Optimal Simulation of Buffer Overflows in Queueing Systems with Application to ATM Switches*

Michael R. Frater [†] Jean Walrand[‡] Brian D. O. Anderson[†]

* Work supported by the Australian Telecommunications and Electronics Research Board (ATERB).

[†] Department of Systems Engineering, Australian National University, G.P.O. Box 4, Canberra A.C.T., 2601

[‡] Dept Electrical Engineering and Computer Science, and Electronics Research Laboratory, University of California, Berkeley, CA 94720, USA

Abstract

Simply because of their rarity, the estimation of the statistics of buffer overflows in queueing systems via direct simulation is often very expensive in computer time. Past work on fast simulation using importance sampling has concentrated on systems with Poisson arrival processes and exponentially distributed service times. The main reason for this has been the analytic tractability of these distributions. However, in practical systems, such as ATM switches, service times are often deterministic and constant. This paper demonstrates how one can generate an asymptotically optimal simulation system (in the sense of variance) for queues with deterministic service times and a variety of arrival processes.

1 Introduction

In a queueing system with finite buffers, some proportion of customers arriving at any queue are lost due to buffer overflows. While this number will be small in a properly dimensioned system, it is of interest because there is often a large cost associated with such a loss. However, the very rarity of the event of losing a customer makes direct simulation very costly in terms of computer time, if not impossible. For some simple systems, such as the M/M/1 queue, it is possible to analytically calculate the mean time between overflows, and simulation is unnecessary. However, for more complex systems, it is not generally possible to calculate the recurrence times of buffer overflows.

The use of importance sampling for estimating the statistics of buffer overflows in queueing networks is addressed by [1, 2]. The emphasis in these works is M/M/1 queues and Jackson networks, and it is shown how one can find an asymptotically optimal simulation system for simulating buffer overflows in these systems. However, in practical configurations, it is more usual to have systems whose service time is both deterministic and constant. In this paper, we concentrate on queues with deterministic service times.

Section 2 describes the problem, and summarizes a method that can be used for generating an asymptotically optimal simulation system for arbitrary arrival and service distributions, but with specific reference to queues with deterministic service times. In Section 3, this technique is applied to the simple example of an M/M/1 queue. Specific results for queues with deterministic service times are presented in Sections 4 and 5. The arrival processes analysed are Poisson and Markov-modulated Poisson. A systematic method for generalizing the results for a Poisson arrival stream to perform sub-optimally efficient simulation (but still with a significant speedup factor over direct simulation) of a system with a Markov modulated arrival process is also given. Some simulation results are presented in Section 6.

2 Problem Formulation

2.1 The Model

We consider a queue with a finite buffer of size N, a deterministic server, and some arrival process with average rate λ, as shown in Figure 1. We will assume without loss of generality that the virtual service rate is 1. By sampling the number of customers in the queue immediately after each virtual service, we can form a discrete-time Markov chain whose

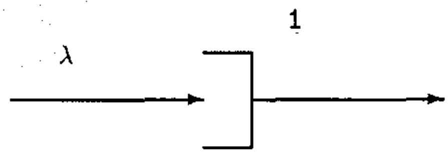

Figure 1: Queue with a deterministic server and arrival process with average rate λ.

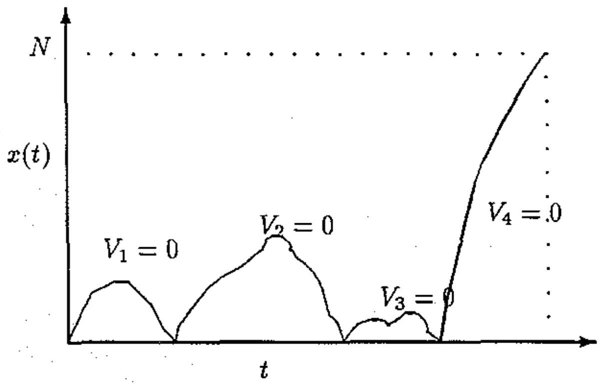

Figure 2: Typical trajectory of a queue.

state is the number of customers in the queue.

If we run the system starting from the empty state, we will observe a trajectory something like that shown in Figure 2. We will use the term *cycle* to denote each piece of trajectory starting with the queue empty and ending with the first time that either the buffer is empty again or overflows. Let τ be the time for an overflow to occur, starting with the buffer empty, and α be the probability that a cycle ends in an overflow. Then we have:

$$\mathrm{E}[\tau] = \frac{\mathrm{E}[J_k]}{\alpha} \tag{1}$$

where J_k is the length of cycle k [1]. (We note that this assumes that the cycles are independent, as is common.)

The expected length of a cycle $\mathrm{E}[J_k]$ is of moderate size. Hence, this quantity is estimated easily via direct simulation. However, the probability that a cycle ends in an overflow α will be very small when overflows are rare. We will use importance sampling to create a system from which α can be estimated much more quickly than is possible via direct simulation.

2.2 Importance Sampling

The idea in importance sampling is as follows. Suppose that we are interested in certain (rare) events in a system S that we can simulate on a digital computer. Instead of simulating S, we simulate a second system $\overline{S}$, which has the property that the events in S and $\overline{S}$ correspond in some way. In particular, to the rare events A in S correspond events $\overline{A}$ in $\overline{S}$ (which may be the same as the events A). The correspondence is such that

1. the events $\overline{A}$ in $\overline{S}$ are more frequent that the events A in S, and
2. the connection between S and $\overline{S}$ allows one to infer $\mathrm{P}(A)$ if one knows $\overline{\mathrm{P}}(\overline{A})$. ($\overline{\mathrm{P}}(\overline{A})$ is the probability of the event $\overline{A}$ in $\overline{S}$.)

Let $V_k = \mathbf{1}_{\{\text{buffer overflows in cycle } k\}}$. Then in our original system S we have:

$$\mathrm{E}[V_k] = \alpha \tag{2}$$

Let L_k denote the likelihood ratio $\frac{dP}{d\overline{P}}$ during cycle k, i.e. the ratio of the probabilities of the trajectories under the measures P and $\overline{P}$ in S and $\overline{S}$. We observe that the L_k are i.i.d. and

$$\overline{\mathrm{E}}[L_k V_k] = \mathrm{E}[V_k] = \alpha \tag{3}$$

Hence, if we simulate the system $\overline{S}$ for p cycles, we can estimate the probability that a cycle ends in an overflow α from:

$$\widehat{\alpha} = \frac{L_1 V_1 + L_2 V_2 + \ldots + L_p V_p}{p} \tag{4}$$

Now we have not yet suggested how the system $\overline{S}$ might be chosen in order to ensure that a good speedup is obtained, or better still, to maximize the speedup obtained. Nor have we defined precisely what we mean by speedup. In many ways, we have replaced one difficult problem (finding the probability of overflow) with another.

2.3 Optimal Simulation - Large Deviations

The problem of finding the best system to use in importance sampling can be posed as an optimization problem as follows. Let A be a rare event for a system S, with $\alpha = \mathrm{P}(A) \ll 1$. For a direct Monte Carlo simulation involving n independent experiments, we could estimate α via:

$$\widehat{\alpha}_n = \frac{1}{n}\sum_{i=1}^{n} \mathbf{1}_A(\omega_i) \qquad (5)$$

where the ω_i are the i.i.d. outcomes of the experiments, and $\mathbf{1}_A$ takes value 1 when the event A has occurred, and zero otherwise. The variance of $\widehat{\alpha}_n$ is easily computed as

$$\mathrm{E}[\alpha - \widehat{\alpha}_n]^2 = \frac{1}{n}(\alpha - \alpha^2) \qquad (6)$$

Alternatively, consider a probability measure $\overline{P}$ associated with a system $\overline{S}$, with P absolutely continuous with respect to $\overline{P}$, such that the same event spaces apply for S and $\overline{S}$. Using $\overline{S}$ we can obtain a second estimate

$$\widehat{\widehat{\alpha}}_n = \frac{1}{n}\sum_{i=1}^{n} 1_A(\overline{\omega}_i)L(\overline{\omega}_i) \qquad (7)$$

where $L = \frac{dP}{d\overline{P}}$ and the $\overline{\omega}_i$ are the i.i.d. outcomes of n experiments using $\overline{S}$. The variance of $\widehat{\widehat{\alpha}}$ is different to (6), and is obtainable as

$$\frac{1}{n}(\int_A L^2(\omega)d\overline{P}(\omega) - \alpha^2) \qquad (8)$$

We want this to be as accurate as possible. So we want to adjust all the probabilities in S to new ones in $\overline{S}$ so that

$$(\sigma^*)^2 = \int_A L^2(\omega)d\overline{P}(\omega) \qquad (9)$$

is minimized. This corresponds to minimizing the time necessary for simulation. In fact, the system $\overline{S}$ that we will find will be asymptotically optimal in the limit as the buffer size tends to infinity.

Given a system $\overline{S}$ minimizing $(\sigma^*)^2$, we can use (4) to find the value of α for the original system S from (much faster) simulation performed on $\overline{S}$.

Let $x(k)$ be the state of a Markov chain formed by sampling the system S, which is now assumed to be a queue. Recall that in Section 2.1, we indicated that with a deterministic server, sampling occurs just after each virtual service; in general, the precise details of the sampling will vary with the arrival and service distributions, and will be outlined separately later in the text. We assume that the state-transition equation for $x(\cdot)$ can be written in the form

$$x(k+1) = x(k) + w(k) \qquad (10)$$

where $w(\cdot)$ is a random process, defined such that $\mathrm{E}[w(k) \mid x(k)] < 0$ for $x(k) > 0$, and with an appropriate boundary condition at $x(k) = 0$ to prevent $x(k+1)$ becoming negative. Hence the Markov chain is asymptotically stable in the sense that, on average, its state will tend towards zero. In solving the optimization problem, we will ignore the boundary condition at $x(\cdot) = 0$ that stops the state going negative.

Let $F(\cdot)$ be the jump distribution of the Markov chain $x(\cdot)$ associated with the system S we wish to simulate. Its Cramer transform[1] $h(y)$ is given by:

$$h(y) = \inf_{s\in\mathbb{R}} \left[sy - \log \int_{-\infty}^{\infty} e^{sz}\mathrm{d}F(z) \right]. \qquad (11)$$

Let

$$V(T, y_0, \ldots, y_{T-1}) = \sum_{k=0}^{T-1} h(y_k) \qquad (12)$$

where y_k is the value of y at time k. It has been shown [4] that in the limit of small noise, minimizing $(\sigma^*)^2$ in (9) is equivalent to minimizing $V(\cdot,\cdot)$ with respect to T and the y_k, subject to the constraint

$$\sum_{k=0}^{T-1} y_k = N \qquad (13)$$

It turns out that, so long as the distribution function $F(\cdot)$ is time and state invariant, the optimal value of y_k is a constant, and does not vary with k. Let $y^\star$ be the optimal value of the y_k. Then $y^\star$ is the unique positive solution of:

$$h(y^\star) = y^\star \frac{d}{dy} h(y^\star) \qquad (14)$$

A proof for this result is given in [5]. Also, the constraint (13) implies $y^\star > 0$. Because $h(\cdot)$ is convex, there is always exactly one positive solution of (14).

It has been shown (see e.g. [1, 4] for a continuous-time form) that this value of $y^\star$ is the average rate of increase of the asymptotically optimal simulation system of S. That is, if we denote the state of the optimal simulation system by $\overline{x}$ and its state-transition equation is:

$$\overline{x}(k+1) = \overline{x}(k) + \overline{w}(k) \qquad (15)$$

[1] More information on the Cramer transform can be found in, for example, [1, 2, 3].

then

$$\mathrm{E}[\overline{w}(k)] = y^\star \qquad (16)$$

(In examples below, we shall indicate how to find the distribution of $\overline{w}$.) Note that because $y^\star > 0$, the simulation system is unstable in the sense that its state will, on average, increase with time.

It is (14) that will be used in the following sections to find the parameters of the optimal simulation system for a number of types of queue with deterministic service times. The simple example of the M/M/1 queue is given in the next section.

3 M/M/1 Queue

Before presenting the original results of the paper, dealing with queues with deterministic service times, we will first summarize how the results of the previous section, and in particular (14), can be applied to find an optimal simulation system for buffer overflows in an M/M/1 queue.

Given an M/M/1 queue, with Poisson arrival stream at average rate λ and exponentially distributed service times with parameter μ, we wish to find a new system that we can use to find the probability of overflow α with the least cost in simulation time. We assume without loss of generality that $\lambda + \mu = 1$. We form a Markov chain by sampling the state of the buffer immediately after each arrival or service takes place. Ignoring the boundary condition at $x = 0$, the transition function can be written:

$$x(k+1) = x(k) + \begin{cases} 1 & \text{probability } \lambda \\ -1 & \text{probability } \mu \end{cases} \qquad (17)$$

The Cramer transform of the Bernoulli jump distribution associated with this queue is [1]:

$$h(y) = \frac{1}{2}\left[(1+y)\log\frac{1+y}{2\lambda} + (1-y)\log\frac{1-y}{2\mu}\right] \qquad (18)$$

Substituting for $h(\cdot)$ in (14), and rejecting the solution with $y < 0$, it turns out that:

$$y^\star = \mu - \lambda \qquad (19)$$

Now, if our optimal simulation queue has arrival rate $\lambda^\star$ and service rate $\mu^\star$, the average rate of increase of this system is $y = \lambda^\star - \mu^\star$. If we assume, without loss of generality, that $\lambda^\star + \mu^\star = 1$, then we have, on account of the claim immediately following (12), as well as (19):

$$\lambda^\star = \mu \qquad (20)$$
$$\mu^\star = \lambda \qquad (21)$$

which corresponds to swapping the arrival and service rates in passing from the original system to the optimal simulation system. This is well known as the optimal simulation system for simulating buffer overflows in an M/M/1 queue, see [1].

For an M/M/1 queue, the likelihood ratio for the k'th cycle is given by

$$L_k = \left(\frac{\lambda}{\mu}\right)^N \qquad (22)$$

This follows by an easy calculation set out in [1].

4 M/D/1 Queue

An M/D/1 queue can be used to model an ATM switch with output buffering [6][2]. If we sample the output of an M/D/1 queue, having Poisson arrival stream with rate λ and deterministic service rate 1, immediately after each service, the probability that the state of the queue has increased by z is (for $z \geq -1$):

$$\frac{1}{(z+1)!}\lambda^{z+1}e^{-\lambda} \qquad (23)$$

The associated jump distribution is evidently:

$$\mathrm{d}F(z) = \sum_{i=-1}^{\infty} \frac{1}{(z+1)!}\lambda^{z+1}e^{-\lambda}\delta(z-i)\mathrm{d}z \qquad (24)$$

We can now use (18) to find the Cramer transform of the distribution:

$$h(y) = (y+1)\log\frac{y+1}{\lambda} + \lambda - (y+1) \qquad (25)$$

and hence the derivative of $h(\cdot)$ is:

$$h'(y) = \log\frac{y+1}{\lambda} \qquad (26)$$

Hence, the average rate of increase of the state of the optimal simulation system (i.e. $y^\star$, the optimal value of y) is the unique positive solution of:

$$\log\frac{y^\star+1}{\lambda} + \lambda - (y^\star+1) = 0 \qquad (27)$$

[2] As is shown in [6], an ATM switch with output buffering has superior delay characteristics to a switch with input buffering, and hence is of more practical interest.

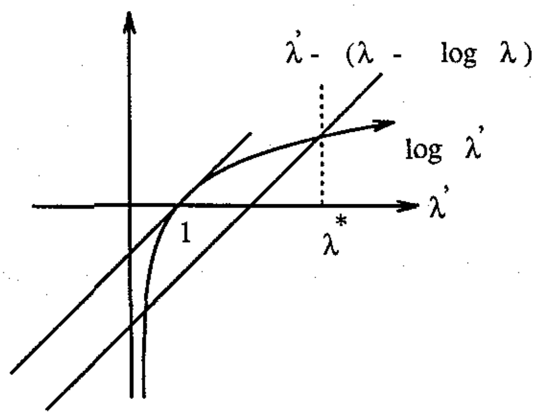

Figure 3: Plots of $\log \lambda^*$ and $\lambda^* - (\lambda - \log \lambda)$.

Now, we can choose the deterministic service rate in the simulation system as we please, provided that we do it in such a way that (27) has a solution. For convenience, we will make the service rate in the simulation system the same as that in the original system, i.e. 1. Hence, the optimal value of the arrival rate in the simulation system is $\lambda^* = y^* + 1$, and $y^* > 0$ requires $\lambda^* > 1$. Replacing $(y^* + 1)$ by λ^* in (27), it can be seen from the plots in Figure 3 that there is exactly one solution for λ^* for which the arrival rate is greater than the service rate.

If we take a second order Taylor series expansion of the logarithm terms in (27), we can find an approximate explicit equation for λ^*:

$$\lambda^* = 2 - \lambda \tag{28}$$

for λ close to 1.

For an M/D/1 queue, the likelihood ratio for the k'th cycle is readily found as

$$L_k = \left(\frac{\lambda}{\lambda^*}\right)^N \tag{29}$$

5 Markov Modulation

The Markov-modulated Poisson arrival process is a generalization of the Poisson, in which at any given time the arrival stream is Poisson, but where the average rate of the stream varies with time according the the state of some Markov process. While it has been possible to directly obtain solutions to the optimization problem for the Poisson arrival process, the case of a Markov-modulated Poisson arrival process presents significantly greater difficulties. In this section, a general sub-optimal heuristic approach, which should still provide good speedup, will be discussed. It will be shown that this heuristic does actually provide optimality where the average arrival rate is a deterministic function of the number of customers in the queue. This latter case is of practical interest, because it can be used to model a multiple-priority traffic system.

5.1 Heuristic Approach

However, while global optimality appears to be difficult to achieve, since at any given time the input stream is Poisson, we can easily achieve a form of local optimality by taking the solution of (27) for the simulation system, i.e. at each time step, we behave as if the average arrival rate is constant, and hence pretend that we are simulating an M/D/1 queue. The following assumptions appear to be implicit in claiming that this is a "good" thing to do globally:

- that there is no explicit time dependence in the modulating process;
- that overflows occur due to "malicious noise sequences" in the Poisson arrival stream, rather than due to the behaviour of the modulating Markov process.

There is no reason to suspect that this approach will provide asymptotic optimality in general, but it is to be expected that significant speedup over direct simulation would be obtained.

5.2 n-Priority Arrival Process

Given an M/D/1 queue for which the arrival rate is a deterministic function of the instantaneous state of the queue,[3] we wish to find a system with which we can simulate buffer overflows with the maximum efficiency.

As before, we will use two processes, one stochastic $x(\cdot)$ and the other deterministic $z(\cdot)$, defined as follows:

$$\begin{aligned} x(k+1) &= x(k) + w(x(k), \omega_k) && (30)\\ z(k+1) &= z(k) + y_k && (31) \end{aligned}$$

[3] This is a special case of Markov modulation, since the state of the queue is a Markov process.

As was done in (10) above, we assume that the random process w is such that $x(\cdot)$ tends to zero on average, i.e. $\mathrm{E}[w \mid x] < 0$ for all $x > 0$, and that there is an appropriate boundary condition that prevents $x(\cdot)$ becoming negative. However, in this case, the conditional distribution of w is a deterministic function of the number of customers in the queue. We will use the solution of an optimal control problem on (31) to find an optimally efficient system for simulating overflows in (30).

Let $h_i(\cdot)$ to be the Cramer transform of the distribution of w when $x(\cdot)$ takes value i, i.e. there are i customers in the queue. We now define the cost function $V(\cdot,\cdot)$:

$$V(T, y_0, \ldots, y_{T-1}) = \sum_{k=0}^{T-1} h_{\mathrm{i}(z(k))}(y_k), \qquad (32)$$

where $\mathrm{i}(z)$ is the integer part of z. We wish to minimize this cost function, subject to the same constraints as were used previously:

$$\begin{aligned} z(0) &= 0 & (33)\\ z(T) &= N & (34) \end{aligned}$$

Since the average rate of arrivals is a deterministic function of the instantaneous state of the queue, we can rewrite the cost function:

$$\begin{aligned} V(T, y_0, \ldots, y_{T-1}) &= \sum_{k=0}^{T-1} h_{\mathrm{i}(z(k))}(y_k) & (35)\\ &= \sum_{j=0}^{N} \sum_{k=T_{j-1}}^{T_j - 1} h_j(y_k) & (36)\\ &= \sum_{j=0}^{N} V_j & (37) \end{aligned}$$

where T_j is the time that the integer part of the state $z(\cdot)$ is first j, starting at $z = 0$, and V_j is the least cost required to move the system $z(\cdot)$ from state $z(T_j)$ to another state greater than or equal to $\mathrm{i}(z(T_j)) + 1$, which is $j + 1$. Because the value of $\mathrm{i}(\cdot)$ is discrete, it is possible that the discretised optimal trajectory $\mathrm{i}(z)$ will not in fact pass through all integer values between 0 and N, if, for example, the optimal values of some y_k are greater than 1. For convenience, we define $V_j = 0$ for states j that are skipped by the discretised optimal trajectory, and also $T_j = T_{j-1}$.

It is clear that $V(\cdot,\cdot)$ is minimized by individually minimizing the V_i. Each of these minimization

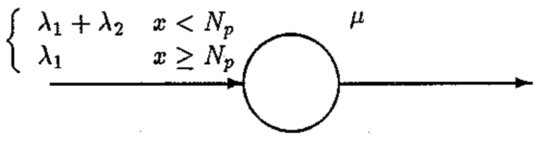

Figure 4: M/M/1 queue with two-priority arrival process

problems is the same as the problem treated in previous sections for the M/M/1 and M/D/1 queues, with the buffer size set to 1, and can be solved using (14), as shown previously.

We note that in this special case, the heuristic approach outlined in Section 5.1 does in fact generate a simulation system that is actually optimal.

As a simple example of the application of this result, we will show how to generate the optimal simulation system for an M/M/1 queue with a two priority arrival process, as shown in Figure 4. The higher priority process, which may correspond to signalling traffic in the network, has rate λ_1. The lower priority arrival process, which has rate λ_2, is blocked whenever the state of the queue exceeds some value, say N_P. As before, the service process has rate μ.

The result above says that we can deal separately with the two different arrival rates seen by the queue. For a simple M/M/1 queue with constant average arrival rate, we generate its optimal simulation system by exchanging the arrival and service rates. Therefore, for the two priority process, we do the same. Hence, in the optimal simulation for the two-priority arrival process, the arrival rate is always μ. When the state of the queue is less than or equal to N_p, the service rate is $\lambda_1 + \lambda_2$. When the state of the queue is greater than N_p, the service rate is λ_1. This system is shown in Figure 5.

6 Simulation Results

A number of simulations were performed on M/D/1 queues in order to ascertain both the amount of speedup that is obtained by estimating the overflow probability α via simulation of the optimal systems described above, rather than direct simulation, and also the amount of computer time required for the

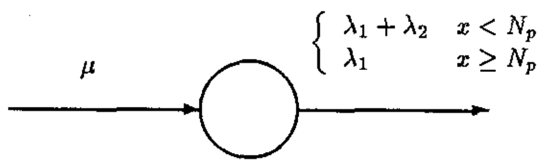

Figure 5: Optimal Simulation system for M/M/1 queue with two-priority arrival process

fast simulation. All simulations were run until the relative standard deviation was less than 0.1, corresponding to a 95 % confidence of the estimation error being less than 20 %.

Table 1 shows the increases in speed obtained, as well as the number of simulation steps required for the various simulations. On the Sparcstation 1 on which these simulations were run, each simulation step took approximately 2.5 microseconds. Hence, it can be seen that if the buffer size grows large, then direct simulation becomes impossible very quickly. Even for the small buffer size used in these simulations (10), direct simulation of the two most lightly loaded cases would have required over one hour of CPU time, and over one billion simulation steps.

Table 1: Simulation results. The buffer size is 10 in all cases.

λ	λ^*	Simulation Times		Speedup
		Direct[4]	Fast	
0.1	3.71	$\gg 10^9$	37	$\gg 10^7$
0.3	2.36	$\gg 10^9$	80	$\gg 10^7$
0.5	1.76	1.3×10^8	142	9.2×10^5
0.7	1.37	760922	1466	520

The table illustrates that the more "stable" the original system is, i.e. the smaller λ and the larger the recurrence time of buffer overflows, the more "unstable" is the importance sampling system (and thus the shorter is the simulation time, and hence the greater the speedup factor. This phenomenon appears quite general in its occurrence in importance sampling applied to queues.

7 Conclusion

This paper has presented methods using importance sampling and Large Deviations theory for performing asymptotically optimal simulation of queues with deterministic servers, and a variety of different arrival processes.

These ideas have yet to be extended to networks of queues, and the relationship between the arrival and service processes and the optimal control problem used to generate the optimal simulation system needs to be more fully explored.

References

[1] S. Parekh and J. Walrand, "A quick simulation of excessive backlogs in networks of queues," *IEEE Trans. Automatic Control*, vol. 34, pp. 54 - 66, January 1989.

[2] M. R. Frater and B. D. O. Anderson, "Fast estimation of the statistics of excessive backlogs in tandem networks of queues," *Australian Telecommunication Research*, vol. 23, pp. 49 - 55, May 1989.

[3] S. R. S. Varadhan, *Large Deviations and Applications*. Philadelphia, PA: Society for Industrial and Applied Mathematics, 1984.

[4] M. Cottrell, J. C. Fort, and G. Malgouyres, "Large deviations and rare events in the study of stochastic algorithms," *IEEE Trans. Automatic Control*, vol. AC-28, pp. 907 - 918, September 1983.

[5] M. R. Frater, J. Walrand, and B. D. O. Anderson, "Optimally efficient simulation of buffer overflows in queues with deterministic service times," *Australian Telecommunication Research*. To appear.

[6] M. J. Karol, M. G. Hluchyj, and S. P. Morgan, "Input versus output queueing on a space-division packet switch," *IEEE Trans. Communications*, vol. COM-35, December 1987.

[7] D. P. Bertsekas and R. G. Gallager, *Data Networks*. Englewood Cliffs, New Jersey: Prentice-Hall, 1987.