$	0.00±?
60	α_1	$2.15 \times 10^{-11} \pm 7.3\%$	0.00±?
	α_2	$1.21 \times 10^{-9} \pm 3.3\%$	0.00±?%

Table 1: Estimates of steady-state loss probabilities for Priority Rule 1 applied to single queue example

B	Loss Prob.	Import. Samp. Estimate	Standard Sim. Estimate
10	α_1	$3.39 \times 10^{-5} \pm 6.4\%$	0.00±?
	α_2	$3.19 \times 10^{-2} \pm 2.9\%$	$3.21 \times 10^{-2} \pm 4.0\%$
20	α_1	$6.95 \times 10^{-11} \pm 6.2\%$	0.00±?
	α_2	$8.70 \times 10^{-4} \pm 3.3\%$	$8.80 \times 10^{-4} \pm 20\%$
40	α_1	$6.40 \times 10^{-22} \pm 7.9\%$	0.00±?
	α_2	$1.03 \times 10^{-6} \pm 3.4\%$	0.00±?
60	α_1	$5.45 \times 10^{-33} \pm 6.5\%$	0.00±?
	α_2	$1.24 \times 10^{-9} \pm 3.3\%$	0.00±?

Table 2: Estimates of steady-state loss probabilities for Priority Rule 2 applied to single queue example.

B	Loss Prob.	Import. Samp. Estimate	Standard Sim. Estimate
10	α_1	$4.00 \times 10^{-5} \pm 4.2\%$	$5.16 \times 10^{-5} \pm 47\%$
	α_2	$3.99 \times 10^{-2} \pm 3.6\%$	$4.02 \times 10^{-2} \pm 5.7\%$
25	α_1	$2.92 \times 10^{-11} \pm 7.4\%$	0.00±?
	α_2	$6.03 \times 10^{-4} \pm 3.9\%$	$6.81 \times 10^{-3} \pm 21\%$
50	α_1	$4.07 \times 10^{-21} \pm 9.9\%$	0.00±?
	α_2	$6.43 \times 10^{-7} \pm 3.9\%$	0.00±?

Table 3: Estimates of steady-state loss probabilities for Priority Rule 2 applied to intree network example