

Variance Reduction in Simulation of Multiclass Processing Networks

Shane G. Henderson*

School of Operations Research and Industrial Engineering, Cornell University
Ithaca, NY 14850

Sean P. Meyn†

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

May 5, 2005

Abstract

We use simulation to estimate the steady-state performance of a stable multiclass queueing network. Standard estimators have been seen to perform poorly when the network is heavily loaded. We introduce two new simulation estimators. The first provides substantial variance reductions in moderately-loaded networks at very little additional computational cost. The second estimator provides substantial variance reductions in heavy traffic, again for a small additional computational cost. Both methods employ the variance reduction method of control variates, and differ in terms of how the control variates are constructed.

*Work supported in part by NSF grant DMI-0224884

†Work supported in part by NSF grants DMI-0224884 and ECS 940372

1 Introduction

A *multiclass* queueing network is a network of service stations through which multiple classes of customers move. Each customer class can have different service-time characteristics at a single service station. Multiclass queueing networks are of great interest in a large variety of applications (Bertsekas and Gallager, 1987; Lavenberg, 1983; Buzacott and Shanthikumar, 1993; Gershwin, 1993) because of their tremendous modeling flexibility. Perhaps the most common reason for modeling a system using a multiclass queueing network is to try to determine a suitable operating policy for the network. An *operating policy* is a policy that determines which customers should be worked on at which times. For example, if there are multiple customer classes at a single service station, then which class should the station work on?

In order to make comparisons between operating policies, one must define a suitable performance measure, such as expected steady-state work in process, or expected steady-state throughput, etc. For broad classes of networks one can compute certain performance measures analytically (Jackson, 1963; Baskett et al., 1975; Kelly, 1979; Harrison and Williams, 1987, 1990), or one can turn to numerical computation (Schweitzer, 1984; Neuts, 1994; Dai and Harrison, 1992; Shen et al., 2002). In general, however, these approaches are either infeasible, or intractable due to the high complexity of network models. Some results are available in the form of bounds on performance measures through the construction of linear programs (Kumar and Kumar, 1994; Bertsimas et al., 1994; Kumar and Meyn, 1996; Schwerer, 2001; Morrison and Kumar, 1999). Unfortunately, these bounds are often quite loose, and so it can be difficult to compare operating policies based on such bounds alone. It is natural then, to turn to simulation.

Given that one is going to simulate many different operating policies, it is important that any simulation return relatively accurate answers as quickly as possible. This suggests the need for variance-reduction techniques that can increase the accuracy in simulation results for a given computational budget. Another reason for desiring efficient simulation techniques is that the network under consideration is often moderately to heavily loaded, in the sense that some of the resources of the network are close to full utilization. It has

been noted that, in such settings, simulation can take a tremendously long time to return precise estimates of performance (Whitt, 1989; Asmussen, 1992).

So we are strongly motivated to seek special variance-reduction techniques for multi-class networks. In this paper we develop two such variance-reduction techniques. Both are based on the approximating martingale process method (Henderson and Glynn, 2002; Henderson, 1997), which is a specialization of the method of control variates; see, for example, Law and Kelton (2000) for an introduction to control variates. This paper is an outgrowth of Henderson and Meyn (1997). The methods introduced there and here have since seen further development in Henderson et al. (2003) and Borkar and Meyn (2003). The theoretical results on the order of the variance constants given here have been considerably extended in Meyn (2005,b). These papers also describe further insights on network behavior that have been uncovered since the work presented in this paper was completed. Furthermore, Kim and Henderson (2004) and Kim and Henderson (2005) have since introduced a new family of variance reduction techniques that may lead to even greater variance reductions than those seen in this paper.

Although our presentation concentrates on the estimation of the mean steady-state number of customers (of all classes) in the system, our methods may be tailored to the steady-state estimation of any linear function of the individual customer-class populations. In particular, we can also estimate, for instance, the mean steady-state number of customers of a particular class present in the system.

In Section 2 we describe our model of a multiclass-queueing system. We also review *Poisson's equation* and explain its importance in our context. In particular, we wish to approximate the solution to Poisson's equation to construct efficient simulators.

In Section 3, we explore quadratic forms as approximations to the solution to Poisson's equation. Computational results are given for the resulting simulation estimator, which we term the *quadratic* estimator. The quadratic estimator significantly outperforms a more standard simulation estimator in lightly to moderately loaded networks. In heavily loaded networks, the difference between the performance of the two estimators closes, although the quadratic estimator still provides variance reductions that would significantly reduce the computational effort involved in exploring a class of operating policies.

In Section 4, we explore alternative approximations to the solution to Poisson's equation based on the concept of a fluid limit. The resulting simulation estimator, the *fluid* estimator, yields significant variance reductions in heavily loaded networks, and modest variance reductions in less heavily loaded networks. For a given simulation run-length, it is slightly more expensive to compute than the standard estimator, so that the issue of variance reduction versus computational effort needs to be considered (Glynn and Whitt, 1992). We discuss the choice of simulation estimator for a given network in Section 5.

2 Multiclass Queueing Networks

Consider a system consisting of d stations (or machines) and ℓ classes of customers (or jobs). Class i customers require service at station $s(i)$. Upon completion of service at station $s(i)$, a class i customer becomes a class j customer with probability R_{ij} , and exits the system with probability

$$R_{i0} \triangleq 1 - \sum_{j=1}^{\ell} R_{ij}.$$

The service times for class i customers are assumed to form an i.i.d. (independent and identically distributed) sequence of exponentially distributed r.v.'s (random variables) with mean μ_i^{-1} . Class i customers arrive exogenously to station $s(i)$ according to a Poisson process with rate λ_i (which may be zero). The Poisson-arrival processes and the service-time processes are mutually independent. We let μ denote the ℓ -dimensional vector of service rates, and λ the d -dimensional vector of arrival rates. Unless otherwise stated, all vectors are assumed to be column vectors.

We require that the *routing matrix* $R = (R_{ij} : 1 \leq i, j \leq \ell)$ be transient, so that the following inverse exists:

$$(I - R)^{-1} = \sum_{k=0}^{\infty} R^k.$$

This ensures that all customers that enter the system will eventually leave, and we can also be assured that there is a unique solution $\gamma \geq 0$ to the traffic equations,

$$0 = \lambda - \gamma + R'\gamma, \tag{1}$$




Figure 1: A multiclass network with $d = 2$ stations and $\ell = 3$ customer classes.

where R' denotes the transpose of the matrix R . We assume throughout that $\gamma_i > 0$ for all i .

Figure 1 illustrates an example with $d = 2$ stations, $\ell = 3$ customer classes, and a single exogenous arrival process (two of the arrival rates are zero). Customer classes 1 and 3 are served at Station 1, so that $s(1) = s(3) = 1$. Similarly $s(2) = 2$. The routing values R_{ij} are zero except for $R_{12} = R_{23} = 1$.

Let $X_i(t)$ denote the number of class i customers present in the system at time t , and let $X(t) = (X_i(t) : 1 \leq i \leq \ell)$ be the vector of customer populations at time t . Let $V_i(t)$ denote the fraction of station $s(i)$'s effort allocated to serving customers of class i at time t , and let $V(t)$ denote the corresponding vector quantity. We must have $V_i(t) \geq 0$ for all i and t , and $\sum_{i:s(i)=s} V_i(t) \leq 1$ for all s and t , i.e., a station s can never allocate negative effort, or more than unit effort, to the classes $\{i : s(i) = s\}$ that are served at the station. For the network in Figure 1 for example, suppose that at time t Machine 1 is serving Class-3 jobs, and Machine 2 is empty. Then $V(t) = (0, 0, 1)'$.

We require the operating policy adopted by stations to be *stationary*, *non-idling*, and $0 - 1$. We next define and explain each of these terms.

By *stationary*, we mean that $V(t)$ is a deterministic function of $X(t)$, so that the workload allocations depend only on the current customer-class levels. This allows the modeling of preemptive priority policies, for example, but precludes the modeling of policies such as FIFO which rely on additional information such as the order in which customers arrive to a station.

By non-idling, we mean that if $\sum_{i:s(i)=s} X_i(t) > 0$, then $\sum_{i:s(i)=s} V_i(t) = 1$, i.e., if there are customers present at station s at time t , then the station allocates all of its effort at that time.

By 0 – 1, we mean that at any given station, at most one class receives service at any given time. This assumption is applied *only* for notational convenience. The estimators we derive may also be applied to networks controlled by randomized or processor-sharing policies.

With the above structure in place, we may conclude that $X = (X(t) : t \geq 0)$ is a time-homogeneous continuous-time Markov chain. However, we prefer to work in discrete time because it simplifies the analysis. So we first rescale time so that $e'(\lambda + \mu) = 1$, where e denotes a vector of ones. Then we uniformize; see, e.g., Ross (1996, p. 282). Uniformization is a process that allows us to study the continuous-time process X through a related discrete-time process $Y = (Y(n) : n \geq 0)$. Define $\tau_0 = 0$ and let the times $\{\tau_n : n \geq 1\}$ correspond to epochs when either arrivals, real service completions, or virtual service completions occur in the uniformized process. For $n \geq 0$, let $Y(n) = X(\tau_n)$ and $W(n) = V(\tau_n)$. Then the process $Y = (Y(n) : n \geq 0)$ is a discrete-time Markov chain evolving on a (countable) state space S that is a subset of $\{0, 1, 2, \dots\}^\ell$.

Example 1: For the M/M/1 queue, Y is a Markov chain on $\{0, 1, 2, \dots\}$ with transition matrix P , where

$$P_{ij} = (\mu + \lambda)^{-1} \begin{cases} \lambda & \text{if } j = i + 1, \\ \mu & \text{if } j = \max(i - 1, 0), \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

■

Our goal is to estimate α , the steady-state mean number of customers in the system, i.e., the steady-state mean of $|Y(0)| \triangleq e'Y(0)$ (i.e. $|\cdot|$ is the L_1 norm). To ensure (among other things) that α exists and is finite, we make a certain assumption **(A)** below. The assumption **(A)** is known as a Lyapunov, or Foster-Lyapunov, condition. Intuitively, the function V represents energy, and (3) indicates that there is an expected loss in energy for states y that are “large.” This then ensures that energy never gets too large, and so

the chain remains stable.

For any function V on \mathbb{R}_+^ℓ we define

$$PV(y) \triangleq \mathbb{E}_y V(Y(1)), \quad \Delta_V(y) = PV(y) - V(y), \quad y \in S,$$

where $\mathbb{E}_y(\cdot) \triangleq \mathbb{E}(\cdot | Y(0) = y)$. Intuitively, $PV(y)$ ($\Delta_V(y)$) represents the expected energy (expected change in energy) one step from now, assuming that the chain is currently in state y .

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

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

$$\delta(|y|^5 + 1) \leq V(y) \leq \delta^{-1}(|y|^5 + 1); \text{ and} \tag{2}$$

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

$$\Delta_V(y) = PV(y) - V(y) \leq -|y|^4 + \eta. \tag{3}$$

Continuation of Example 1: For the M/M/1 queue, we may take $V(y) = b_0 y^5$, with b_0 a sufficiently large constant, and then **(A)** is satisfied as long as $\rho \triangleq \lambda/\mu < 1$. \blacksquare

In general, this condition will be satisfied if the stability-linear program of Kumar and Meyn (1996) admits a solution, generating a co-positive $\ell \times \ell$ matrix Q . In this case, the function V may be taken as $V(y) = (y'Qy)^{5/2}$.

Alternatively, if a fluid model is stable, and we define $\kappa = \inf\{n \geq 1 : Y(n) = 0\}$, then the function

$$V(y) = \mathbb{E}_y \sum_{k=0}^{\kappa} |Y(k)|^4$$

is bounded as in Condition 1 (Dai and Meyn, 1995), and this function is known to satisfy Condition 2 (Meyn and Tweedie, 1993, p. 338).

Under the assumption **(A)**, the chain Y possesses a unique stationary distribution π , and $\mathbb{E}_\pi |Y(0)|^4 \leq \eta < \infty$ (Meyn and Tweedie, 1993, p. 330), where

$$\mathbb{E}_\pi(\cdot) \triangleq \int_S \mathbb{E}(\cdot | Y(0) = y) \pi(dy).$$

Hence, in particular, the steady-state mean number of customers $\alpha = \mathbb{E}_\pi |Y(0)| < \infty$.

The Lyapunov condition **(A)** is stronger than is strictly necessary to ensure that $\alpha = \mathbb{E}_\pi |Y(0)|$ is finite. A “tighter” requirement is that there is a function $h : \mathbb{R}_+^\ell \rightarrow \mathbb{R}_+$ satisfying

$$\Delta_h(y) = Ph(y) - h(y) \leq -|y| + \eta \quad (4)$$

for all $y \in S$ and some constant $\eta > 0$. Given such a function, Theorem 14.3.7 of Meyn and Tweedie (1993) allows us to conclude that $\alpha = E_\pi |Y(0)| \leq \eta$.

One might then ask whether the inequality in (4) can be made an equality, thereby yielding a tighter upper bound η on α . In such a case we would have

$$\Delta_{h^*}(y) = Ph^*(y) - h^*(y) = -|y| + \eta. \quad (5)$$

This equation is known as *Poisson’s equation*.

If h^* is a solution to Poisson’s equation, then it is easy to see that $h^* + c$ is also a solution for any constant c . In fact, Proposition 17.4.1 of Meyn and Tweedie (1993) shows that any two π -integrable solutions to Poisson’s equation *must* differ by an additive constant, in the sense that for all y in a set A with $\pi(A) = 1$,

$$h_1(y) - \pi(h_1) = h_2(y) - \pi(h_2)$$

where, for a real-valued function $g : S \rightarrow \mathbb{R}$, we denote $\pi(g) = \int_S g(y)\pi(dy)$.

One may estimate α using $\alpha(n) \triangleq |\bar{Y}(n)|$, where $\bar{Y}(n) \triangleq n^{-1} \sum_{i=0}^{n-1} Y(i)$, the mean number of customers in the system up to time n . The following result is a special case of Theorem 17.0.1 of Meyn and Tweedie (1993), and shows that the estimator $\alpha(n)$ is consistent, and satisfies a central limit theorem.

Theorem 1 *Suppose that **(A)** holds. Then $\alpha(n) \rightarrow \alpha$ almost surely (a.s.), and furthermore,*

$$n^{1/2}(\alpha(n) - \alpha) \Rightarrow \sigma N(0, 1),$$

as $n \rightarrow \infty$. The time-average variance constant (TAVC) σ^2 is given by $\sigma^2 = \mathbb{E}[h(Y)(|Y| - \mathbb{E}|Y|)] - \text{Var}|Y|$, where Y is distributed according to the stationary distribution π , and h solves Poisson’s equation; see (5).

As noted in the introduction, it has been observed that simulation can take a very long time to yield accurate answers for heavily loaded networks (Asmussen, 1992; Whitt, 1989). This problem is exhibited in our framework through the TAVC. Our simulation experiments indicate that the TAVC grows rapidly as the network becomes heavily loaded. Our next result lends further weight to these observations.

Consider a multiclass queueing system consisting of a single station and (possibly) multiple customer classes. Suppose that the service rates μ and arrival rates λ are such that

$$e' \mathcal{M}^{-1} (I - R')^{-1} \lambda = 1,$$

where $\mathcal{M} = \text{diag}(\mu)$. (This corresponds to a situation where the resources of the network are exactly matched by the demand.) Now consider a family of queueing systems indexed by $\rho \in (0, 1)$, where the ρ th system has arrival rate vector $\lambda(\rho) = \rho\lambda$. For the sake of clarity, we occasionally suppress dependence on ρ .

For a given vector of buffer levels y , let $f(y) = d'y$ be a measure of the work in the system, where

$$d' = e' \mathcal{M}^{-1} (I - R')^{-1} = e' Q,$$

and $Q = \mathcal{M}^{-1} (I - R')^{-1}$. Intuitively, $f(y)$ measures that total expected amount of processing required to completely serve all of the customers presently in the system as given by y .

Let $\sigma_f^2(\rho)$ be the TAVC associated with the estimator $f(\bar{Y}(n)) = d'\bar{Y}(n)$, and let $\sigma_i^2(\rho)$ be the TAVC associated with $\bar{Y}_i(n)$ ($i = 1, \dots, \ell$). If Σ_ρ denotes the time-average covariance matrix of Y , then $\sigma_f^2(\rho) = d'\Sigma_\rho d$ and $\sigma_i^2(\rho) = e'_i \Sigma_\rho e_i$, where e_i denotes the i th basis vector. The proof of the following result may be found in the appendix.

Theorem 2 *Consider the family of multiclass queueing systems above under any non-idling work-allocation policy. Then the following are true for each $\rho < 1$.*

1. *Assumption **(A)** holds with the function $b_\rho f^5$ for some sufficiently large constant b_ρ , and the TAVCs $\sigma_f^2(\rho)$ and $\sigma_i^2(\rho)$ ($i = 1, \dots, \ell$) are finite.*

2. The solution to Poisson's equation for the estimator $f(\bar{Y}(n))$ is given by

$$h(y; \rho) = \frac{f^2(y)}{2(1-\rho)} + c'_\rho y,$$

where the vector c_ρ is of the order $(1-\rho)^{-1}$.

Furthermore, there exist constants $A, B > 0$ (independent of ρ) such that for ρ sufficiently close to 1,

$$\frac{A}{(1-\rho)^4} \leq \sigma_f^2(\rho) \leq \frac{B}{(1-\rho)^4},$$

and finally

$$\text{trace } \Sigma_\rho = \sum_{i=1}^{\ell} \sigma_i^2(\rho) \geq \frac{A'}{(1-\rho)^4}$$

for some constant A' (again, independent of ρ).

Theorem 2 shows that $\text{trace } \Sigma_\rho$ is of the order $(1-\rho)^{-4}$ as $\rho \rightarrow 1$, and this suggests, although it does not necessarily prove, that the TAVC $e'\Sigma_\rho e$ of the standard estimator $\alpha(n)$ is of the same order (see also Meyn (2005) where this precise order is verified for the TAVC in a diffusion model.) We have already mentioned that our simulation experiments also indicate that $\alpha(n)$ has high variance in heavily congested networks. Therefore, there is strong motivation for identifying alternative estimators to $\alpha(n)$ that can improve performance in heavy traffic. The key to these estimators is Poisson's equation.

If h^* is π -integrable, then by taking expectations with respect to π in (5), we see that $\alpha = \eta$. Therefore, if the solution to Poisson's equation is known, then so is α . In general then, we cannot expect to know the solution to Poisson's equation. But what if an *approximation* is known?

If the approximation h is π -integrable, then $\pi(\Delta_h) = 0$, i.e., $\Delta_h(Y_n)$ has steady-state mean 0, and so one might consider using Δ_h in building a simulation control variate for estimating α . In particular, one might consider using the controlled estimator

$$\alpha_c(n) = \alpha(n) + \frac{\beta}{n} \sum_{i=0}^{n-1} \Delta_h(Y(i)), \quad (6)$$

where β is an adjustable constant.

If $h = h^*$ and we take $\beta = 1$, then $\alpha_c(n) = \alpha$, and we obtain a zero-variance estimator of α . In general, we can expect useful variance reductions using the estimator $\alpha_c(n)$ provided that h is a suitable approximation to the solution to Poisson's equation. We will discuss the choice of the constant β later. For more details on this approach to variance reduction see Henderson and Glynn (2002); Henderson (1997); Henderson et al. (2003); Borkar and Meyn (2003) and Meyn (2005).

So how should one go about determining an approximation to the solution h^* to Poisson's equation? It is known that for any 'reasonable' policy, the function h^* is equivalent to a quadratic, in the sense of (2) (see Kumar and Meyn (1996); Meyn (1997, 2001, 2005) and Theorem 3 below). So it is reasonable to search for a quadratic function h that approximately solves (5).

3 A Quadratic Approximation

We begin this section by demonstrating the general ideas of the approach on the stable M/M/1 queue.

Continuation of Example 1: Recall that the solution to Poisson's equation is equivalent to a quadratic. So it is reasonable to approximate the solution h^* to Poisson's equation by a quadratic function.

In the linear case $h(y) = y$, we have that $\Delta_h(y) = \lambda - \mu w$, where $w = \mathbb{I}(y > 0)$ ($\mathbb{I}(\cdot)$ is the indicator function that is 1 if its argument is true and 0 otherwise) represents the only non-idling policy: The server works when customers are present, and idles when customers are not present. Taking expectations with respect to π , we see that the expected fraction of time that the server is working is $\mathbb{E}_\pi W(0) = \lambda/\mu$, and this is the result expected by work conservation.

Taking the pure quadratic $h(y) = ay^2$, we have

$$\Delta_h(y) = 2ay(\lambda - \mu) + a\lambda + a\mu w. \quad (7)$$

We now choose $a = (\mu - \lambda)^{-1}/2$ to ensure that the coefficient of y is -1 . We could then use the right-hand side of (7) as a control variate as in (6). However, it is instructive (and

useful) to adopt a slightly different approach where we avoid estimation of the variable w , and replace it by its known steady-state mean $\mathbb{E}_\pi W(0) = \pi(w) = \lambda/\mu$. We may then use the controlled estimator $\bar{Y}(n) + \beta(-\bar{Y}(n) + 2\lambda a)$.

This estimator is in fact equal to the estimator $\bar{Y}(n) + \beta\Delta_{h^*}(\bar{Y}(n))$, with $h^*(y) = a(y^2 + y)$, the solution to Poisson's equation. Hence if $\beta = 1$, this estimator has zero variance. \blacksquare

We turn now to the general case. We assume throughout this section that the assumption **(A)** is in place, so that the network is stable, and all steady-state expectations that we use exist. A similar development for reentrant lines may be found in Henderson and Meyn (1997).

Since the solution to Poisson's equation is equivalent to a quadratic, it is reasonable to use $h(y) = y'Qy + q'y$, for some symmetric matrix Q and vector q . This approach was originally proposed in Kumar and Meyn (1996), based on the prior approaches to bounding network performance presented in Kumar and Kumar (1994); Bertsimas et al. (1994) and Meyn and Down (1994).

We first consider the linear part of the function, and then turn to the quadratic terms. Consider the function $h_j(y) = y_j$ for some j . Then, letting w denote the work allocation vector corresponding to the vector y , we have

$$\Delta_{h_j}(y) = \lambda_j - \mu_j w_j + \sum_i \mu_i w_i R_{ij}. \quad (8)$$

Denote by \bar{x} the vector $(\mu_i \mathbb{E}_\pi W_i(0) : 1 \leq i \leq \ell)$. Taking expectations with respect to π in (8) and writing the equations (one for each j) in vector form, we obtain

$$0 = \lambda - \bar{x} + R'\bar{x}.$$

Since the solution to (1) is unique, we conclude that $\mathbb{E}_\pi W_j(0) = \gamma_j/\mu_j$. Thus by considering linear functions we have shown that the usual traffic conditions hold.

Now consider the function $h_{jk}(y) = y_j y_k$ for $j \neq k$. Then,

$$\begin{aligned} \Delta_{h_{jk}}(y) &= \lambda_j y_k + \lambda_k y_j - \mu_j w_j y_k - \mu_k w_k y_j \\ &+ \sum_i \mu_i w_i (R_{ij} y_k + R_{ik} y_j) - \mu_j w_j R_{jk} - \mu_k w_k R_{kj}. \end{aligned} \quad (9)$$

Notice that (9) is a nonlinear expression in y , due to the presence of the w terms. We would prefer to work with linear expressions. To this end, introduce the variables $Z_{ij}(n) = W_i(n)Y_j(n)$, and let $\bar{z}_{ij} = \mathbb{E}_\pi W_i(0)Y_j(0)$ for $1 \leq i, j \leq \ell$. Under our assumptions on the policy it follows that $Z(n)$ is a fixed, deterministic function of $Y(n)$. Furthermore, let $\bar{y}_j = \mathbb{E}_\pi Y_j(0)$. Taking expectations with respect to π in (9), we obtain

$$0 = \lambda_j \bar{y}_k + \lambda_k \bar{y}_j - \mu_j \bar{z}_{jk} - \mu_k \bar{z}_{kj} + \sum_i \mu_i (R_{ij} \bar{z}_{ik} + R_{ik} \bar{z}_{ij}) - \gamma_j R_{jk} - \gamma_k R_{kj}. \quad (10)$$

Now, the non-idling condition implies that whenever $y_j > 0$ so that work is present at station $s(j)$,

$$\sum_{i:s(i)=s(j)} w_i = 1,$$

so that station j allocates all of its effort. Hence, for *any* value of y_j , including 0,

$$y_j = \sum_{i:s(i)=s(j)} w_i y_j,$$

and consequently,

$$\bar{y}_j = \sum_{i:s(i)=s(j)} \bar{z}_{ij}. \quad (11)$$

Therefore, if we let \bar{z} be a column vector containing the \bar{z}_{jk} 's, then the expression (10) may be written as $u'_{jk} \bar{z} = c_{jk}$, for a suitably defined column vector u_{jk} and constant c_{jk} .

By considering the function $h_{jj}(y) = y_j^2$, we obtain

$$0 = 2\gamma_j + 2\lambda_j \bar{y}_j - 2\mu_j \bar{z}_{jj} + 2 \sum_i \mu_i R_{ij} \bar{z}_{ij},$$

for $j = 1, \dots, \ell$, and again these equations can be written as $u'_{jj} \bar{z} = c_{jj}$ for suitably defined u_{jj} and c_{jj} . We obtain one equation for each $1 \leq j \leq k \leq \ell$, so that in all there are $\ell(\ell + 1)/2$ equations of the form $u'_{jk} \bar{z} = c_{jk}$. If we now write the vectors $\{u_{jk}\}$ as columns in a matrix U , and the values $\{c_{jk}\}$ in a vector c , these equations can be written as $U' \bar{z} = c$. The matrix U has ℓ^2 rows, and $\ell(\ell + 1)/2$ columns.

Although we began with expressions involving the Markov chain Y , we are now working with the Markov chain $Z = (Z(n) : n \geq 0)$. Therefore, it is useful to express the function $|y|$ as $p'z$, for some vector p . In particular, $p_{ij} = 1$ if $s(i) = s(j)$, and 0 otherwise. In view of (11), the estimator $\alpha(n)$ may be written as $p' \bar{Z}(n)$, where $\bar{Z}(n) = n^{-1} \sum_{i=0}^{n-1} Z(i)$.

So define the *quadratic* estimator as

$$\begin{aligned}\alpha_q(n) &\triangleq p'\bar{Z}(n) + \beta\nu'(U'\bar{Z}(n) - c) \\ &= (p + \beta U\nu)'\bar{Z}(n) - \beta\nu'c,\end{aligned}\tag{12}$$

where ν is a vector of coefficients.

This is again of the form

$$\alpha_q(n) = |\bar{Y}(n)| + \beta\Delta_h(\bar{Y}(n)),$$

where h is a quadratic, $h(y) = y'Qy + \zeta'y$ for some matrix Q and vector ζ . However in the case of networks, we can only hope that h *approximately* solves Poisson's equation, in the sense that $Ph(y) - h(y) = -d(y) + \alpha_d$, with $d(\cdot) \approx |\cdot|$.

Let us assume (for now) that $\beta = 1$. The variance of (12) is then given by

$$(p + U\nu)'\Lambda_n(p + U\nu),\tag{13}$$

where Λ_n is the covariance matrix of $\bar{Z}(n)$. But under appropriate initial conditions and assuming **(A)** holds, $\Lambda_n \sim \Lambda/n$, where Λ is the time-average covariance matrix for $\bar{Z}(n)$. Then (13) is asymptotically given by

$$n^{-1}\|p + U\nu\|_{\Lambda}^2 \triangleq n^{-1}(p + U\nu)'\Lambda(p + U\nu).\tag{14}$$

Standard control variate methodology suggests that one could estimate the covariance matrix Λ (or $U'\Lambda U$), and then choose ν to minimize the vector norm (14) (Law and Kelton, 2000, p. 609). However, as cautioned in Law and Kelton (2000), there is the danger of a variance increase associated with the additional estimation of the covariance matrix. The “loss factor” is discussed in Lavenberg and Welch (1981) and Nelson (1990) for terminating simulations, and in Loh (1994) for steady-state simulation. The loss factor can become an issue when many control variables are used (as is potentially the case here).

Rather than attempt to determine an optimal selection of control variables from those at our disposal, we instead choose to avoid the issue altogether by preselecting ν , and then using a standard approach to select the single parameter β . See Henderson and Meyn (1997) for further discussion related to this point.

From (14), it is “optimal” to choose ν to minimize $\|U\nu + p\|_\Lambda$. Since Λ is unknown prior to the simulation, we instead choose ν to minimize $\|U\nu + p\|_i$ for some L_i norm $\|\cdot\|$. This problem can be solved using linear programming if i is chosen to be 1 or ∞ , or using least-squares methods if $i = 2$.

In addition, for some workload policies, for example preemptive priority policies, it is known that $\bar{z}_{ij} = 0$ for some i and j . In this case, we may modify the norm used in the minimization slightly to ignore the cost coefficient of \bar{z}_{ij} . See Henderson and Meyn (1997) for further discussion of this point, and Kumar and Meyn (1996) for the related concept of *auxiliary constraints*.

It remains to explain how β is selected. Given a response X , a control variable C , and the form of the controlled estimator $X + \beta C$, it is well known that the value of β that minimizes the controlled variance is $\beta^* = -\text{Cov}(X, C) / \text{Var } C$. In our case, the response X is $p' \bar{Z}(n)$ and the control C is $\nu'(U' \bar{Z}(n) - c)$. One may use any reasonable approach to estimate β^* from the simulation. In particular, Loh (1994) discusses how this may be done in a steady-state context using both regenerative and batch-means approaches. The process Z is regenerative with regeneration times defined by the hitting times of the state 0, but the regenerative cycles can be expected to be very long. So we suggest instead using a batch-means approach to estimating β^* . The required calculations are taken from Loh (1994), and are summarized in the appendix.

Theorem 5, also in the appendix, gives the relevant asymptotic theory for the quadratic estimator. Basically, under the assumption **(A)**, the quadratic estimator converges in probability, and is asymptotically t -distributed when suitably normalized. The convergence mode being “in probability” results from the fact that the estimator for β^* is only weakly consistent. If however, a strongly consistent estimator for β were used, or if β were chosen to be a constant (e.g., 1), then the quadratic estimator would be strongly consistent.

In summary, to estimate α :

1. Choose ν to minimize $\|U\nu + p\|$ for some suitable norm.
2. Simulate the Markov chain Z up until time n . (This amounts to simulating Y since

Z is a deterministic function of Y .)

3. Compute $\beta(n)$, the estimate of β^* .
4. Compute the estimator $\alpha_q(n) = p'\bar{Z}(n) + \beta(n)\nu'(U'\bar{Z}(n) - c)$.

Simulation results for the quadratic estimator and three separate queueing networks are given in Henderson and Meyn (1997). Simulation results for the network in Figure 1 operating under the FBFS (first buffer, first-served) preemptive priority policy are given in Table 1. We minimized $\|U\nu + p\|_2$ to determine ν , and ignored the additional information that $\bar{z}_{31} = 0$ under the chosen service policy. The results are representative of all the other networks we experimented with.

Table 1: Simulation Results for the Two-Station Three-Buffer Example (to two significant figures).

	Standard		Quadratic		
ρ_2	Mean	Var	Mean	Var	Reduction
0.2	0.48	2.1E-4	0.48	1.7E-6	120
0.4	1.26	1.4E-3	1.26	2.7E-5	52
0.6	2.8	1.1E-2	2.8	5.1E-4	22
0.8	6.9	0.19	6.9	2.6E-2	7.1
0.9	14	2.0	14	0.64	3.1
0.95	25	13	25	7.1	1.9
0.99	70	99	70	95	1.0

We took $\mu_1 = \mu_3 = 22$, $\mu_2 = 10$, and then chose λ to ensure that $\rho_2 \triangleq \lambda/\mu_2$ was as specified in the table. Before conducting the experiments, time was rescaled so that $\lambda + \sum_i \mu_i = 1$. Results for the standard estimator are provided in the first two columns, and those for the quadratic estimator are in the next two columns. The simulations were run for 100,000 time steps using 20 batches, and we repeated the experiments 200 times to obtain an estimate of the error in the estimators. For each estimator we supply the estimated mean and variance (over the 200 runs). The column labeled “Reduction” gives

the ratio of the observed variances. The results for $\rho = 0.99$ are subject to some suspicion, owing to the fact that our batch means exhibited correlation.

We see that for lightly to moderately loaded systems, variance reduction factors on the order of between 120 and 7 are observed. The largest variance reduction occurs in light traffic, while smaller variance reductions are obtained in moderately loaded systems. These variance reductions are certainly useful, and come at very little additional computational cost, since the only real additional computational cost in computing the estimator $\alpha_q(n)$ as opposed to the standard estimator $\alpha(n)$ is the solution of the optimization problem to choose ν . But this problem is solved once only before the simulation begins, and takes a (very) small amount of time to solve relative to the computational effort devoted to the simulation.

As discussed earlier, we are more interested in the performance of these simulation estimators in moderately to heavily loaded systems. The results in Table 1 suggest that our estimator is less effective (relative to the standard estimator) in heavy traffic. In particular, for (very) heavily loaded systems, only modest variance reductions are seen.

It is conceivable that the disappointing performance of the controlled estimator in heavy traffic is due to the heuristic of choosing the multipliers ν prior to the simulation, as opposed to attempting to make an “optimal” choice. To test this idea, we estimated the maximum possible variance reduction using the estimator $\alpha_q(n)$. For several networks, and for various traffic loadings, we estimated the time-average covariance matrix Λ . As discussed earlier, the maximum possible variance reduction (asymptotically) is obtained by selecting ν to minimize $\|U\nu + p\|_\Lambda$. The results for the network in Figure 1 are presented in Table 2.

The columns headed “Standard” and “Quadratic” give an estimate of the asymptotic variance for the standard and best possible quadratic estimators respectively. The column headed “Reduction” gives the ratio of these two values. The simulation run lengths required to get reasonable estimates of Λ for $\rho_2 > 0.9$ were infeasibly large, and so we omitted these values.

Comparing these results with those of Table 1, we see that the potential variance reductions appear to decrease as congestion increases. However, substantial variance reductions

Table 2: The best possible performance for the quadratic estimator in the two-station three-buffer example (2 significant figures).

ρ_2	Standard	Quadratic	Reduction
0.2	14	5.7E-3	2500
0.4	88	0.13	680
0.6	670	2.0	340
0.8	1.7E+4	41	420
0.9	7.1E+4	320	220

may yet be possible with a carefully chosen weighting vector ν . In view of the “loss factor,” the question of how best to choose ν to achieve greater variance reduction using the quadratic estimator is an interesting open question. In the absence of a more-effective candidate than the one we have suggested, it would seem that a different approach to generating control variates is warranted for heavily loaded systems.

4 Fluid Models and Stability

To motivate our second approach to generating control variates for the simulation of multiclass queueing networks, we consider an alternative expression (Meyn and Tweedie, 1993, p. 432) for the solution h^* to Poisson’s equation (5), namely

$$h^*(y) = \mathbb{E}_y \sum_{k=0}^{\kappa} (|Y(k)| - \alpha), \quad (15)$$

where $\kappa = \inf\{n \geq 0 : Y(n) = 0\}$.

Although it is difficult to compute (15) exactly, we can certainly approximate it through the use of fluid models. As before, we introduce the key ideas through the M/M/1 queue.

Continuation of Example 1: Consider a uniformized (and discrete-time) process $Y = (Y(n) : n \geq 0)$ describing the queue length in an M/M/1 queue. We have the recursion

$$Y(n+1) = (Y(n) + I(n+1))^+,$$

for $n \geq 0$, where $I = (I(n) : n \geq 1)$ is a Bernoulli, i.i.d. process: $\lambda = P(I(n) = 1)$ is the arrival rate, and $\mu = P(I(n) = -1)$ is the service rate. Time has been normalized so that $\lambda + \mu = 1$.




Figure 2: (a) A sample path Y of the M/M/1 queue with $\rho = \lambda/\mu = 0.9$, and $Y(0) = 400$. (b) A solution to the differential equation $\dot{\phi} = \mathbb{I}(\phi > 0)(\lambda - \mu)$ starting from the same initial condition.

To construct an approximation to the solution to Poisson’s equation, first note that in heavy traffic, the network will typically be somewhat congested, and so we are primarily concerned with “large” states. So it may pay to consider the process starting from a large initial condition. In the left-hand side of Figure 2 we see one such simulation.

One approach to computing the solution to Poisson’s equation is to compute (15). While this is easy for the M/M/1 queue, such computation can be formidable for more complex network models. However, consider the right-hand side of Figure 2 which shows a sample path of the deterministic fluid, or leaky bucket model. This satisfies the differential equation $\dot{\phi} = \mathbb{I}(\phi > 0)(\lambda - \mu)$, where $\mathbb{I}(\cdot)$ is the indicator function that is 1 if its argument is true and 0 otherwise. The behavior of the two processes looks similar when viewed on

this large spatial/temporal scale. It appears that a good approximation is $h^*(x) \approx$

$$\begin{aligned} h(x) &\triangleq \int_0^\infty \phi(t) dt, \quad \phi(0) = x, \\ &= \frac{1}{2} \frac{x^2}{\mu - \lambda}. \end{aligned} \tag{16}$$

This is the same approximation arrived at in Section 3. ■

Of course, the M/M/1 queue is a very special case of a multiclass queueing network, and so it is worthwhile investigating this approximation more carefully before adopting it wholesale. We return now to the case of a general multiclass queueing network.

The dynamics of the process Y can be described by a random linear system after a slight extension of the previous definitions. Define a sequence of i.i.d. random matrices $\{I(n) : n \geq 1\}$ on $\{0, 1\}^{(\ell+1)^2}$, with $\mathbb{P}\{\sum_j \sum_k I_{j,k}(n) = 1\} = 1$, and $\mathbb{E}[I_{j,k}(n)] = \mu_j R_{jk}$, where μ_j denotes the service rate for class j customers. Note that exactly one element of $I(n)$ is positive for each n . These random variables indicate which event in the uniformized process $Y = (Y(n) : n \geq 0)$ is to occur. The variable $I_{jk}(n) = 1$ if and only if a class j job completes service and moves to station k . It is convenient to capture the exogenous arrival processes within the same framework. An exogenous arrival is indicated by $j = 0$, and a departure from the system is indicated by $k = 0$. For $j = 0$, let $\mu_0 \triangleq \sum_{k=1}^\ell \lambda_k$ denote the overall arrival rate of customers to the system. Define $R_{0,0} = 0$, and for $1 \leq k \leq \ell$, $R_{0,k} = \lambda_k / \mu_0$. Thus, we pool all customer arrivals into one stream with rate μ_0 , and an arriving customer is allocated to one of the ℓ classes according to the appropriate probability.

For $1 \leq j \leq \ell$, let $W_j(n) = 1$ if station $s(j)$ is allocating its entire effort to customers of class j at time n , and 0 otherwise. As before we require that $W_j(n)$ is a deterministic function of $Y(n)$ for all $n \geq 0$ and all $j = 1, \dots, \ell$. We define $W_0(n) = 1$ for all n , indicating that the exogenous arrival process is always active. For $1 \leq k \leq \ell$, let e^k denote the k th basis vector in \mathbb{R}^ℓ , and set $e^0 \triangleq 0$. The random linear system can then be defined as

$$Y(n+1) = Y(n) + \sum_{j=0}^\ell \sum_{k=0}^\ell I_{j,k}(n+1) [-e^j + e^k] W_j(n), \tag{17}$$

where the state process Y denotes the vector of customer classes in the system as before.

To define the fluid model associated with this network we suppose that the initial condition is large so that $m = |Y(0)| \gg 1$. We then construct a continuous time process $\phi^y(t)$ as follows: If tm is an integer, we set

$$\phi^y(t) = \frac{1}{m}Y(mt),$$

where $Y(0) = y$ and $|y| = m$. For all other $t \geq 0$, we define $\phi^y(t)$ by linear interpolation, so that it is continuous and piecewise linear in t . Note that $|\phi^y(0)| = 1$, and that ϕ^y is Lipschitz continuous (see, e.g., Apostol (1969, p. 229) for a definition of Lipschitz continuity.) The collection of all “fluid limits” is defined by

$$\mathcal{L} \triangleq \bigcap_{m=1}^{\infty} \overline{\{\phi^y : |y| > m\}}$$

where the overbar denotes weak closure. The set \mathcal{L} depends upon the particular policy chosen, and for many policies such as preemptive priority policies, it is a family of purely deterministic functions.

Any process $\phi \in \mathcal{L}$ evolves on the state space \mathbb{R}_+^ℓ and, for a wide class of scheduling policies, satisfies a differential equation of the form

$$\frac{d}{dt}\phi(t) = \sum_{j=0}^{\ell} \sum_{k=0}^{\ell} \mu_j R_{jk} [-e^j + e^k] u_j(t) \quad (18)$$

where the function $u(\cdot)$ is analogous to the discrete control, and satisfies similar constraints (see the M/M/1 queue model described earlier, or Dai (1995) and Dai and Weiss (1996) for more general examples). In many cases the differential equation (18) admits a unique solution, from any initial condition, even though typically in practice the control u is a discontinuous function of the state ϕ (consider again any priority policy).

It is now known that stability of (17) is closely connected with the stability of the fluid model (Dai, 1995; Kumar and Meyn, 1996; Dai and Meyn, 1995). The fluid model \mathcal{L} is called L_p -stable if

$$\lim_{t \rightarrow \infty} \sup_{\phi \in \mathcal{L}} \mathbb{E}[|\phi(t)|^p] = 0.$$

Let T_0 denote the first hitting time $\inf\{t \geq 0 : \phi(t) = 0\}$. It is shown in Meyn (1997) that $\sup_{\phi \in \mathcal{L}} \mathbb{E}[T_0] < \infty$ when the model is L_2 -stable. Hence, when \mathcal{L} is non-random, L_2 -stability is equivalent to *stability* in the sense of Dai (1995): There is some time T such

that $\phi(t) = 0$ for $t \geq T$, $\phi \in \mathcal{L}$. For example, in the M/M/1 queue with $\lambda < \mu$, the queue eventually hits 0 as seen in Figure 2.

The following result is a minor generalization of results from Kumar and Meyn (1996); Dai and Meyn (1995). Its proof is omitted.

Theorem 3 *The following two stability criteria are equivalent for the network under any non-idling policy, and any $p \geq 2$.*

(i) *There is a function V , and a constant $b < \infty$ satisfying*

$$PV(y) - V(y) \leq -|y|^{p-1} + b$$

where for some $\delta > 0$,

$$\delta(1 + |y|^p) \leq V(y) \leq \delta^{-1}(1 + |y|^p), \quad y \in S. \quad (19)$$

(ii) *The fluid model \mathcal{L} is L_p -stable.*

■

Thus, L_p stability can be verified through the Lyapunov condition **(A)**. Using this result it is possible to show that the solution to Poisson's equation is asymptotically equal to a value function for the associated fluid model, provided that the fluid model is L_2 -stable. It can be shown that many policies for the fluid model are piecewise constant on a finite set of cones in \mathbb{R}_+^ℓ . This is certainly the case for buffer priority policies, and also holds for L_1 optimal policies (for a discussion see Weiss (1995)). It then follows that for such policies the fluid value function V is piecewise quadratic. The proof of the following result appears in the appendix.

Theorem 4 *Suppose that for a given non-idling policy w , the fluid model \mathcal{L} is L_2 -stable and non-random. Suppose moreover that limits are unique, in the sense that $\phi_1(0) \neq \phi_2(0)$ for any two distinct $\phi_i \in \mathcal{L}$.*

Then a solution h^ to Poisson's equation exists, and*

$$\limsup_{|y| \rightarrow \infty} \left| \frac{h^*(y)}{V(y)} - 1 \right| = 0,$$

where

$$V(y) = |y|^2 \int_0^\infty |\phi(t)| dt, \quad \phi(0) = \frac{y}{|y|}.$$

Hence, under the conditions of Theorem 4, a solution h to Poisson's equation is intimately related to the fluid value function V . This result then strongly motivates the use of a fluid value function as an approximation for the solution to Poisson's equation.

Another way to motivate the fluid approximation is to note that (from (17))

$$\begin{aligned} Y(n+1) &= Y(n) + \sum_{j,k} \mu_j R_{jk} (-e^j + e^k) W_j(n) \\ &+ \sum_{j,k} (I_{jk}(n+1) - \mu_j R_{jk}) (-e^j + e^k) W_j(n) \\ &= Y(n) + BW(n) + D(n+1) \\ &= Y(0) + \sum_{i=0}^n BW(i) + M(n+1), \end{aligned} \tag{20}$$

where B is an $(\ell+1) \times (\ell+1)$ matrix. The process $M(\cdot)$ is a vector-valued martingale with respect to the natural filtration, and $D(i)$ is the martingale difference $M(i) - M(i-1)$. (See, e.g., Ross 1996 for an introduction to martingales.) It is straightforward to check that $ED(i)'D(i)$ is bounded in i (by b say), so that $EM(n)'M(n) \leq bn$ for all $n \geq 0$. Hence, the network is essentially a deterministic fluid model with a 'disturbance' M . When the initial condition $Y(0)$ is large, then the state dominates this disturbance, and hence the network behavior appears deterministic.

We therefore have strong motivation for approximating the solution to Poisson's equation h^* by V , where

$$V(y) \triangleq \int_0^\infty \phi(t) dt,$$

and ϕ solves the differential equation (18) with $\phi(0) = y$. The *fluid estimator* of α is then

$$\alpha_f(n) = |\bar{Y}(n)| + \frac{\beta}{n} \sum_{k=0}^{n-1} \Delta_V(Y(k)). \tag{21}$$

The parameter β is again a constant that may be chosen to attempt to minimize the variance of the fluid estimator. We use the methodology outlined in the appendix to estimate the optimal β . Consequently, the asymptotic results for the quadratic estimator

also apply here, namely that under the assumption **(A)**, the fluid estimator is weakly consistent, and is asymptotically t -distributed when suitably normalized.

Clearly, to implement the fluid estimator we need to be able to compute Δ_V . For any function V , we have that

$$\begin{aligned} PV(y) &= \mathbb{E}_y V(Y(1)) \\ &= \sum_{j=0}^{\ell} \sum_{k=0}^{\ell} \mu_j R_{jk} V(y - e^j + e^k) W_j(y), \end{aligned}$$

so that it suffices to be able to compute $V(y)$ for $y \in S$.

Solving the differential equation (18) to find ϕ is not difficult when the fluid control u is piecewise constant on a finite set of cones in \mathbb{R}_+^ℓ since in this case ϕ is piecewise linear. Integrating ϕ to find the fluid value function V is then straightforward. In other words, for a specific model, some preliminary work has to be done to give code that can compute V , but this is usually not a difficult step. Algorithms for computation or approximation of V are described in Eng and Meyn (1996) and Borkar and Meyn (2003).

It should be apparent from the above discussion that computing the control Δ_V for the estimator (21) may be moderately time-consuming (computationally speaking) relative to the time taken to simply simulate the process Y . Be that as it may, it is certainly the case that the time taken to compute the control is relatively insensitive to the congestion in the system.

In Table 3 we present simulation results for the fluid estimator on the network of Figure 1. (Similar results were obtained for all other multiclass queueing networks that we tried.) The entries in Table 3 have the same interpretation as those in Table 1. In particular, the column headed ‘‘Reduction’’ represents the variance reduction factor over the standard estimator.

The best value of β was found to be close to unity in each of the simulations, particularly at high loads where it was found to be within $\pm 5\%$ of unity.

Observe that for low traffic intensities, the fluid estimator yields reasonable variance reductions over the standard estimator. However, because it is more expensive to compute than the standard estimator, these results are not particularly encouraging. But as the system becomes more and more congested, the fluid estimator yields large variance

Table 3: Simulation results for the two-station three-buffer example (2 significant figures). The interpretation of the values given is the same as in Table 1.

ρ_2	Mean	Var	Reduction
0.2	0.47	6.1E-5	3.5
0.4	1.3	4.0E-4	3.5
0.6	2.8	3.5E-3	3.1
0.8	6.9	4.3E-2	4.4
0.9	14	0.17	12
0.95	26	0.23	56
0.99	110	0.98	100

reductions over the standard estimator, meaning that the extra computational effort per iteration is certainly worthwhile. For very high traffic intensities, the fluid estimator *significantly* outperforms both the standard estimator and the quadratic estimator, and so we have achieved our goal of deriving an estimator that can be effective in heavy traffic.

5 Conclusions

We have given two simulation estimators for estimating a linear function of the steady-state customer class population.

The quadratic estimator produces very useful variance reductions in light to moderate traffic at very little additional computational cost. We recommend that it be used in simulations of such lightly loaded networks. The quadratic estimator is less effective in simulations of heavily loaded networks, but could potentially provide useful variance reductions in this regime if a better choice of weighting vector ν can be employed.

The fluid estimator provides modest variance reduction in light to moderate traffic, but appears to be very effective in heavy traffic. There is an additional computational overhead in computing the fluid estimator, but this overhead is (roughly) independent of the load on the network. Hence we may conclude that in heavily loaded systems, the fluid

estimator should yield significant computational improvements, and should therefore be used.

One might conclude from the above discussion that the quadratic and fluid estimators could be combined using the method of multiple control variates to yield a single “combined” estimator. However, we believe that it is unlikely that a combined estimator would yield significant improvements over the use of either the quadratic estimator (in light traffic) or the fluid estimator (in heavy traffic). In light traffic, we expect that the additional reductions in variance would be negated by the increased computational effort. And in heavy traffic we expect that any additional variance reduction would be modest, owing to the weaker performance of the quadratic estimator in this regime.

Finally, we note that Veatch (1995) explores bounded perturbations of the fluid value function in approximate dynamic programming. Related techniques are considered in current research to refine the fluid estimator.

References

- T. M. Apostol. *Calculus, Volume 2*. Wiley, New York, 2nd edition, 1969.
- S. Asmussen. Queueing simulation in heavy traffic. *Mathematics of Operations Research*, 17:84–111, 1992.
- F. Baskett, K. M. Chandy, R. R. Muntz, and F. G. Palacios. Open, closed, and mixed networks of queues with different classes of customers. *Journal of the ACM*, 22:248–260, 1975.
- D. Bertsekas and R. Gallager. *Data Networks*. Prentice Hall, Englewood Cliffs, NJ, 1987.
- D. Bertsimas, I. Paschalidis, and J. N. Tsitsiklis. Optimization of multiclass queueing networks: polyhedral and nonlinear characterizations of achievable performance. *Ann. Appl. Probab.*, 4:43–75, 1994.
- P. Billingsley. *Convergence of Probability Measures*. Wiley, New York, 1968.

- V. S. Borkar and S. P. Meyn. Value functions and performance evaluation in stochastic network models. 42nd IEEE Conference on Decision and Control, December 9-12, 2003 Hyatt Regency Maui, Hawaii, USA.
- J. A. Buzacott and J. G. Shanthikumar. *Stochastic Models of Manufacturing Systems*. Prentice Hall, Englewood Cliffs, NJ, 1993.
- J. G. Dai. On positive Harris recurrence of multiclass queueing networks: a unified approach via fluid limit models. *Ann. Appl. Probab.*, 5(1):49–77, 1995.
- J. G. Dai and J. M. Harrison. Reflected Brownian motion in an orthant: numerical methods for steady-state analysis. *Ann. Appl. Probab.*, 2(1):65–86, 1992.
- J. G. Dai and S. P. Meyn. Stability and convergence of moments for multiclass queueing networks via fluid limit models. *IEEE Transactions on Automatic Control*, 40:1889–1904, 1995.
- J. G. Dai and G. E. Weiss. Stability and instability of fluid models for reentrant lines. *Math. Operations Res.*, 21:115–134, 1996.
- S. P. Meyn and D. G. Down. Stability of Generalized Jackson Networks. *Annals Appl. Probab.*, 4:124–148, 1994.
- J. Eng, D. Humphrey and S. P. Meyn. Fluid network models: Linear programs for control and performance bounds. In J. C. J. Gertler and M. Peshkin, editors, *Proceedings of the 13th IFAC World Congress*, volume B, pages 19–24, San Francisco, California, 1996.
- S. B. Gershwin. *Manufacturing Systems Engineering*. Prentice Hall, Englewood Cliffs, NJ, 1993.
- P. W. Glynn and W. Whitt. The asymptotic efficiency of simulation estimators. *Operations Research*, 40:505–520, 1992.
- J. M. Harrison and R. J. Williams. Multidimensional reflected Brownian motions having exponential stationary distributions. *Ann. Probab.*, 15(1):115–137, 1987.

- J. M. Harrison and R. J. Williams. On the quasireversibility of a multiclass Brownian service station. *Ann. Probab.*, 18(3):1249–1268, 1990.
- S. G. Henderson. *Variance Reduction Via an Approximating Markov Process*. PhD thesis, Department of Operations Research, Stanford University, Stanford, California, USA, 1997.
- S. G. Henderson and P. W. Glynn. Approximating martingales for variance reduction in Markov process simulation. *Mathematics of Operations Research*, 27:253–271, 2002.
- S. G. Henderson and S. P. Meyn. Efficient simulation of multiclass queueing networks. In S. Andradottir, K. J. Healy, D. H. Withers, and B. L. Nelson, editors, *Proceedings of the 1997 Winter Simulation Conference*, pages 216–223, Piscataway NJ, 1997. IEEE.
- S. G. Henderson, S. P. Meyn, and V. Tadić. Performance evaluation and policy selection in multiclass networks. *Discrete Event Dynamic Systems*, 13:149–189, 2003. Special issue on learning and optimization methods.
- J. R. Jackson. Jobshop-like queueing systems. *Management Science*, 10:131–142, 1963.
- F. P. Kelly. *Reversibility and Stochastic Networks*. Wiley, New York, NY, 1979.
- S. Kim and S. G. Henderson. Adaptive control variates. In R. Ingalls, M. Rossetti, J. Smith, and B. Peters, editors, *Proceedings of the 2004 Winter Simulation Conference*, pages 621–629, Piscataway NJ, 2004. IEEE.
- S. Kim and S. G. Henderson. Adaptive control variates. 2005. Submitted for publication.
- P. R. Kumar and S. P. Meyn. Duality and linear programs for stability and performance analysis queueing networks and scheduling policies. *IEEE Transactions on Automatic Control*, 41(1):4–17, 1996.
- S. Kumar and P. R. Kumar. Performance bounds for queueing networks and scheduling policies. *IEEE Transactions on Automatic Control*, 39:1600–1611, 1994.
- S. S. Lavenberg. *Computer Performance Modeling Handbook*. Academic Press, New York, NY, 1983.

- S. S. Lavenberg and P. D. Welch. A perspective on the use of control variables to increase the efficiency of Monte Carlo simulations. *Management Science*, 27:322–335, 1981.
- A. M. Law and W. D. Kelton. *Simulation Modeling and Analysis*. McGraw-Hill, New York, 3rd edition, 2000.
- W. W. Loh. *On the Method of Control Variates*. PhD thesis, Department of Operations Research, Stanford University, Stanford, CA, 1994.
- S. P. Meyn. Stability and optimization of queueing networks and their fluid models. In *Mathematics of Stochastic Manufacturing Systems (Williamsburg, VA, 1996)*, pages 175–199. Amer. Math. Soc., Providence, RI, 1997.
- S. P. Meyn. Sequencing and routing in multiclass queueing networks. Part I: Feedback regulation. *SIAM J. Control Optim.*, 40(3):741–776, 2001.
- S. P. Meyn. Workload models for stochastic networks: Value functions and performance evaluation. To appear, *IEEE Transactions on Automatic Control*, 2005.
- S. P. Meyn. Large deviation asymptotics and control variates for simulating large functions. Submitted for publication, 2004.
- S. P. Meyn and R. L. Tweedie. *Markov Chains and Stochastic Stability*. Springer-Verlag, London, 1993.
- J. R. Morrison and P. R. Kumar. New linear program performance bounds for queueing networks. *Journal of Optimization Theory and Applications*, 100:575–597, 1999.
- B. L. Nelson. Control variate remedies. *Operations Research*, 38:974–992, 1990.
- M. F. Neuts. *Matrix-Geometric Solutions in Stochastic Models*. Dover Publications Inc., New York, 1994. An Algorithmic Approach, Corrected reprint of the 1981 original.
- S. M. Ross. *Stochastic Processes*. Wiley, New York, 2nd edition, 1996.
- P. J. Schweitzer. Aggregation methods for large Markov chains. In *Mathematical Computer Performance and Reliability (Pisa, 1983)*, pages 275–286. North-Holland, Amsterdam, 1984.

- E. Schwerer. A linear programming approach to the steady-state analysis of reflected Brownian motion. *Stochastic Models*, 17:341–368, 2001.
- X. Shen, H. Chen, J. G. Dai, and W. Dai. The finite element method for computing the stationary distribution of an SRBM in a hypercube with applications to finite buffer queueing networks. *Queueing Systems*, 42(1):33–62, 2002.
- M. H. Veatch. Approximate dynamic programming for networks: Fluid models and constraint reduction. Submitted for publication, 2004.
- G. Weiss. Optimal draining of a fluid re-entrant line. In F. Kelly and R. Williams, editors, *Volume 71 of IMA volumes in Mathematics and its Applications*, pages 91–103, New York, 1995. Springer-Verlag.
- W. Whitt. Planning queueing simulations. *Management Science*, 35:1341–1366, 1989.

Appendix

Proof of Theorem 2

It is straightforward to show that

$$\Delta_{f^2}(y) = -2(1 - \rho)f(y) + e'Q[L + Q^{-1}W - \mathcal{M}WR + \text{diag}(e'\mathcal{M}WR)]Q'e, \quad (22)$$

where W is the diagonal matrix containing the work allocation vector w corresponding to y , and $L = \text{diag}(\lambda)$. Furthermore, for $m \geq 3$,

$$\Delta_{f^m}(y) = -m(1 - \rho)f^{m-1}(y) + \text{lower order terms}. \quad (23)$$

It follows from (23) with $m = 5$ that **(A)** holds. This implies (Theorem 17.5.3 of Meyn and Tweedie (1993)) that the TAVC's are finite, and furthermore, that

$$\sigma_f^2(\rho) = \lim_{n \rightarrow \infty} n \text{Var}(d'\bar{Y}(n)), \text{ and}$$

$$\sigma_i^2(\rho) = \lim_{n \rightarrow \infty} n \text{Var}(\bar{Y}_i(n)).$$

The fact that h is of the form given in the theorem follows from (22). Observe that the function $f^2/(2(1-\rho))$ is “almost” the solution to Poisson’s equation. It needs to be adjusted slightly to remove the terms in the RHS of (22) involving the work allocation vector w . These terms are of the form $f'w$, where the coefficients in f are bounded in ρ . Therefore, the solution to Poisson’s equation is as given in the theorem.

So then the TAVC $\sigma_f^2(\rho)$ is given by

$$\mathbb{E} h(Y)(f(Y) - d'\bar{y}) = \frac{\text{Cov}(f^2(Y), f(Y))}{2(1-\rho)} - \mathbb{E}(c'_\rho Y (f(Y) - d'\bar{y})), \quad (24)$$

where Y is distributed according to the stationary distribution π and $\bar{y} = \mathbb{E} Y$.

Since $\mathbb{E} \Delta_{f^m}(Y) = 0$ for $m = 0, \dots, 4$, it follows from (22) that $\mathbb{E} f(Y)$ is of the order $(1-\rho)^{-1}$ as $\rho \rightarrow 1$. Then, by induction using (23), $\mathbb{E} f^m(Y)$ is of the order $(1-\rho)^{-m}$ as $\rho \rightarrow 1$ for $m = 1, \dots, 4$.

The second term on the right-hand side of (24) is therefore of the order $(1-\rho)^{-3}$ as $\rho \rightarrow 1$. As for the first term, we have the easily proved inequality that for any non-negative r.v. X with $\mathbb{E} X^3 < \infty$,

$$\text{Cov}(X^2, X) \geq \frac{\text{Var} X}{\mathbb{E} X^2} \mathbb{E} X^3.$$

Applying this inequality to the first term on the right-hand side of (24), and noting that $\text{Var} f(Y)$ is of the same order as $\mathbb{E} f(Y)^2$ as $\rho \rightarrow 1$, we obtain the required result that $\sigma_f^2(\rho)$ is of the order $(1-\rho)^{-4}$ as $\rho \rightarrow 1$.

The last statement of the theorem follows from the fact that

$$\begin{aligned} \sigma_f^2(\rho) &= \lim_{n \rightarrow \infty} n \text{Var}(d'\bar{Y}(n)) \\ &= \lim_{n \rightarrow \infty} n \text{Var}\left(\sum_{i=1}^{\ell} d_i \bar{Y}_i(n)\right) \\ &\leq \lim_{n \rightarrow \infty} \ell n \sum_{i=1}^{\ell} \text{Var}(d_i \bar{Y}_i(n)) \\ &= \ell \sum_{i=1}^{\ell} d_i^2 \sigma_i^2(\rho). \end{aligned}$$

■

Control Variates in Steady-State Simulation

We repeat formulae from Loh (1994) for estimating the control variate parameter β that is used in both the quadratic estimator (12) and the fluid estimator (21) via the batch means method of simulation output analysis. To encapsulate both estimators and avoid repetition, we give the formulae for the case where a real-valued stochastic process $X = (X(n) : n \geq 0)$ is simulated, and a real-valued control $C = (C(n) : n \geq 0)$ is recorded.

Let the b batches each consist of m observations, so that the simulation run-length $n = mb$. If X_i and C_i are the i th batch means of the process and control respectively, then for $0 \leq i \leq b - 1$, we have

$$X_i = \frac{1}{m} \sum_{j=im}^{(i+1)m-1} X(j) \text{ and } C_i = \frac{1}{m} \sum_{j=im}^{(i+1)m-1} C(j).$$

Let \bar{X}_n and \bar{C}_n denote the overall (sample) means of the process and control respectively.

Define

$$\begin{aligned} V_{XX}(n) &= \frac{1}{b-1} \sum_{i=0}^{b-1} (X_i - \bar{X}_n)^2, \\ V_{CC}(n) &= \frac{1}{b-1} \sum_{i=0}^{b-1} (C_i - \bar{C}_n)^2, \text{ and} \\ V_{XC}(n) &= \frac{1}{b-1} \sum_{i=0}^{b-1} (X_i - \bar{X}_n)(C_i - \bar{C}_n). \end{aligned}$$

Define $\beta = -V_{XC}/V_{CC}$, and let $\alpha_n = \bar{X}_n + \beta\bar{C}_n$ be the controlled estimator.

Finally, let

$$R^2(n) = \frac{b-1}{b-2} \left(V_{XX}(n) - \frac{V_{XC}(n)^2}{V_{CC}(n)} \right)$$

and

$$S^2(n) = R^2(n) \left(\frac{1}{b} + \frac{1}{b-1} \frac{\bar{C}_n^2}{V_{CC}} \right). \quad (25)$$

Using the above computational process, we can construct the quadratic estimator $\alpha_q(n)$ and the fluid estimator $\alpha_f(n)$. The following result describes the asymptotic behavior of these estimators.

Theorem 5 (Loh 1994) Under the assumption **(A)**, the estimators $\alpha_q(n)$ and $\alpha_f(n)$ converge in probability to α , and for $j = q, f$,

$$\frac{\alpha_j(n) - \alpha}{S_j(n)} \Rightarrow T_{b-2},$$

where T_{b-2} has the Student's t -distribution with $b - 2$ degrees of freedom, and $S_j(n)$ is defined in the obvious way through (25).

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

5.1 Proof of Theorem 4

First note that for any $\phi \in \mathcal{L}$ we have $|\phi(0)| = 1$, and under the assumptions of the theorem there is a $T > 0$ such that

$$V(y) = |y|^2 \int_0^T |\phi(t)| dt, \quad \phi(0) = \frac{y}{|y|}, \quad \phi \in \mathcal{L}.$$

We shall fix such a T throughout the proof.

From Theorem 3, a Lyapunov function exists that satisfies (4), which is equivalent to a quadratic in the sense of (19). It follows that $\pi(c) < \infty$, where $c(y) = |y|$, and that a solution to Poisson's equation exists which is bounded from above by a quadratic, and uniformly bounded from below (Meyn and Tweedie, 1993, p. 432).

We can then take the solution to Poisson's equation and iterate as follows: $P^n h = h - \sum_{i=0}^{n-1} P^i \bar{c}$, where $\bar{c}(y) = |y| - \alpha$. Let $m = |y|$ and take $n = [mT] =$ the integer part of mT to give

$$\frac{\mathbb{E}_y[h(Y(mT))]}{m^2} = \frac{h(y)}{m^2} - \frac{\mathbb{E}_y\left[\sum_{i=0}^{[mT]-1} \frac{|Y(i)|}{m}\right]}{m} - \frac{T\alpha}{m}.$$

Since h is bounded above by a quadratic, there is a $K < \infty$ such that

$$\left| \frac{[h(Y(mT))]}{m^2} \right| \leq K \left(1 + \frac{|Y(mT)|^2}{m^2} \right)$$

The random variable on the right hand side is uniformly bounded by $K(1 + 1/m + T)^2$ for all initial y since at most one customer can arrive during each time slot. It then follows

from weak convergence (see, e.g., Billingsley (1968)) and the definition of T that

$$\limsup_{|y| \rightarrow \infty} \mathbb{E}_y \left| \frac{[h(Y(mT))]}{m^2} \right| = 0$$

Moreover, again by Lipschitz continuity of the fluid model we have

$$\limsup_{|y| \rightarrow \infty} \left| \mathbb{E}_y \left[\frac{1}{m} \sum_{i=0}^{[mT]-1} \frac{|Y(i)|}{m} \right] - V\left(\frac{y}{m}\right) \right| = 0.$$

Putting these results together we see that

$$\limsup_{|y| \rightarrow \infty} \left| \frac{h(y)}{m^2} - V\left(\frac{y}{m}\right) \right| = 0,$$

proving the result. ■

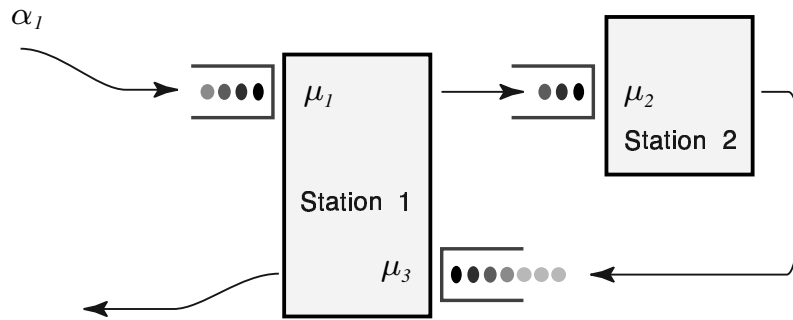


Figure 1: A multiclass network with $d = 2$ stations and $\ell = 3$ customer classes.




Figure 2: (a) A sample path Y of the M/M/1 queue with $\rho = \lambda/\mu = 0.9$, and $Y(0) = 400$.
 (b) A solution to the differential equation $\dot{\phi} = \mathbb{I}(\phi > 0)(\lambda - \mu)$ starting from the same initial condition.

Table 1: Simulation Results for the Two-Station Three-Buffer Example (to two significant figures).

	Standard		Quadratic		
ρ_2	Mean	Var	Mean	Var	Reduction
0.2	0.48	2.1E-4	0.48	1.7E-6	120
0.4	1.26	1.4E-3	1.26	2.7E-5	52
0.6	2.8	1.1E-2	2.8	5.1E-4	22
0.8	6.9	0.19	6.9	2.6E-2	7.1
0.9	14	2.0	14	0.64	3.1
0.95	25	13	25	7.1	1.9
0.99	70	99	70	95	1.0

Table 2: The best possible performance for the quadratic estimator in the two-station three-buffer example (2 significant figures).

ρ_2	Standard	Quadratic	Reduction
0.2	14	5.7E-3	2500
0.4	88	0.13	680
0.6	670	2.0	340
0.8	1.7E+4	41	420
0.9	7.1E+4	320	220

Table 3: Simulation results for the two-station three-buffer example (2 significant figures).
 The interpretation of the values given is the same as in Table 1.

ρ_2	Mean	Var	Reduction
0.2	0.47	6.1E-5	3.5
0.4	1.3	4.0E-4	3.5
0.6	2.8	3.5E-3	3.1
0.8	6.9	4.3E-2	4.4
0.9	14	0.17	12
0.95	26	0.23	56
0.99	110	0.98	100