

EFFICIENT SIMULATION OF MULTICLASS QUEUEING NETWORKS

Shane G. Henderson

Department of Industrial and
Operations Engineering
The University of Michigan
1205 Beal Avenue
Ann Arbor, MI 48109-2117, U.S.A.

Sean P. Meyn

Department of Electrical and Computer Engineering
University of Illinois-Urbana/Champaign
Urbana, IL 61801, U.S.A.

ABSTRACT

General multiclass queueing systems are extremely difficult to analyze. A great deal of effort has been devoted to examining the question of stability and performance of such networks. However, the simulation of multiclass queueing networks as a tool for performance evaluation has received little attention.

We generate useful control variates for the steady-state simulation of multiclass queueing networks with Markovian structure. The resulting variance reductions greatly outweigh the cost of solving a minimization problem prior to the simulation, as evidenced through numerical examples.

1 INTRODUCTION

Queueing networks are an extremely useful class of models that have found application in many areas including communication networks (Bertsekas and Gallager 1987), computer systems (Lavenberg 1983) and manufacturing systems (Buzacott and Shanthikumar 1993, Gershwin 1993). Much of this success is perhaps due to the inherent tractability of a large class of queueing networks (Kelly 1979, Baskett et al. 1975). However, for many queueing networks, the questions of stability and performance evaluation are extremely challenging problems. In this paper we assume stability, and consider the performance evaluation of a special class of queueing networks known as re-entrant lines (see §2). By performance we mean the expected steady-state number of jobs in the re-entrant line. Our results are easily extended to cover a broader class of systems, and other performance measures (see §8).

Stability analysis and performance evaluation of queueing networks have received a great deal of attention recently. This is partly due to several examples that demonstrate that the usual conditions (traffic intensity less than one at each station) are not sufficient for stability, even under the well-known FIFO

(first in-first out) service discipline (see, for example, Rybko and Stolyar 1993, Bramson 1993).

Methods for establishing the stability of queueing networks have been developed by several authors, based on fluid limits (Dai 1995, Chen 1995), and Lyapunov functions (for example Kumar and Meyn 1995, 1996). Mathematical programming approaches for obtaining upper and lower performance bounds for a queueing network that is assumed stable were simultaneously developed by Bertsimas et al. (1994) and Kumar and Kumar (1994a). Ou and Wein (1992) developed process approximations, the performance of which lower bound the original system's performance. The lower bounds are evaluated by simulating the approximation.

In this paper we simulate the network, but use ideas from the mathematical programming techniques of Bertsimas et al. (1994) and Kumar and Kumar (1994a) to develop effective control variates for the purposes of variance reduction. The basic idea behind control variate schemes is to add random variables with zero mean to an estimator in an attempt to reduce its variance. Because the random variables have zero mean, no bias is introduced. For a general introduction to the variance reduction method of control variates, see Law and Kelton (1991) or Bratley, Fox and Schrage (1987).

Our approach may be viewed as an adaptation of the Approximating Markov Process method (Henderson 1997), a variance reduction technique developed to improve the efficiency of steady-state simulations of Markov processes. The basic idea in Henderson (1997) is to approximate the solution to Poisson's equation (see §5), and then construct an estimator that incorporates the approximation. Andradóttir, Heyman and Ott (1993) also used approximations to the solution to Poisson's equation in simulating finite state-space discrete-time Markov chains.

This paper is organized as follows. In §2 the re-entrant line model is introduced. Then, in §3, control variates are derived and the estimators are proposed.

Asymptotic behaviour of the estimators is explored in §4. In §5, we discuss the choice of multipliers for the control variates derived in §3. In order to obtain confidence intervals (C.I.'s) based on the estimators, we must be able to estimate certain variances and covariances. This is done using the method of batch means, as described in §6. Several examples are presented in §7 and §8 summarizes the paper and discusses future research directions.

All vectors are assumed to be column vectors.

2- THE BASIC OPEN RE-ENTRANT LINE

For simplicity, we present our results in the context of re-entrant lines, which are special cases of multiclass queueing networks. The approach can be extended in a straight-forward manner to more general systems (such as those with probabilistic routes etc., as described in Kumar and Kumar 1994a).

Consider a system consisting of s stations (or machines). Jobs arrive to the system according to a Poisson process at rate λ . All jobs follow a deterministic route through the system. Jobs first visit station $\sigma(1) \in \{1, \dots, s\}$, and are considered to be stored in buffer 1 while either waiting for or receiving service. After completing service at station $\sigma(1)$, jobs then follow the route $\sigma(2), \sigma(3), \dots, \sigma(\ell)$, and are stored in buffers $2, 3, \dots, \ell$ respectively. After completing service at buffer ℓ , jobs leave the system. Service times for jobs in buffer i are exponentially distributed with mean μ_i . All service and interarrival times are independent. We assume that the traffic intensity at each station is less than one, i.e., $\sum_{i:\sigma(i)=\sigma} \lambda/\mu_i < 1, \forall \sigma$.

Let $X_i(t)$ denote the number of jobs present at time t in buffer i , and let $X(t) = (X_1(t), \dots, X_\ell(t))$ be the vector of buffer contents at time t . Let $W_i(t)$ be the fraction of station $\sigma(i)$'s effort devoted to processing jobs in buffer i at time t , and let $W(t) = (W_1(t), \dots, W_\ell(t))$ be the vector of station effort at time t . Then $W_i(t) \geq 0$ for all i and t , and for all stations σ and all t , $\sum_{i:\sigma(i)=\sigma} W_i(t) \leq 1$.

We require the scheduling policy adopted by stations in allocating their effort to be *stationary* and *non-idling*. By stationary, we mean that $W(t)$ is a function of $X(t)$ alone. This precludes such disciplines as FIFO (which rely on additional information such as the order in which jobs arrive), but allows (for example) pre-emptive priority and processor sharing policies. By non-idling, we mean that if $\sum_{i:\sigma(i)=\sigma} X_i(t) > 0$ (there are jobs present at station σ), then $\sum_{i:\sigma(i)=\sigma} W_i(t) = 1$ (the station is allocating all of its effort). Because the station may allocate effort to empty buffers, this constraint does not *exactly* specify the non-idling condition, but it is a reasonable

approximation.

The stationary scheduling policy assumption allows us to conclude that $X = (X(t) : t \geq 0)$ is a continuous-time Markov chain. Now rescale time so that $\lambda + \sum_{i=1}^{\ell} \mu_i = 1$ and uniformize; see Lippman (1975). Let the times $\{\tau_n\}$ correspond to either arrivals, real service completions, or virtual service completions. The control variates that we derive in the next section are a linear function of the products $W_i(\tau_n)X_j(\tau_n)$, and so are not necessarily a linear function of the $X_j(\tau_n)$'s. Therefore, define the process $Z = (Z(n) : n \geq 0)$, where $Z(n)$ is a column vector of the form

$$(Z_{ij}(n) = W_i(\tau_n)X_j(\tau_n) : 1 \leq i, j \leq \ell).$$

The process Z is a Markov chain, and it is important to note that $X(\tau_n)$ is a deterministic function of $Z(n)$, as can be seen from the equality (for each j)

$$X_j(\tau_n) = \sum_{i:\sigma(i)=\sigma(j)} Z_{ij}(n).\psi \quad (1)$$

Remark: Notice that for all σ , $\sum_{i:\sigma(i)=\sigma} W_i(\tau_n) \leq 1$, so that for $\sigma \neq \sigma(j)$

$$X_j(\tau_n) \geq \sum_{i:\sigma(i)=\sigma} Z_{ij}(n).\psi \quad (2)$$

Let the state space of Z be S . Our basic assumption on the process Z is that

(A) there exists a function $V : \mathbb{R}_+^{\ell^2} \rightarrow \mathbb{R}_+$ satisfying

1. V is equivalent to a quintic in the sense that for some $\gamma < 1$,

$$\gamma(\|z\|^5 + 1) \leq V(z) \leq \gamma^{-1}(\|z\|^5 + 1); \text{ and}$$

2. for some $\eta > 0$, and all z ,

$$PV(z) \leq V(z) - \|z\|^4 + \eta,$$

$$\text{where } PV(z) \triangleq E[V(Z(1)) | Z(0) = z].$$

This condition will be satisfied if the stability LP of Kumar and Meyn (1995) admits a solution. In this case, the function V takes the form $V(z) = (z^T Q z)^{5/2}$, where Q is a matrix of suitable dimension.

Alternatively, if a fluid model is stable, it is shown in Dai and Meyn (1995) that if we define $\tau_0 = \inf\{n \geq 1 : Z(n) = 0\}$, then the function

$$V(z) = E \left[\sum_{k=0}^{\tau_0} \|Z(k)\|^4 \middle| Z(0) = z \right]$$

is bounded as in condition 1, and this function is known to satisfy condition 2 (Meyn and Tweedie 1993, p. 338).

Under **(A)**, the chain admits a unique invariant probability distribution π , and $E_\pi \|Z(0)\|^4 < \infty$ (Meyn and Tweedie 1993, p. 330), where

$$E_\pi(\cdot) \triangleq \int_S E(\cdot | Z(0) = z) \pi(dz).$$

3- ESTIMATORS AND CONTROLS

Recall that we wish to estimate the mean steady-state number of customers in the system. From (1)

$$\begin{aligned} \sum_{j=1}^{\ell} X_j(\tau_n) &= \sum_{j=1}^{\ell} \sum_{i:\sigma(i)=\sigma(j)} Z_{ij}(n) \\ &= c^T Z(n), \psi \end{aligned} \quad (3)$$

where c is a vector of ones and zeros, and c^T denotes the transpose of the vector c . We may define the mean steady-state number of customers in the system as $c^T \bar{z}$, where $\bar{z}_{i,j} \triangleq E_\pi Z_{i,j}(0)$ and \bar{z} is the column vector of $\bar{z}_{i,j}$'s.

Define the *standard* estimator $q_0(n)$ of $c^T \bar{z}$ to be $c^T \bar{Z}(n)$, where

$$\bar{Z}(n) \triangleq n^{-1} \sum_{k=0}^{n-1} Z(k). \psi \quad (4)$$

We now introduce a collection of random variables with steady-state mean zero that can be used as control variates. Recall that π is the steady-state distribution of Z . If $E_\pi |h(Z(0))| < \infty$, then $E_\pi (Ph - h)(Z(0)) = 0$, so that $Ph - h$ has steady-state mean zero, suggesting possible control variate applicability.

First consider the function

$$h_j(z) \triangleq \sum_{i:\sigma(i)=\sigma(j)} z_{ij} = x_j,$$

where x is the vector of buffer levels corresponding to the state z . Let w be the vector of work allocations corresponding to z . In one transition, either an arrival may occur to buffer 1 (with probability λ), a real service completion may occur at buffer k (with probability $\mu_k w_k$), or a virtual service completion may occur at buffer k (with probability $\mu_k(1 - w_k)$). Therefore,

$$Ph_j(z) - h_j(z) = \mu_{j-1} w_{j-1} - \mu_j w_j, \psi \quad (5)$$

where $\mu_0 \triangleq \lambda$ and $w_0 \triangleq 1$. Condition **(A)** implies that $E_\pi |h_j(Z(0))| < \infty$, so that (5) has zero steady-state mean for all j , and it follows that

$$\bar{w}_j \triangleq E_\pi W_j(0) = \lambda / \mu_j.$$

We have shown that the mean steady-state work allocation at buffer j is λ / μ_j , as expected by work conservation.

The same approach may be used for functions of the form $h_{jk}(z) = x_j x_k$, where once again x and w are the buffer size and workload allocation vectors corresponding to z . Suppose first that $k \geq j + 2$. By considering the possible transitions as before,

$$\begin{aligned} Ph_{jk}(z) - h_{jk}(z) &= \mu_{j-1} w_{j-1} x_k - \mu_j w_j x_k \\ &\quad + \mu_{k-1} w_{k-1} x_j - \mu_k w_k x_j. \end{aligned}$$

Taking steady-state expectations yields

$$\begin{aligned} 0 &= \mu_{j-1} \bar{z}_{j-1,k} - \mu_j \bar{z}_{j,k} + \mu_{k-1} \bar{z}_{k-1,j} - \mu_k \bar{z}_{k,j} \\ &= r_{jk}^T \bar{z}, \end{aligned}$$

where \bar{z} is the vector of steady-state means (of dimension ℓ^2) defined at the start of this section, and r_{jk} is a vector of coefficients.

By considering other $j \leq k$, and noting that $\bar{w}_i = \lambda / \mu_i$, we determine that $r_{jk}^T \bar{z} = \gamma_{jk}$, where

$$\gamma_{jk} = \begin{cases} -2\lambda, \psi & \text{if } j = k, \\ \lambda, \psi & \text{if } k = j + 1, \\ 0, \psi & \text{if } k > j + 1, \end{cases}$$

and r_{jk} is defined as follows. Let e_{jk} be the vector that is zero except for a one in the (j, k) th position, so that $e_{jk}^T z = z_{jk}$. Then for $k = j$, $k = j + 1$ and $k > j + 1$, the vector r_{jk} is given by

$$\begin{aligned} &2\mu_{j-1} e_{j-1,j} - 2\mu_j e_{j,j}, \\ &\mu_{j-1} e_{j-1,j+1} - \mu_j e_{j,j+1} + \mu_j e_{j,j} - \mu_{j+1} e_{j+1,j}, \text{ and} \\ &\mu_{j-1} e_{j-1,k} - \mu_j e_{j,k} + \mu_{k-1} e_{k-1,j} - \mu_k e_{k,j} \end{aligned}$$

respectively. There is one equation for each pair (j, k) where $1 \leq j \leq k \leq \ell$, so that there are $\ell(\ell + 1)/2$ equations in all. These equations were derived in Kumar and Kumar (1994a) using the above approach, and using a slightly different approach that requires weaker conditions in Kumar and Meyn (1996). Bertsimas et al. use the same basic ideas to construct inequalities.

Let R be the matrix where the (j, k) th column contains r_{jk} , and let γ be the vector of γ_{jk} 's. In matrix notation, we have shown that $R^T \bar{z} = \gamma$. These relations may be used to obtain control variates in a simulation of the re-entrant line. In fact, one may envisage two approaches at this point.

1. Simulate X up to time n storing the entire sample path. Then estimate the mean steady-state number in the system by

$$q_1(n) = c^T \bar{Z}(n) + \hat{\alpha}(n)^T (R^T \bar{Z}(n) - \gamma),$$

choosing the vector $\hat{\alpha}(n)$ after the simulation is complete in an attempt to minimize the variance of the estimator $q_1(n)$.

- Choose the vector α prior to the simulation, then simulate X , and use the estimator

$$q_2(n) = c^T \bar{Z}(n) + \hat{\beta}(n) \alpha^T (R^T \bar{Z}(n) - \gamma), \quad (6)$$

where $\hat{\beta}(n)$ is chosen after the simulation is complete in an attempt to minimize the variance of the estimator $q_2(n)$.

At first sight, approach 1 may seem somewhat storage intensive, but after noting that we merely need to store the initial state and the sequence of events, it is apparent that the storage requirements would not be overwhelming. However, for a large re-entrant line with many buffers, the number of coefficients (α) needed to be estimated is very large, and this directly affects the statistical efficiency of $q_1(n)$ (see the discussion of the loss factor in Lavenberg and Welch 1981, or a generalization of this result given in Loh 1994). Therefore, we will pursue the use of $q_2(n)$ as an estimator of the mean steady-state system size.

Remark: A hybrid of the two estimators q_1 and q_2 could be constructed in which one first “clumps” the control variates into (say) $m < \ell(\ell + 1)/2$ variates, then performs the simulation, and finally estimates an m -vector of control coefficients to obtain the estimator. However, it is unclear how to perform the initial clumping, so we do not pursue such an approach here.

4 ESTIMATOR ASYMPTOTICS

We would like to investigate the asymptotic properties of the estimators outlined in the previous section. In particular, do they satisfy laws of large numbers and central limit theorems? To answer these questions, first observe that assumption **(A)** implies that Z satisfies a functional central limit theorem (FCLT).

Proposition 1 Let $M(n) \triangleq \sum_{k=0}^{n-1} Z(k)$. Under the assumption **(A)**,

$$n^{1/2} \left(\frac{M(n)}{n} - \bar{z} \right) \Rightarrow B_1(\cdot) \quad (7)$$

as $n \rightarrow \infty$, where $B_1(\cdot)$ is an ℓ^2 -dimensional Brownian motion with zero drift and positive semi-definite covariance matrix Λ , and \Rightarrow denotes weak convergence.

Proof: Let $V_2 = V^{2/5}$, where V is defined as in **(A)**. By concavity,

$$\begin{aligned} PV_2(z) &\leq (PV(z))^{2/5} \\ &\leq (V(z) - \|z\|^4 + \eta)^{2/5} \\ &\leq V_2(z) + (\eta - \|z\|^4)(V(z))^{-3/5} \\ &\leq V_2(z) - \epsilon \|z\| + \eta_2 \end{aligned}$$

for some ϵ and η_2 , where we have used the assumption that V is equivalent to a quintic in the last step.

It follows from this bound and **(A)** that the moment conditions of Meyn and Tweedie (1993) §17.5 are satisfied. Theorems 17.4.4 and 17.5.3 of Meyn and Tweedie (1993) complete the proof. \square

Remark: The FCLT may be viewed as an extension of the ordinary central limit theorem to an entire process. In particular, the ordinary central limit theorem results by looking at the FCLT result at time 1.

We may now describe the asymptotic behaviour of the standard estimator $q_0(n)$.

Theorem 1 Under the assumption **(A)**,

- $q_0(n) \rightarrow c^T \bar{z}$ a.s. as $n \rightarrow \infty$; and
- $n^{1/2}(q_0(n) - c^T \bar{z}) \Rightarrow \sigma_0 N(0, 1)$, as $n \rightarrow \infty$ where $\sigma_0^2 = c^T \Lambda c$ and $N(0, 1)$ denotes a standard normal random variable.

Proof: The first result is a consequence of the strong law of large numbers for Markov chains; see Meyn and Tweedie (1993), p. 424. The second result is a direct consequence of the FCLT and the fact that $q_0(n)$ is a linear function of $M(n)$. \square

In order to describe the asymptotic behaviour of $q_2(n)$, we first note that the FCLT (7) implies another FCLT of the form

$$n^{1/2} \left(\begin{array}{c} n^{-1} c^T M(n) - (c^T \bar{z}) \\ n^{-1} \alpha^T R^T M(n) - (\alpha^T \gamma) \end{array} \right) \Rightarrow B_2(\cdot), \quad (8)$$

as $n \rightarrow \infty$, where B_2 is a zero drift Brownian motion with covariance matrix

$$\Gamma = \left(\begin{array}{cc} c^T \Lambda c & c^T \Lambda R \alpha \\ c^T \Lambda R \alpha & \alpha^T R^T \Lambda R \alpha \end{array} \right).$$

Theorem 2 Suppose that assumption **(A)** holds, and that $\hat{\beta}(n) \rightarrow \beta$ a.s. as $n \rightarrow \infty$, where β is a deterministic constant. Then

- $q_2(n) \rightarrow c^T \bar{z}$ a.s. as $n \rightarrow \infty$; and
- $n^{1/2}(q_2(n) - c^T \bar{z}) \Rightarrow \sigma_2 N(0, 1)$ as $n \rightarrow \infty$, where $\sigma_2^2 = \Gamma_{11} + 2\beta \Gamma_{12} + \beta^2 \Gamma_{22}$.

Proof: The first result follows from the fact that $\bar{Z}_n \rightarrow \bar{z}$ a.s. as $n \rightarrow \infty$ and $\alpha^T R^T \bar{z} - \alpha^T \gamma = 0$ for any α . The second result is a consequence of the FCLT (8) and the converging together lemma (Billingsley 1986, p. 349). \square

Remark: To minimize the time-average variance constant (TAVC) σ_2^2 we should take $\beta^* = -\Gamma_{12}/\Gamma_{22}$ unless Γ_{22} is zero, in which case $\beta^* = 0$. Since Γ will rarely (if ever) be known exactly, β^* must be estimated. We will return to this point in §6.

5- CHOOSING α

In this section we show how to choose the parameters α so that the TAVC of $q_2(n)$ is minimized. Our goal in doing so is to provide some insight as to a “good” choice of α . Of course, the parameter β is also at our disposal, but merely scales α , so may be ignored (set to 1) in determining the optimal α . Because the optimal choice of α will prove to depend on the covariance structure of Z , it cannot, in general, be determined prior to the simulation.

Consider the estimator $q_2(n)$ with $\hat{\beta}(n) = 1$ for all n . From Theorem 2, we see that

$$\begin{aligned}\sigma_2^2 &= (c + R\alpha)^T \Lambda (c + R\alpha) \\ &= \|\Lambda^{1/2}(R\alpha + c)\|_2^2,\end{aligned}$$

where $\Lambda^{1/2}$ is a Cholesky factor of Λ (i.e., $\Lambda^{1/2}$ is upper triangular and $(\Lambda^{1/2})^T \Lambda^{1/2} = \Lambda$). Ideally, α would be chosen to minimize σ_2^2 , so that we should choose α to minimize the above 2-norm.

Prior to the simulation, both R and c are known, but Λ is not. In general then, we cannot choose the optimal α 's prior to the simulation. However, we now have some insight into how to choose α . A reasonable approach is to attempt to minimize $\|R\alpha + c\|$ for some norm, and then use the estimator $q_2(n)$ defined in (6). A reasonable choice of norm is the 2-norm, because this will yield the optimal α when Λ is the identity matrix, and intuitively should yield good results if Λ is “close” to the identity. However, it is conceivable that other norms may also prove useful.

We may also exploit prior knowledge of Λ . For example, consider the class of pre-emptive priority policies. In these policies, the buffers at a station are ordered according to priority, and a job in one of the buffers receives service only if the buffers of higher priority at that station are empty. Therefore, if $\sigma(i) = \sigma(j)$, and j has a higher priority than i , then $Z_{ij} = W_i X_j = 0$ (Kumar and Kumar 1994a). We immediately see that the column of $\Lambda^{1/2}$ corresponding to the variable Z_{ij} must be zero, so that the (i, j) element of $R\alpha + c$ corresponding to that column has no effect on the TAVC. Hence, we minimize $\|R\alpha + c\|$ but exclude the (i, j) element from the norm calculation. It is conceivable that similar methods might be used to exploit knowledge about other policies.

The problem of minimizing the 1 or ∞ norm of $R\alpha + c$ may be formulated as a linear program. By appealing to the normal equations, one may also minimize the 2-norm using linear programming, or of course, solve this least squares problem by some other method. Thus, our approach is as follows.

1. Solve a linear program or least squares problem to obtain α .

2. Simulate the Markov chain Z up to time n .
3. Compute some reasonable $\hat{\beta}(n)$.
4. Compute the estimator $q_2(n)$ as given in (6).

Remark: The minimization problem above may be viewed as approximating the solution to Poisson's equation (Henderson 1997). For positive recurrent discrete-time processes, Poisson's equation may be written

$$Ph(z) - h(z) = -(c^T z - c^T \bar{z}) \quad \forall z \in S. \psi \quad (9)$$

If this equation could be solved, then we would obtain a zero variance estimator for $c^T \bar{z}$ via $c^T Z(0) + Ph(Z(0)) - h(Z(0))$. We may view the problem of minimizing $\|R\alpha + c\|$ as finding a function

$$h_\alpha \triangleq \sum_{i,j} \alpha_{ij} h_{ij}$$

that approximately solves (9). To see this, observe that

$$Ph_\alpha(z) - h_\alpha(z) \approx \alpha^T (R^T z - \gamma),$$

so that if $R\alpha = -c$, then h_α approximately solves (9).

6- BATCH MEANS

Recall that it is desirable to estimate $\beta^* = -\Gamma_{12}/\Gamma_{22}$ when using the estimator $q_2(n)$. One approach to doing so would be to use the regenerative method (see Law and Kelton 1991), with regeneration times defined by the return times to a distinguished state. However, it is conceivable that for large systems, the time between returns to a distinguished state would be prohibitive, so that very large run lengths would be required to obtain even a single regenerative cycle.

Therefore, we choose to use the method of batch means for estimating β^* . Following the recommendations of Nelson (1989) on combining batch means and control variates, we used a fixed number of batches $b = 30$. However, in doing so we no longer obtain a consistent estimator of β^* , so that the asymptotic results of §4 may not apply (and in fact, they do not). In order to be precise, we first outline how the estimator $q_2(n)$ is constructed, which basically comes down to how $\hat{\beta}(n)$ is constructed, and then provide the relevant asymptotic results. The following computational procedure is adapted from Loh (1994), p. 33.

1. Select α .
2. Simulate Z up to time n and compute (for $i = 1, \dots, b$) the batch means

$$\bar{Z}^i = \frac{b}{n} \sum_{k=\lceil n(i-1)/b \rceil}^{\lceil ni/b - 1 \rceil} Z(k).$$

3. Compute $\bar{Z}(n) = b^{-1} \sum_{i=1}^b \bar{Z}^i$.

4. Compute

$$\begin{aligned} V_{11}(n) &= (b-1)^{-1} \sum_{i=1}^b (c^T \bar{Z}^i - c^T \bar{Z}(n))^2, \\ V_{22}(n) &= b^{-1} \sum_{i=1}^b (\alpha^T R^T \bar{Z}^i - \alpha^T R^T \bar{Z}(n))^2, \\ V_{12}(n) &= (b-1)^{-1} \sum_{i=1}^b (c^T \bar{Z}^i - c^T \bar{Z}(n)) \\ &\quad \times (\alpha^T R^T \bar{Z}^i - \alpha^T R^T \bar{Z}(n)), \\ V^2(n) &= \frac{b-1}{b-2} (V_{11}(n) - \frac{V_{12}(n)^2}{V_{22}(n)}) \text{ and} \\ S^2(n) &= V^2(n) \left(\frac{1}{b} + \frac{(\alpha^T R^T \bar{Z}(n) - \alpha^T \gamma)^2}{b-1} \right). \end{aligned}$$

5. Set $\hat{\beta}(n) = -V_{12}(n)/V_{22}(n)$.

6. Set $q_2(n) = c^T \bar{Z}(n) + \hat{\beta}(n)(\alpha^T R^T \bar{Z}(n) - \alpha^T \gamma)$.

Theorem 3 *If $q_2(n)$ is computed according to the computational procedure outlined above, then under the assumption (A),*

1. $q_2(n)$ converges in probability to $c^T \bar{z}$ as $n \rightarrow \infty$; and
2. $\frac{q_2(n) - c^T \bar{z}}{S(n)} \Rightarrow t_{b-2}$, as $n \rightarrow \infty$, where t_{b-2} is a random variable with the Student's t distribution with $b-2$ degrees of freedom.

Proof: The second result is proved in §1.3.2 of Loh (1994). In addition, Proposition 1.5 of Loh (1994) shows that $nS^2(n)$ converges in distribution to a finite-valued random variable as $n \rightarrow \infty$, so the first result follows. \square

Suppose that $P(t_{b-2} < t^*) = 0.95$. Since the constant σ_0^2 of Theorem 1 may be estimated by $nV_{11}(n)/(b-1)$, approximate 95% C.I.'s for $c^T \bar{z}$ are given by

$$q_0(n) \pm t^* \sqrt{\frac{V_{11}(n)}{b-1}} \text{ and } q_2(n) \pm t^* S(n). \psi \quad (10)$$

7- EXAMPLES

In this section we present three examples that provide a comparison between the standard estimator $q_0(n)$ and the controlled estimator $q_2(n)$. We will base the comparisons on the expected halfwidths of the C.I.'s given in (10).

The major component of the additional work required to use the new estimator $q_2(n)$ is the solution

of the minimization problem prior to the simulation. For the examples presented below we minimized the 2-norm problem. In all cases, the minimization took less than 20% of the time required for each of the simulation runs. Furthermore, in practice one would run the simulations for a longer period than was feasible in collecting these numerical results, so that the minimization would be an even smaller portion of the overall computational time. This observation and the reductions in halfwidth demonstrated below make the controlled estimator very appealing.

Example 1 The symmetric tandem queue may be considered to be the simplest non-trivial re-entrant line. Jobs arrive to the first station at rate λ , are served at rate μ , move to the second station, are again served at rate μ , and then depart from the system.

We ran simulations for several utilization levels to see how effective the variance reductions are over various system loads. For each utilization level, we ran 100 simulations, calculating $q_0(n)$ and $q_2(n)$, and estimates of their C.I. halfwidths. The C.I. halfwidths were estimated using the batch means method with 30 batches, where each batch consisted of 10000 simulated time units. In an attempt to remove initialization bias, the first batch of each simulation was discarded.

The results are presented in Table 1. The first column contains the traffic intensity $\rho = \lambda/\mu$, (where $\lambda + 2\mu = 1$), and the next two columns are 95% C.I.'s for the expected halfwidths of $q_0(n)$ and $q_2(n)$. For example, for $\rho = 0.5$, the expected C.I. halfwidths are 0.046 and 0.015. This represents variance reduction by a factor of $(0.046/0.015)^2 \approx 9$ (note that the halfwidths are proportional to the square-roots of the variances).

Table 1: Estimates of CI Halfwidths for Example 1

$\rho\psi$	Standard-	Controlled
1/3-	0.0185 \pm 0.0005-	0.0045 \pm 0.0001
1/2-	0.046 \pm 0.001-	0.015 \pm 0.0005
2/3-	0.134 \pm 0.003-	0.057 \pm 0.002
9/10-	1.80 \pm 0.07-	1.00 \pm 0.05
19/20-	5.9 \pm 0.3	3.7 \pm 0.3

Notice that the method provides useful variance reductions for all traffic intensities ρ . In heavy traffic, the difference between the two methods appears to decrease. This is perhaps disappointing, considering that heavy traffic is exactly the regime that causes estimation difficulties, owing to large variances (Whitt 1989, Asmussen 1992).

Example 2 The second example that we consider is a system consisting of 2 stations and 3 buffers. Jobs

arrive to buffer 1 at station 1, then proceed to buffer 2 at station 2, and complete their excursions by revisiting station 1 at buffer 3. We assume that this system operates under the LBFS (last buffer first served) preemptive priority policy, which is known to be stable as long as $\lambda < \mu_2, \mu_1 + \mu_3$ (Kumar and Kumar 1994b). We took $\mu_1 = \mu_3$.

We performed the same experiments as in the previous example. The results are presented in Table 2. The first 3 columns are the unnormalized rates $\lambda, \mu_1 = \mu_3$ and μ_2 . The final 2 columns are 95% C.I.'s for the expected half-width of the standard and new estimators. The traffic intensity at both stations increases with the rows of the table.

Once again the variance reductions are of a very useful size, but decrease in effectiveness as the system enters heavy traffic.

Table 2: Estimates of CI Halfwidths for Example 2

λ	μ_1	μ_2	Standard	Controlled
1	4	3	0.042 ± 0.001	0.00229 ± 0.00006
1	3	3	0.096 ± 0.003	0.0043 ± 0.0001
2	5	3	0.34 ± 0.01	0.0272 ± 0.0008
9	20	11	1.32 ± 0.05	0.84 ± 0.03

Example 3 Our final example demonstrates that the methodology is effective for larger systems. The system consists of 10 stations and 59 buffers as shown in Figure 1. Every station operates under the LBFS preemptive priority policy. As before, the rates are scaled so that the sum of the service rates and the arrival rates is 1. The unscaled arrival rate is $\lambda = 1$, and the unscaled service rates are identical at each station, and equal to 16, 16, 16, 18, 18, 20, 20, 20, 20 and 20. C.I.'s for the standard and controlled mean C.I. halfwidths are 0.68 ± 0.02 and 0.32 ± 0.03 respectively. Thus we obtain a factor of approximately 4 in variance reduction.

To conclude this section we remark that we could also have minimized the 1 or ∞ norm of $R\alpha - c$, instead of the 2-norm as above. For the examples we tried (including those presented here), similar variance reductions were achieved.

8- SUMMARY AND DISCUSSION

The methodology presented in this paper generates effective control variates for simulations of re-entrant lines. Our control variate approach appears to be quite effective for all traffic loadings, but is more-so for moderate to light loadings. Although we have restricted attention to the mean steady-state number




Figure 1: The Re-Entrant Line of Example 3.

of jobs in the system, we could also apply the method to any linear function of \bar{z} . In particular, we could use the same approach in estimating the mean steady-state number of jobs in a single buffer.

The extension of this methodology to more general multiclass queueing systems (such as those considered in Bertsimas et al. 1994) is straight-forward. Simply uniformize the process Z , and consider functions of the form $Ph - h$, where h is a quadratic function of the buffer sizes, exactly as in §3.

It is of interest to see if it is possible to construct an estimator that performs far better than the standard estimator in heavy traffic, which is traditionally the most difficult performance regime in which to simulate. While the estimator constructed in this paper realized useful variance reductions, the reductions in heavy traffic were not as substantial as (for instance) those seen for the Approximating Markov Process method applied to the single-server queue (Henderson 1997). This remains an open problem.

ACKNOWLEDGMENTS

We thank Steve Chick for helpful suggestions that improved the presentation of this paper. Work supported in part by NSF Grant ECS 940372, 1994 University of Illinois Research Board Grant *Beckman* 1-6-49749, and JSEP grant N00014-90-J-1270. This work was completed with the assistance of equipment granted through the IBM Shared University Research program and managed by the Computing and Communications Services Office at the University of Illinois at Urbana-Champaign.

REFERENCES

- Andradóttir, S., D. P. Heyman, and T. J. Ott. 1993. Variance reduction through smoothing and control variates for Markov chain simulations. *ACM Transactions on Modeling and Computer Simulation* 3:167–189.
- Asmussen, S. 1992. Queueing simulation in heavy traffic. *Mathematics of Operations Research* 17:84–111.
- Baskett, F., K. M. Chandy, R. R. Muntz, and F. G. Palacios. 1975. Open, closed, and mixed networks of queues with different classes of customers. *Journal of the ACM* 22:248–260.
- Bertsekas, D., and R. Gallager. 1987. *Data networks*. Englewood Cliffs, NJ: Prentice Hall.
- Bertsimas, D., I. Ch. Paschalidis, and J. N. Tsitsiklis. 1994. Optimization of multiclass queueing networks: polyhedral and nonlinear characterizations of achievable performance. *Annals of Applied Probability* 4:43–75.
- Billingsley, P. 1986. *Probability and measure, 2nd ed.* Wiley, New York.
- Bramson, M. 1994. Instability of FIFO queueing networks. *Annals of Applied Probability* 4:414–431.
- Bratley, P., B. L. Fox, and E. L. Schrage. 1987. *A guide to simulation, 2nd Ed.* Springer-Verlag.
- Buzacott J. A., and J. G. Shanthikumar. 1993. *Stochastic models of manufacturing systems*. Prentice Hall.
- Chen, H. 1995. Fluid approximations and stability of multiclass queueing networks: work-conserving disciplines. *Annals of Applied Probability* 5:637–665.
- Dai, J. G. 1995. On positive Harris recurrence of multiclass queueing networks: a unified approach via fluid limit models. *Annals of Applied Probability* 5:49–77.
- Gershwin S. B. 1993. *Manufacturing systems engineering*. Englewood Cliffs, NJ: Prentice Hall.
- Henderson S. G. 1997. *Variance reduction via an approximating Markov process*. Ph.D. Thesis, Department of Operations Research, Stanford University, CA.
<http://www-or.stanford.edu/~simlib/theses.html>
- Kelly, F. P. 1979. *Reversibility and stochastic networks*. New York, NY: John Wiley and Sons.
- Kumar, S., and P. R. Kumar. 1994a. Performance bounds for queueing networks and scheduling policies. *IEEE Transactions on Automatic Control* 39:1600–1611.
- Kumar, S., and P. R. Kumar. 1994b. Fluctuation smoothing policies are stable for stochastic reentrant lines. *33rd IEEE Proceedings Conference on Decision and Control* 1476–1480.
- Kumar, P. R., and S. P. Meyn. 1995. Stability of queueing networks and scheduling policies. *IEEE Transactions on Automatic Control* 40:251–260.
- Kumar, P. R., and S. P. Meyn 1996. Duality and linear programs for stability and performance analysis of queueing networks and scheduling policies. *IEEE Transactions on Automatic Control* 41:4–17.
- Lavenberg, S. 1983. Editor. *Computer performance modeling handbook*. New York: Academic Press.
- Lavenberg, S. S., and P. D. Welch. 1981. A perspective on the use of control variables to increase the efficiency of Monte Carlo simulations. *Management Science* 27:322–335.
- Law, A. M., and W. D. Kelton. 1991. *Simulation modeling and analysis, 2nd ed.* McGraw-Hill.
- Lippman, S. 1975. Applying a new device in the optimization of exponential queueing systems. *Operations Research* 23:687–710.
- Loh, W. W. 1994. *On the method of control variates*. Ph.D. Thesis. Department of Operations Research, Stanford University, Stanford CA.
<http://www-or.stanford.edu/~simlib/theses.html>
- Meyn, S. P., and R. L. Tweedie. 1993. *Markov chains and stochastic stability*. Springer-Verlag.
- Nelson, B. L. 1989. Batch size effects on the efficiency of control variates in simulation. *European Journal of Operational Research* 43:184–196.
- Ou, J., and L. M. Wein. 1992. Performance bounds for scheduling queueing networks. *Annals of Applied Probability* 2:460–480.
- Rybko, A. N., and A. L. Stolyar. 1993. On the ergodicity of random processes that describe the functioning of open queueing networks. *Problems of Information Transmission* 28:199–220.
- Whitt, W. 1989. Planning queueing simulations. *Management Science* 35:1341–1366.

AUTHOR BIOGRAPHIES

SHANE G. HENDERSON is an Assistant Professor in the Department of Industrial and Operations Engineering at the University of Michigan, Ann Arbor. He received a Ph.D. in Operations Research from Stanford University in 1997. His research interests include discrete-event simulation and queueing systems.

SEAN P. MEYN is an Associate Professor in the Department of Electrical and Computer Engineering, and a Research Associate Professor in the Coordinated Science Laboratory at the University of Illinois, Urbana-Champaign. His research interests include queueing networks, Markov chains, and systems and control theory.