Proceedings of the 1993 Winter Simulation Conference
G.W. Evans, M. Mollaghasemi, E.C. Russell, W.E. Biles (eds.)

THE SURROGATE ESTIMATION APPROACH FOR SENSITIVITY ANALYSIS IN QUEUEING NETWORKS

Felisa J. Vázquez-Abad
Dép. d'informatique et recherche opérationnelle
Université de Montréal, C.P. 6128 Succ. A
Montréal, Québec H3C 3J7
CANADA

Harold J. Kushner
Division of Applied Mathematics
Brown University, Box F.
Providence, R.I. 02912
U.S.A.

ABSTRACT

The construction of good gradient estimators, known as sensitivity analysis, has recently been the subject of many studies. Gradient estimators may be used to optimize the performance of stochastic processes. In this paper we propose a method for sensitivity estimation that can be used for the optimization of stochastic discrete event dynamic systems (DEDS).

Most of the current methods for gradient estimation have limited applicability for complex problems. The surrogate estimation approach that we propose uses the system's dynamics and heuristic relations in the parameters to construct the desired gradient estimators using local sensitivity estimators.

We present an example of routing in an open data network. Some of the most successful methods for gradient estimation–such as the IPA method–cannot be applied directly and other methods are inappropriate for real time operation. We show how the estimation of the gradient of the stationary average sojourn time with respect to the routing probabilities can be decomposed in terms of local sensitivities. Each node needs to estimate the derivatives of the average queue length with respect to its own arrival rate. The computation is thus distributed and the estimation of the local sensitivities is a much simpler problem, suited for IPA. The amount of calculations required for our approach is proportional to the number of nodes in the network. Those required for direct estimation grow with the number of nodes, the number of outgoing links and the number of destinations.

Our simulations indicate that surrogate estimation is very efficient, even when some of the required assumptions for IPA estimation are not satisfied.

1 INTRODUCTION

Telecommunication systems can be modeled as queueing networks, where performance optimization can be in principle achieved by adaptively choosing some control variables. Many of the successful techniques for adaptive control require good estimators of the gradients of the sensitivity of the system's performance with respect to the control variable. In this paper we deal with the estimation of such gradients, studying the problem of routing in an open network.

It is common practice to use an approximate expression of the derivatives that is obtained under Kleinrock's independence assumption (see Bertsekas, Gallager, 1987; Tsitsiklis, Bertsekas, 1986, and Cotton, Mason, 1991). Mean flow data is then estimated and used in the approximation equations, as if the network were a Jackson network with a product form solution. As our results show, the method that we use can be substantially better than the Jackson network approximation, which yields very poor estimates.

An alternative to the Jackson approximation is to estimate the desired gradients from the measurements of the process. Several methods have been proposed in order to estimate the sensitivity with respect to certain control variables for queueing networks (see L'Ecuyer, 1991 and references therein). Most of the available theoretical results are concerned with the sensitivity of the system's performance with respect to service or arrival rates as control variables and the proofs are obtained mostly using regenerative systems and in many cases analyzing only one server queue. The implementation of such methods to more complex problems such as routing in a data network is very difficult and in large networks the decentralized nature of the controllers may render direct estimation inappropriate.

In this work we develop a method that can be used to get sensitivity estimators for a broad class of parameters. We call it the surrogate estimation approach. It combines some features of the mean flow data approach to implement sensitivity estimators for simpler subsystems of the network. It is relatively easy to implement and it actually reduces the operation count required to estimate the desired gradients, compared to direct estimation. This method aims at redefining the problem to calculate sensitivities of local quantities with respect to local variables at several subsystems of the queueing network. Then this information, broadcast to the controllers, is used to get the estimation of the desired gradients. Thus the computation is distributed. Surrogate estimation was first used in Ho, Cao, 1985 in the context of routing in a closed network, where the derivatives with respect to routing probabilities are calculated indirectly via the derivatives with respect to service rates.

In section 2 we discuss our model example of optimal routing in an open network, where the performance measure can be written in a separable form. The surrogate estimation approach is introduced in section 3. We show how the derivatives with respect to the routing parameters can be estimated using local derivative estimators with respect to the arrival rates at each node. We also discuss the advantages of the surrogate estimation method.

In section 4 we present the implementation of two infinitesimal perturbation analysis (IPA) estimators to the surrogate estimation approach. Simulation results are included in section 5. This shows the robustness of surrogate estimation, which works much better than the estimators obtained under the assumption that the system is a Jackson network.

2 ROUTING IN AN OPEN NETWORK

Problem Formulation. We shall describe this problem in its simplest form in order to focus on the main ideas of our approach, but most of the simplifying assumptions can be easily extended to more complex models. There are N_0 nodes in the network. We consider for simplicity that all arriving customers share the same destination and that there are no loops in the possible paths from any origin to the destination. We assume that interarrival times, service times and routing decisions are mutually independent random variables. For each node n_α, $\alpha = 1, \ldots, N_0$, the arrival processes are Poisson with parameter λ_α. The sequence of service times for each server are identically distributed with mean values denoted by $1/\nu_\alpha < \infty$. Let $\theta = \{\theta_{\alpha\beta};\ \alpha, \beta = 1, \ldots, N_0\}$ denote the set of routing parameters, where $\theta_{\alpha\beta}$ is the steady state fraction of the flow from n_α which goes to n_β. We focus on randomized routing, where a customer that completes service at node n_α is routed to node n_β with probability $\theta_{\alpha\beta}$ independently of past routing choices. Define r_α as the limit of the total number of arrivals to n_α per unit time. This is the mean arrival rate of customers to n_α. We assume that the routing parameters are such that $\nu_\alpha > r_\alpha$.

Under the usual ergodicity assumptions (Sigman, 1990) the condition $\nu_\alpha > r_\alpha$ ensures stability of the process and we can write:

$$r_\beta = \lambda_\beta + \sum_{\alpha=1}^{N_0} \theta_{\alpha\beta}\, r_\alpha \tag{1}$$

The problem under study is the minimization of a performance function $L(\theta)$ by choosing the control variables θ. A sequential optimization procedure can be used to update the control variable θ in the following way. Suppose that over the time interval $[kT, (k+1)T)$ the value of the control is fixed at $\theta^{(k)}$ and that an estimate $Z_k(\theta^{(k)})$ of the gradient $\nabla_\theta L(\theta^{(k)})$ can be obtained. Then a simple stochastic approximation procedure adjusts the parameter θ at times kT via the gradient search algorithm:

$$\theta^{(k+1)} = \theta^{(k)} - \epsilon_k Z_k(\theta^{(k)}) \tag{2}$$

where ϵ_k is an appropriate sequence of gains for the stochastic approximation algorithm. This type of algorithm is natural for system optimization and is also usable for network design. See, for example, L'Ecuyer, Giroux, and Glynn (1990), and Suri, Leung (1991), where the adjustable parameters are related to the speed of the link or service time.

Algorithms of the form (2) are widely considered and their performance relies on the construction of good estimators of $\nabla_\theta L(\theta^{(k)})$ over the time interval $[kT, (k+1)T)$, where the control variable is fixed. This is the problem that we study.

The Performance Measure. Call $N(t)$ the number of customers in the system at time t, $a(t)$ the total number of arrivals to the system within $[0, t]$, T_k the sojourn time of the k^{th} customer, and $Q_\alpha(t)$ the queue length at node n_α at time t, including customers in service.

Under the stability assumptions of the network, the limits $N = \lim_{T\to\infty} \frac{1}{T}\int_0^T N(t)\,dt$, $T = \lim_{K\to\infty} \frac{1}{K}\sum_{k=1}^{K} T_k$ and $\lambda = \lim_{t\to\infty} \frac{a(t)}{t}$ exist and coincide with the stationary mathematical expectation of the corresponding quantities $N(t)$, T_k, and $a(t)/t$. Denote by Q_α the pathwise average queue length at node n_α:

$$Q_\alpha = \lim_{T\to\infty} \frac{1}{T}\int_0^T Q_\alpha(t)\,dt \tag{3}$$

It follows from Little's theorem that under general service disciplines $N = \lambda T$. By definition, we also have that $N(t) = \sum_\alpha Q_\alpha(t)$ for each t, therefore:

$$T = \frac{1}{\lambda}\sum_\alpha Q_\alpha \tag{4}$$

Typical optimization goals in queueing networks include the maximization of stationary throughputs, minimization of stationary sojourn time and minimization of rejection probabilities.

The total throughput λ is the stationary rate of customers that enter the network. In an open network with infinite buffers and no rejections this quantity is independent of how the routing is decided within the system, provided that the routing probabilities ensure stability. Therefore in this problem the only performance criterion mentioned above that depends on the routing variables is the stationary average sojourn time per customer.

Equation (4) decomposes this function as a separable function in terms of local quantities Q_α at each node. The decomposition approach of Rubinstein (1991) considers the sojourn times at the different queues to approximate the overall sojourn time per customer by local functions but the estimates are biased.

In queueing systems it is well known that the dynamics can be described with respect to the customers or with respect to time (Glynn, Whitt, 1989). In our approach we have replaced the customer dependent performance criterion by an equivalent time dependent performance criterion.

3 THE SURROGATE ESTIMATION APPROACH

In order to motivate the main idea of our method we look at (4). The performance function $L(\theta) = \lambda T$ is additive, so that $\nabla_\theta L(\theta) = \sum_\alpha \nabla_\theta Q_\alpha$. If the computation of the gradient of Q_α could be performed locally at each node then the controller would only need to add the different estimators in order to estimate the sensitivity of the performance function. This localization is done via a change of variables argument.

A Change of Variables. Suppose for the moment that the network is Jackson, then the outside arrival process at node n_α is Poisson and the service times are exponentially distributed. Since the routing is randomized, in stationary operation the total arrival process at each node is also Poisson with parameter r_α. Q_α depends only on the first moment of the arrival distribution,so that:

$$\frac{\partial Q_\gamma}{\partial \theta_{\alpha\beta}} = \frac{\partial Q_\gamma}{\partial r_\gamma}\left(\frac{\partial r_\gamma}{\partial \theta_{\alpha\beta}}\right) = \frac{\partial Q_\gamma}{\partial \lambda_\gamma}\left(\frac{\partial r_\gamma}{\partial \theta_{\alpha\beta}}\right) \qquad (5)$$

where we have used (1) and the assumption that there are no loops to get $\partial r_\gamma / \partial \lambda_\gamma = 1$.

The terms $\partial r_\gamma / \partial \theta_{\alpha\beta}$ represent a "weight factor" which can be evaluated from the rates λ_α and the current values of the routing parameters $\theta_{\alpha\beta}$. Let us define the weight factors $w^\gamma_{\alpha\beta} = \frac{\partial r_\gamma}{\partial \theta_{\alpha\beta}}$ and rewrite (5) as:

$$\frac{\partial Q_\gamma}{\partial \theta_{\alpha\beta}} = w^\gamma_{\alpha\beta}\frac{\partial Q_\gamma}{\partial r_\gamma} = w^\gamma_{\alpha\beta}\frac{\partial Q_\gamma}{\partial \lambda_\gamma} \qquad (5')$$

The formulas in (5') provide an intriguing and manageable approach to the problem of getting accurate estimates of $\partial Q_\gamma / \partial \theta_{\alpha\beta}$. We propose using (5') even for the non-Jackson case, where a closed formula for the cost function is usually impossible to obtain.

Notice that only local estimates are needed: each server computes the sensitivity of its own queue size with respect to its local input rates, so that the computation of the gradient is distributed. These local estimates are then weighted appropriately by the factors $w^\gamma_{\alpha\beta}$, which are calculated from the measured mean flow data to get the estimators of the desired derivatives in the left hand side of (5).

It is important to note that for the non-Jackson case the derivative with respect to r_k might not be generally well defined. Nontheless, we will use the change of variables argument of the surrogate estimation approach as a heuristic argument. The extension of the argument to the non-Jackson model will be justified by the quality of our simulation results.

Discussion. If the system were Jackson (see Bertsekas, Gallager, 1987, and Walrand, 1988), then a closed formula for the average steady state queue lengths is:

$$Q_\alpha = \left(\frac{r_\alpha}{\nu_\alpha - r_\alpha}\right) \qquad (6)$$

In Tsitsiklis, Bertsekas (1986), Cotton, Mason (1991) and Bertsekas, Gallager (1987) it is assumed that a formula similar to (6) holds even if the system is not of the Jackson type. In fact, standard optimization algorithms seek to minimize a cost criterion of the form $\sum_\alpha F(r_\alpha/[\nu_\alpha - r_\alpha])$ for some convex function $F(\cdot)$, subject to the constraints $r_\alpha < \nu_\alpha$. For the case of non-Jackson networks, the mean queue length is not usually of the form (6) and these approximations can yield poor estimates, as seen from our simulation results.

If we wanted to solve the problem directly without the surrogate estimation approach, then we would be estimating the derivatives with respect to the routing probabilities. The construction of good derivative estimators is very difficult. It is well known that IPA exhibits extremely good performance when it can be applied directly, but it cannot be used to estimate derivatives with respect to routing probabilities (Gong, 1988; Glasserman, 1991 and Ho, Cao, 1985). As far as we know, other proposed methods don't stand out for good performance, low variance

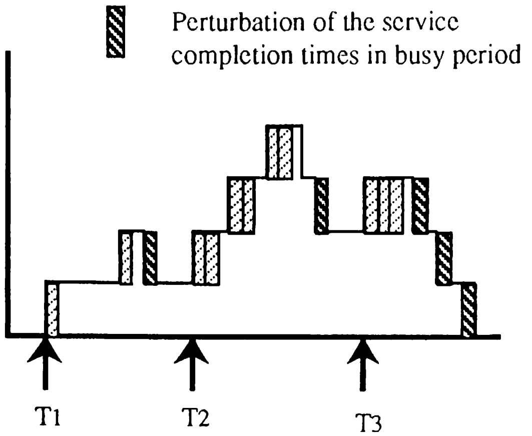

Figure 1a: External Arrival Starts Busy Cycle.

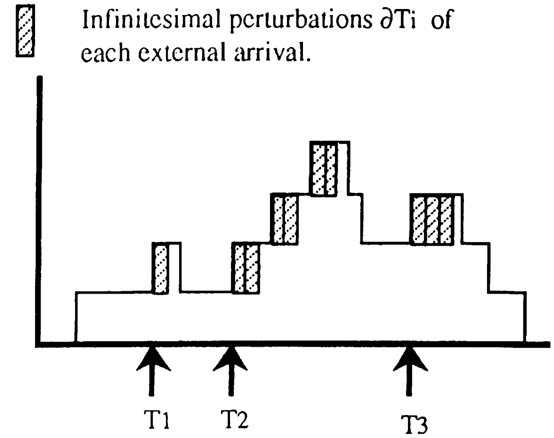

Figure1b: Arrival From Network Starts Busy Cycle.

and ease of implementation. Direct estimation in the general setting involves the construction of estimators of $\partial Q_\gamma/\partial\theta^d_{\alpha\beta}$ for all α,β and γ such that there is a path with destination d from α to γ using the link connecting n_α to n_β. Since the number of routing parameters is the number of outgoing links from each node for each destination in the network, the amount of data processing and memory required may be extremely large.

Using the change of variables for surrogate estimation, we have reduced the problem to that one of estimating the derivatives of $\partial Q_\gamma/\partial r_\gamma$ locally at each node. These estimates are then communicated to all controllers upstream, where they are multiplied by the appropriate weight factor. Since the weight factors involve only the mean rates and the knowledge of the other routing parameters, the computational effort in evaluating the different weight factors at each server is negligible. Therefore the amount of calculations for the estimators of $\partial Q_\gamma/\partial\theta^d_{\alpha\beta}$ grows with the number of nodes in the network, rather than with the product of the number of nodes by the number of links per destination. This is probably the most important property of the surrogate estimation approach for on-line operation.

4 IMPLEMENTATION OF IPA FOR THE EXAMPLE

We shall describe two IPA algorithms to estimate the local sensitivities using (5'): the first one estimates a finite horizon approximation of $\partial Q_\alpha/\partial\lambda_\alpha$. The second one estimates $\partial Q_\alpha/\partial r_\alpha$ and benefits from the regenerative structure of the local busy cycles to achieve lower bias and variance than the first one. Throughout this section we shall fix a server and drop the subscript α from our notation, since we need only work with one server at a time.

4.1 Sensitivity with respect to the External Arrival Rate. Since the external arrival sequences are Poisson, IPA can be applied. Define:

$$C_T(\lambda+\delta\lambda) \equiv \int_0^T Q_t(\lambda+\delta\lambda)\,dt, \quad \text{and}$$

$$\delta C_T = \frac{C_T(\lambda_0+\delta\lambda) - C_T(\lambda_0)}{\delta\lambda}$$

We estimate $\partial C_T(\lambda)/\partial\lambda$ as a finite horizon estimator of (3). It was shown in Vázquez-Abad, Kushner (1990) under general assumptions for the service time distributions and routing strategies that

$$E(\delta C_T(\lambda)) - \frac{\partial}{\partial\lambda}EC_T(\lambda)$$

For the construction of the estimate refer to Figure 1. Let N_m denote the number of external arrivals to the nominal queue during the m^{th} busy period and N'_m the total number of service completions. Let S_1 denote the class of busy periods on $[0,T)$ for which the first arrival is not an external arrival, and S_2 the complement class, for which the busy period starts with an external arrival. Call $K(T)$ the total number of busy periods within $(0,T]$.

Let τ_{m1} denote the time between the first external arrival during the m^{th} busy period and the previous external arrival to the chosen server. For $i>1$, let τ_{mi} denote the time between arrivals of the $(i-1)^{th}$ and i^{th} external arrivals in the m^{th} busy period. Then we have the following pathwise limit of δC_T:

$$Q^{(1)}(T) = \frac{1}{T}\left\{\sum_{m=1}^{K(T)}\sum_{i=1}^{N_m}(N_m - i + 1)\frac{\tau_{mi}}{\lambda_0} - \sum_{m\in S_2} N'_m\frac{\tau_{m1}}{\lambda_0}\right\} \tag{7}$$

This is a consistent estimate of $\frac{\partial}{\partial\lambda}\frac{1}{T}E\int_0^T Q_t(\lambda)\,dt$ that might be a biased estimate of (3). But the bias decreases as $T\to\infty$, Vázquez-Abad, Kushner (1990). However, as $T\to\infty$, the variance of the estimator increases due to the cumulative effects in (7) of the propagation of the perturbation from one busy period to the next. We shall describe in section 5 the implementation of $Q^{(1)}(T)$ for the network problem using two sets of parameters T for each node.

4.2 Sensitivity with respect to the Total Arrival Rate. We now estimate the derivative with respect to the total arrival rate r at the chosen server. Call τ_i the time between the arrival of customer i and $i+1$ at that node. To apply standard IPA to the whole interarrival sequence requires the distribution of this sequence, which is generally unknown. We use the assumption:

$$\frac{d\tau_i}{dr}=-\frac{\tau_i}{r} \tag{8}$$

as if we had an $M/G/1$ queue. This is, of course, a heuristic argument and in general (8) is not true. In a way, using (8) to estimate the local sensitivities together with (5') seems to be very similar to the Jackson approximation assumptions. As will be shown from our simulations, the results using pathwise estimation of the gradients via (5') and (10) below are very robust even for non-Jackson models, while the estimation via mean flow data and (6) can be very poor. We would like to acknowledge an anonymous referee who pointed out the heuristic argument for the application of IPA to this sequence.

When all arrivals are taken into account, the propagation of the perturbation of the first arrival to the busy cycle cancels the perturbations of the service time completions. Therefore the perturbations do not propagate from one busy period to another. An unbiased estimate for the steady state average can be constructed using a regenerative approach for the single queue (see also Glynn, L'Ecuyer, Adès, 1991) as follows:

$$\frac{d}{dr}\left[\lim_{T\to\infty}\frac{1}{T}\int_0^T Q(t)\,dt\right]=\frac{d}{dr}\frac{E\int_0^\tau Q(t)\,dt}{E\tau}=$$
$$=\frac{1}{E\tau}\frac{d}{dr}E\int_0^\tau Q(t)\,dt-\frac{E\int_0^\tau Q(t)\,dt}{E\tau}\left(\frac{1}{E\tau}\frac{d}{dr}E\tau\right) \tag{9}$$

where $\tau=\sum_{i=1}^N \tau_i$ is the duration of the first busy period. Call N the total number of customers within the first busy period and assume that the first customer arrives at time 0, then the epoch of arrival of customer i is $\sum_{j=1}^i \tau_j$. Call S_i the service requirement of customer i. Then the sojourn time of customer i is $X_i=\sum_{j=1}^i S_i-\sum_{j=1}^{i-1}\tau_j$, for $i=1,\ldots,N$. Therefore $\int_o^\tau Q(t)\,dt=\sum_{i=1}^N X_i$. Applying the standard IPA estimates we get:

$$\frac{d}{dr}\int_0^\tau Q(t)\,dt=\frac{1}{r}\sum_{i=1}^N\sum_{j=1}^{i-1}\tau_j,\quad\text{and}$$

$$\frac{d}{dr}\tau=-\frac{1}{r}\sum_{i=1}^N\tau_i=-\frac{\tau}{r}$$

As common in IPA, we exchange the derivatives and expectations in (9) to get:

$$\frac{d}{dr}\left[\lim_{T\to\infty}\frac{1}{T}\int_0^T Q(t)\,dt\right]=\frac{1}{r}\frac{E\sum_{i=1}^N\sum_{j=1}^{i-1}\tau_j}{E\tau}+\frac{1}{r}\frac{E\sum_{i=1}^N X_i}{E\tau}$$

The way to implement this algorithm is to estimate the expectations through sample averages, using K busy periods. Letting the superscript (m) denote the corresponding quantities for the m^{th} cycle, the estimator is:

$$\begin{aligned}Q^{(2)}(K)&=\frac{1}{r}\frac{\frac{1}{K}\sum_{m=1}^K\sum_{i=1}^N\sum_{j=1}^{i-1}\tau_j^{(m)}}{\frac{1}{K}\sum_{i=1}^K\tau^{(m)}}\\&+\frac{1}{r}\frac{\frac{1}{K}\sum_{m=1}^K\sum_{i=1}^N X_i^{(m)}}{\frac{1}{K}\sum_{m=1}^K\tau^{(m)}}\\&=\frac{1}{r}\frac{\sum_{m=1}^K\sum_{i=1}^N\sum_{j=1}^i S_j^{(m)}}{\sum_{m=1}^K\tau^{(m)}}\end{aligned} \tag{10}$$

Remark: We would like to point out that assumption (8) is actually not needed at all for the surrogate estimation: it is used only to implement standard IPA. In Vázquez-Abad, Kushner (1990) a different method was used to estimate an ersätz derivative. This method, now called the finite difference RPA method, does not require knowledge of the distribution of the interarrival times. We simulated various non-Jackson models (including non randomized routing strategies) and obtained comparable results in bias and variance for the regenerative IPA that we have introduced here and the RPA of Vázquez-Abad, Kushner (1990). The SPA method of Gong (1988), which in this case coincides with the RPA method of Brémaud, Vázquez-Abad (1992), could also be used to estimate the local sensitivities with respect to the total arrival rate. For purposes of illustration of the surrogate approach and limitations of space we have chosen to present here only the implementation of

IPA to the local sensitivities with respect to both the outside and the total arrival rates.

5 SIMULATION RESULTS

The network topology that we used for our simulations is depicted in Figure 2, where the $\{\theta_{\alpha\beta}\}$ are the stationary fractions of flows along the respective links.

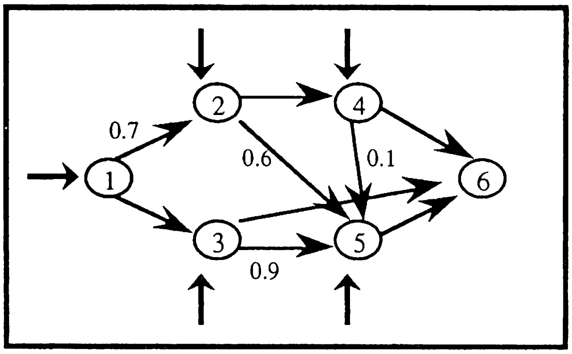

Figure 2: Network for Simulations

The external interarrival times at node n_α are mutually independent with exponential distribution and expectation λ_α^{-1}. The service times are mutually independent and, for each node they are identically distributed with mean ν_α^{-1}. In table 1 we show the network parameters. We estimated the sensitivities of the queue lengths with respect to a single routing parameter $\theta \equiv \theta_{12}$ only to illustrate the salient properties of the estimators. Thus, for the surrogate method, the weight factors are $\partial r_\alpha/\partial\theta$.

Table 1: Mean Network Statistics

α	λ_α	r_α	ν_α	$\partial r_\alpha/\partial\theta$
n_1	0.200	0.200	0.800	0.000
2	0.160	0.300	0.400	0.200
3	0.100	0.160	0.800	−0.200
4	0.300	0.420	0.600	0.080
5	0.150	0.516	0.600	−0.052
6	–	0.910	1.600	0.000

Table 3 shows the results of the surrogate estimation for the two IPA algorithms of the previous section in a Jackson network, where the service times are exponentially distributed. Table 4 presents the results for a non-Jacskon network (see Vázquez-Abad, Kushner, 1990 for the details), where the service time distributions for different parts of the network are be either heavy or light tailed. For the system simulated, nodes n_2 and n_3 are $M/G/1$ queues with randomized service times so the derivatives at these nodes can be evaluated analytically, but not so for nodes n_4 and n_5.

To get the distributed control algorithm we let each node n_α construct its local estimator during time periods of the form $[iT_\alpha, (i+1)T_\alpha)$, with possibly different T_α for different nodes. These estimates are sent to node n_1 at times $(i+1)T_\alpha$. Under a stochastic approximation procedure, θ would be updated according to (2), where Z_k would be constructed using the information from all other nodes received within the updating period $[kT,(k+1)T)$. We assume that $T \geq T_\alpha$ for all α and that the delays in information broadcasting are negligible. We emphasize that during the interval $[kT,(k+1)T)$ the parameter $\theta^{(k)}$ is fixed.

In this study we are concerned primarily with the assessment of the quality of the estimators Z_k that would be used in the sequential approximation procedure (2). To do this, we simulated the network with fixed parameters using a basic interval $T = 20,000$ units of time. The process was simulated for a total time NT with $N = 200$. From the additive form (4) of the performance function and the surrogate estimation formula (5') we can write:

$$Z_k = \sum_\alpha Z_k(\alpha) = \sum_\alpha \left(\frac{\partial r_\alpha}{\partial\theta}\right) D_\alpha(k).$$

$D_\alpha(k)$ is composed of the local sensitivity estimators obtained at times $iT_\alpha \in [kT,(k+1)T)$ and is given by (11) or (12) according to the case. For each of the three estimators described below and each network model, the simulations produced a sample of the estimators $\{D_k(\alpha), Z_k; k = 1,\ldots,N\}$ that we used to calculate the sample means and variances.

IPA(1) and IPA(2). The estimators IPA(1) and IPA(2) are based on formula (7). Each node n_α uses $Q^{(1)}(T_\alpha)$, where T_α is of the form $T/m(\alpha)$ for some integer $m(\alpha)$. Call $d_{ik}(\alpha)$ the estimator produced over the period of time $(k-1)T + (i-1)T_\alpha \leq t < (k-1)T + iT_\alpha$, for $i = 1,\ldots,m(\alpha)$. The statistic $D_k(\alpha)$ corresponding to the whole interval $[kT,(k+1)T)$ is the average:

$$D_k(\alpha) = \frac{1}{m(\alpha)} \sum_{i=1}^{m(\alpha)} d_{ik}(\alpha) = \frac{1}{T}\sum_{i=1}^{m(\alpha)} T_\alpha d_{ik}(\alpha) \quad (11)$$

Since the perturbations propagate from one busy period to the next, larger values of T_α reduce the bias but increase the variance of the the estimators.

The values chosen for the T_α are shown in table 2. IPA(2) uses intervals that are twice as large as those of IPA(1). The results of tables 3 and 4 show that the reduction in bias is barely noticeable in comparison with the growth in variance. Due to the averaging over the subintervals, the sample variances of IPA(1) and IPA(2) are actually smaller than those that would be obtained with $m(\alpha) = 1$.

Table 2: Simulation Parameters

Method	Description			
	T_2	T_3	T_4	T_5
IPA(1)	400	1000	1000	500
IPA(2)	800	2000	2000	1000

Remark: The formula (7) was calculated supposing that the server was idle at $t = T$. If the terminal time is in the middle of a busy period, we simply truncate the sums in $d_{ik}(\alpha)$ as follows. We use only those terms in the first sum corresponding to arrival moments prior to time $(k-1)T + iT_\alpha$. For the last busy period in this interval, we replace N_k^i by the total number of service completions in that busy period up to the terminal time. Analogously, if some customers in the queue at time $kT + (i-1)T_\alpha$ are external arrivals, then we include the associated terms in the sum.

IPA(R). The regenerative estimator (10) is also modified by truncating the numerator of the formula. Let $K_{ik}(T_\alpha)$ denote the number of busy periods that have finished within the local estimation interval $((k-1)T+iT_\alpha, (k-1)T+(i+1)T_\alpha]$. Then by (10) $d_{ik}(\alpha)$ is $Q^2[K_{ik}(T_\alpha)]$. If a busy period starts but does not finish within this interval, it is considered to be part of the following estimation period. Since perturbations do not propagate from one busy period to another in (10), the parameters T_α are irrelevant in the average (11) and:

$$D_k(\alpha) = \frac{1}{r}\frac{\sum_{m=m(k)}^{K(k,T)}\sum_{i=1}^{N}\sum_{j=1}^{i} S_j^{(m)}}{\sum_{m=m(k)}^{K(k,T)}\tau^{(m)}} = \frac{1}{r}\frac{\sum_{m=m(k)}^{K(k,T)}\sum_{i=1}^{N}\sum_{j=1}^{i} S_j^{(m)}}{T} \tag{12}$$

where $m(k)$ is the first busy period that finishes within $[kT, (k+1)T)$ and $K(k,T)$ is the number of local busy cycles that have finished within this time interval, so that $\sum_{m=m(k)}^{K(k,T)}\tau^{(m)} \approx T$. Although we did not make this clear in the notation, $m(k)$ and $K(k,T)$ depend on α. We call IPA(R) the corresponding estimator. Tables 3 and 4 show the difference between the estimators IPA(1) and IPA(2), and the much better estimator IPA(R). We simulated other service distributions and routing strategies and obtained consistently very good estimators with IPA(R).

Remark: In practice, the end effects caused by the truncations are generally of little importance. Typically, a sequence of estimates is taken over a sequence of successive time intervals, say, $[kT, kT+T)$, and one can start the k^{th} estimation at the start of the first busy period of the chosen server after time kT.

Tables 3 and 4 give the sample means and variances (in parenthesis) of the estimators $D_k(\alpha)$ for each of the three estimators. The results for the surrogate estimator $Z_k(\alpha)$ of the derivatives with respect to θ are shown for the IPA(R) only and appear under the column labeled S-IPA. The surrogate estimators using IPA(1) or IPA(2) were very poor in comparison to the estimators using S-IPA and were not tabulated.

The difference between the Jackson and the non-Jackson systems is clear from the tables. Estimation of ν_α and r_α can be done very accurately using the sample means. If we used the mean flow approach with the approximation (6) for the non-Jackson network, we would estimate $\partial Q_\alpha/\partial\theta$ with the quantities under column 7 of table 3. This approximation underestimates the derivatives with respect to θ.

Table 3: Jackson Network

α	Q_α	$\frac{\partial Q_\alpha}{\partial r_\alpha}$	IPA(1)	IPA(2)	IPA(R)	$\frac{\partial Q_\alpha}{\partial\theta}$	S-IPA
2	3.00	40.00	32.9	35.8	43.7	8.00	8.74
			(96.9)	(474.8)	(70.4)		(2.81)
3	0.25	1.95	1.58	2.24	1.95	−0.39	−0.391
			(36.1)	(148.2)	(0.0019)		(0.00)
4	2.33	18.51	17.6	18.4	18.4	1.48	1.47
			(205.3)	(824.2)	(3.62)		(0.023)
5	6.14	85.03	60.0	68.2	82.3	−4.42	−4.282
			(535.7)	(3140)	(353.9)		(0.957)

For the non-Jackson model we do not have closed expressions for the queue lengths at nodes n_4 and n_5. We estimated the finite differences through seven long-run simulations of length NT each. The first simulation was performed at the parameter values of table 1 to evaluate Q_α. The others were performed changing one parameter value at a time, for $\theta \pm (\Delta)\theta$, $\lambda_4 \pm (\Delta)\lambda_4$ and $\lambda_5 \pm (\Delta)\lambda_5$. We used $\Delta = 0.05$ and calculated the sample queue lengths. The values appearing in italics in table 4 below are the finite difference estimates. The other values in the columns for Q_α and $\partial Q_\alpha/\partial\lambda_\alpha$ and those for $\partial Q_\alpha/\partial\theta$ are exact.

Table 4: Non-Jackson Network

α	Q_α	$\frac{\partial Q_\alpha}{\partial \lambda_\alpha}$	IPA(1)	IPA(2)	IPA(R)	$\frac{\partial Q_\alpha}{\partial \theta}$	S-IPA
2	4.4	63.4	48.3	55.3	70.1	12.6	14.1
			(114.4)	(599.3)	(185.44)		(7.41)
3	0.21	1.75	1.79	2.05	1.60	−0.35	−0.32
			(41.1)	(162.1)	(0.00)		(0.00)
4	*3.44*	*29.18*	28.7	29.3	29.6	*2.57*	2.36
			(337)	(1484)	(16.4)		(0.11)
5	*5.05*	*66.04*	50.4	53.6	63.1	*−2.9*	−3.2
			(540)	(2895)	(177)		(0.48)

In order to estimate the derivatives of the queue lengths with respect to other routing parameters, we could use the same IPA estimators with the corresponding weight factors. Since the basic information is anyway broadcasted between the nodes, and since most of the computational effort is spent in the derivative estimation, the savings of the surrogate approach are more dramatic as the size of the network, the possible destinations and the number of outgoing links increase.

6 CONCLUDING REMARKS

Knowledge of the gradient with respect to a control variable can be used in sequential optimization procedures for on-line performance optimization as well as network design. However, in most applications where the system can be modeled as a queueing network, there is no closed expression for the performance measure. Therefore good estimators of the gradients are extremely important. The adaptive control of highly decentralized systems requires single path estimation methods: otherwise the system itself would be operating at different values of the control parameters in order to estimate the desired gradients. When more than one parameter is being updated under a decentralized operation, the combined effects result in extremely large variances.

A common approach to solve this problem is to approximate the solution via an independence assumption and use mean flow data in the corresponding expressions. This approach, known as the Jackson network approximation, may yield very poor estimates of the required gradients, as shown in our simulations.

We have discussed a method for the estimation of such gradients that we believe to be applicable for a broad class of problems. The surrogate estimation approach uses the system's dynamics and heuristic relations among the system's parameters to distribute the computation. Local estimators are then implemented for single queues, a problem which is much easier to solve. The amount of data processing and memory required by direct estimation may be extremely large. The savings in computation time due to surrogate estimation are actually more dramatic as the complexity of the problem increases. This is because the *same* local estimates are used to calculate the surrogate estimator of the sensitivity with respect to all the routing parameters.

The local estimation can in principle be achieved using any of the available methods for sensitivity estimation, many of which work very well in the context of local estimation. We have presented in particular the implementation of IPA for the local sensitivities. Although a mean flow argument is used in the construction of the surrogate estimators, our simulation results indicate that our method is very promising and robust under more general networks than the Jackson.

Future research goals involve the extension of the ideas of surrogate estimation to other classes of problems, including problems where the performance measure is not separable, the control variable is integer-valued and the network has customer classes with different service priorities.

ACKNOWLEDGEMENTS

The work of the first author was partially supported by grant number WFA0139015 through the Women in Faculty Positions Programme of the Natural Sciences and Engineering Research Council of Canada. The work of the second author was partially supported by grant number ARO DAAL-03-92-G-0115.

REFERENCES

Bertsekas, D., Gallager, R. 1987. *Data Networks.* New Jersey: Prentice-Hall.

Brémaud, P., Vázquez-Abad, F. 1992. On the pathwise computation of derivatives with respect to the rate of a point process: the phantom RPA method, *QUESTA*, **10**:249-270.

Cotton, M., Mason, L. 1991. Adaptive isarithmic flow control in fast packet switching networks – heavy traffic case, *ITC Specialist Seminar*, Cracow, Poland, april 1991.

Glasserman, P. 1991 *Gradient estimation via perturbation analysis.* Hingham, Mass: Kluwer Academic publishers.

Glynn, P., L'Ecuyer, P., Adès. M. 1991 Gradient estimation for ratios, in *Proc. 1991 Winter Simulation Conf.*, B. Nelson, W.D. Kelton, G. M. Clark, editors: 986-993.

Gong, W.B. 1988. Smoothed perturbation analysis algorithm for a $G/G/1$ Routing Problem", in *Proc. of the 1988 Winter Simulation Conf.*, M. Abrams, P. Haigh, J. Comfort, editors: 525-531.

Ho, Y. C., Cao, X. 1985 Performance sensitivity to routing changes in queueing networks and flexible manufacturing systems using perturbation analysis, *IEEE Jour. Robotics and Autom.*, **RA-1**, 4: 165-172.

L'Ecuyer, P. 1991 An overview of derivative estimation, in *Proc. of the 1991 Winter Simulation Conf.*, B. Nelson, W.D.Kelton, G. M. Clark, editors: 207-217.

L'Ecuyer, P., Giroux, N., Glynn, P. 1990 Stochastic optimization by simulation: some experiments with a simple steady-state queue, département d'informatique, Université Laval, working paper.

Rubinstein, R. 1991 Decomposable score function estimators for sensitivity analysis and optimization of queueing networks, to appear in *Annals of Oper. Res.*.

Sigman, K. 1990 The stability of open queueing networks, *Stoch. Proc. and their applications.*

Suri, R., Leung, Y.T. 1991 Single run optimization of discrete event simulations: an empirical study using the $M/M/1$ queue, in *Proc. of IIE*, **21**: 35-49.

Tsitsiklis, J., Bertsekas, D. 1986. Distributed asynchronous optimal routing in data networks, *IEEE Trans. Autom. Control*, **AC-31**, 4: 325-332.

Vázquez-Abad, F., Kushner, H. 1990 A surrogate estimation approach for adaptive routing in communication networks, Brown University Report LCDS /CCS #90-2, May 24 1990.

Walrand, J., 1988 *An Introduction to Queueing Networks* New Jersey: Prentice Hall.

AUTHOR BIOGRAPHIES

FELISA J. VÁZQUEZ-ABAD is Professor at the Department of Computer Science and Operations Research at the University of Montreal. She received her B.Sc. in Physics in 1983 and her M.Sc. in Statistics and Operations Research in 1984 from Universidad Nacional Autónoma de México. She received her Ph.D. in Applied Mathematics from Brown University in 1989. Her research interests include the optimization of stochastic processes, adaptive parallel control of DEDS, sensitivity estimation and simulation of stochastic networks. She was a researcher at INRS-Télécommunications from 1990 to 1993, where she acquired some experience in the application of on-line stochastic control to problems in telecommunications. Since June 1993 she is the recipient of an NSERC Women's Faculty Award.

HAROLD J. KUSHNER is a University Professor and Professor of Applied Mathematics and Engineering at Brown University. He received his Ph.D. in Electrical Engineering from the Univerisity of Wisconsin in 1958. He has worked on virtually all aspects of stochastic systems theory and was the recipient of the *IEEE* Control Systems Field Award in 1992.