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

TWO APPROACHES FOR ESTIMATING THE GRADIENT IN FUNCTIONAL FORM

Pierre L'Ecuyer

Département d'IRO
Université de Montréal, C.P. 6128, Succ. A
Montréal, H3C 3J7, CANADA

ABSTRACT

Consider a stochastic model for which the performance measure is defined as a mathematical expectation which depends on a parameter θ. By using a likelihood ratio (i.e., a change of measure), it is often possible to construct an estimator of the performance measure in functional form, i.e., given as a function of θ, and computed from a single simulation run. It is also possible to obtain in functional form an estimator of the gradient with respect to θ. On way of doing that is to combine the likelihood ratio technique with a score function gradient estimator; another way is to combine it with a perturbation analysis gradient estimator. We compare and illustrate those two approaches.

1. LIKELIHOOD RATIOS AND ESTIMATORS IN "FUNCTIONAL" FORM

Consider a stochastic model defined over a probability space $(\Omega, \Sigma, P_\theta)$, where the probability law P_θ is parameterized by $\theta \in \Theta$, and Θ is some bounded subset of $\mathbb{R}^d$. Let the random variable $h(\theta, \omega)$ represent the (finite-horizon) sample "cost" at parameter level θ for a sample point $\omega \in \Omega$. The expected cost at parameter level θ is then

$$\alpha(\theta) = \int_\Omega h(\theta, \omega) dP_\theta(\omega).$$

We want to estimate the gradient $\alpha'(\theta) = d\alpha(\theta)/d\theta$, which is assumed to exist. More specifically, we are seeking an "functional" estimator $\psi(\cdot, \omega)$ of $\alpha'(\cdot)$ which can be obtained from a single simulation run. Here and throughout the paper, the prime denotes the gradient with respect to θ.

Under appropriate conditions, such an estimator can be obtained as follows (see Glynn 1990, L'Ecuyer 1990, Rubinstein and Shapiro 1993). Select a probability measure G that dominates all the P_θ's. Then, α can be rewritten as

$$\begin{aligned}\alpha(\theta) &= \int_\Omega h(\theta, \omega) dP_\theta(\omega) \\ &= \int_\Omega h(\theta, \omega) L(G, \theta, \omega) dG(\omega),\end{aligned}$$

where $L(G, \theta, \omega) = (dP_\theta/dG)(\omega)$ is a *likelihood ratio*. Further, under appropriate regularity (uniform integrability) conditions (see L'Ecuyer 1993 or Rubinstein and Shapiro 1993):

$$\alpha'(\theta) = \int_\Omega \psi(\theta, \omega) dG(\omega), \tag{1}$$

where

$$\begin{aligned}&\psi(\theta, \omega) \\ &= L(G, \theta, \omega) h'(\theta, \omega) + h(\theta, \omega) L'(G, \theta, \omega) \\ &= [h'(\theta, \omega) + h(\theta, \omega) S(\theta, \omega)] L(G, \theta, \omega) \qquad (2)\end{aligned}$$

and

$$S(\theta, \omega) = \frac{L'(G, \theta, \omega)}{L(G, \theta, \omega)} = \frac{\partial}{\partial\theta} \ln L(G, \theta, \omega)$$

is called the *score function* (SF).

If $h(\theta, \omega) \equiv h(\omega)$ does not depend explicitly on θ and P_θ and G have densities f_θ and g, then

$$\begin{aligned}L(G, \theta, \omega) &= f_\theta(\omega)/g(\omega), \\ S(\theta, \omega) &= \frac{\partial}{\partial\theta} \ln f_\theta(\omega),\end{aligned}$$

and (2) becomes

$$\psi(\theta, \omega) = h(\omega) \frac{f_\theta(\omega)}{g(\omega)} \frac{\partial}{\partial\theta} \ln f_\theta(\omega). \tag{3}$$

From a simulation using the density g, one can estimate α and α' all over Θ using (3): ω is generated from g, $h(\omega)$ is computed (during the simulation), and then $\psi(\theta, \omega)$ can be computed by computing the remaining factor at any value of θ of interest. This technique is analyzed in great detail in Rubinstein and Shapiro (1993), where many examples are also given. These authors call the problem of estimating α and α' all over Θ the "What-if" problem.

Note that the score function is what permits one to estimate the gradient $\alpha'(\theta)$ instead of the performance measure $\alpha(\theta)$ itself, while the likelihood ratio is what permits one to estimate α and α' all over Θ (i.e., to solve the "what-if" problem) by a single simulation. For that reason, we shall call this technique the likelihood ratio (LR) method. Estimating expectations in a functional form using a likelihood ratio is also discussed in Glynn and Iglehart (1989), Reiman and Weiss (1989), Rubinstein (1991), and several other papers about *importance sampling* and change of measure.

Of course, in practice, multiple independent replications of the simulation must be performed to reduce the variance of the estimator and compute confidence intervals. One then uses the gradient estimator

$$\bar{\psi}_N(\theta,\omega) = \frac{1}{N}\sum_{j=1}^{N}\psi_j(\theta,\omega), \qquad (4)$$

where $\psi_1(\theta,\omega),\ldots,\psi_N(\theta,\omega)$ are i.i.d. replicates of $\psi(\theta,\omega)$. Confidence intervals can be computed as usual, using the central-limit theorem.

One important application of the LR method is optimization (Rubinstein and Shapiro 1993). Suppose one wishes to minimize $\alpha(\theta)$ with respect to θ. One way of doing that is (roughly) to obtain an estimation of α in a functional form, say from N i.i.d. replications, and then take the optimizer $\hat{\theta}_N$ of the average sample performance measure as an estimator of the optimizer θ_* of the expectation. Under appropriate conditions, it can be shown that $\hat{\theta}_N$ converges towards θ_* at rate $O(N^{-1/2})$ (see Rubinstein and Shapiro 1993). This optimization approach is called the *stochastic counterpart* (SC) method. A variant of it (which is equivalent if the sample gradient estimator is unimodal) is to estimate the gradient in a functional form, find a zero of that gradient estimator, and use it as an estimator of θ_*.

Another approach for gradient estimation is *infinitesimal perturbation analysis* (IPA); see Glasserman (1991). The idea is to define the probability space (Ω,Σ,P_θ) and the random variable $h(\theta,\omega)$ in such a way that only h depends on θ, i.e., the probability measure $P_\theta \equiv P$ is independent of θ. In that case, one can take $G \equiv P$ and the gradient estimator (2) becomes the sample derivative

$$\psi(\theta,\omega) = h'(\theta,\omega). \qquad (5)$$

In several situations, for a given model, it is possible to formulate the problem either in the IPA context (P_θ independent of θ) or in the SF-LR context (with $h(\cdot,\omega)$ independent of θ) by selecting the meaning of ω and defining the probability space appropriately (L'Ecuyer 1990). We will examine an example of that in the next section.

Typically, gradient estimators based on IPA have less variance than those based on SF as in (3). However, (5) contains no likelihood ratio and does not permit (in general) to estimate the gradient in a functional form, because for a given ω, $h(\theta,\omega)$ is typically a very complicated function of θ and the best way to evaluate it at different values of θ is usually to rerun the simulation at each of those different values. Therefore, (5) does not appear very convenient, at first sight, for the SC optimization method.

To combine the IPA estimator (5) with LR, one must *change* the interpretation of ω *after* (5) has been defined. More specifically, when the estimator (5) is derived, the probability space is defined in such a way that $P_\theta \equiv P$ is independent of θ. Afterwards, the random variable $\psi(\theta,\omega)$ can be redefined over a different probability space (i.e., using a different representation), for which P_θ now depends on θ. Assuming that this can be done in such a way that $\psi(\cdot,\omega)$ itself can now be computed in a functional form from a single simulation, one then obtains a gradient estimator in a functional form by multiplying $\psi(\theta,\omega)$ by the appropriate likelihood ratio. In practice, this can be achieved if $\psi(\theta,\omega)$ (under the redefined probability space) depends on θ in a way that is not too complicated. We shall call this functional estimator the IPA-LR gradient estimator and the estimator (3) the SF-LR gradient estimator.

In the next section, we will illustrate the construction of IPA-LR and SF-LR estimators through a $GI/M/1$ queueing example. We shall compare the variance properties of both estimators for that example, give numerical illustrations, and discuss the generalization of those properties to other situations in the following section and in the conclusion.

2. EXAMPLE: A $GI/M/1$ QUEUE

Let $h(\theta,\omega)$ be the average sojourn time for the first t customers in a GI/M/1 queue, initially empty, with mean service time θ. For each i, let S_i, W_i and $X_i = W_i + S_i$ represent the service time, waiting time, and sojourn time of customer i, respectively, and A_i be the inter-arrival time between customers i and $i+1$, whose distribution is assumed independent of θ. The sample performance measure is

$$h(\theta,\omega) = \frac{1}{t}\sum_{i=1}^{t} X_i$$

and we want to estimate $\alpha'(\theta)$, the derivative of the expectation of $h(\theta,\omega)$, with respect to θ.

2.1. The SF-LR Derivative Estimator

Suppose that ω represents $(S_1, A_1, \ldots, A_{t-1}, S_t)$, so that $dP_\theta(\omega)/d\omega$ corresponds to the joint density of $S_1, A_1, \ldots, A_{t-1}, S_t$. In that case, for a fixed ω, $h(\theta,\omega) = h(\omega)$ does not depend on θ (only its distribution does). Let $G = P_{\theta_0}$ for some fixed parameter value $\theta_0 > 0$. Then, in the actual simulation, the service times are generated from the exponential distribution with mean θ_0, the likelihood ratio and score function become

$$L(G,\theta,\omega) = \left(\frac{\theta_0}{\theta}\right)^t \exp\left(\left(\frac{1}{\theta_0} - \frac{1}{\theta}\right)\sum_{i=1}^{t} S_i\right) \quad (6)$$

and

$$S(\theta,\omega) = -\frac{t}{\theta} + \frac{1}{\theta^2}\sum_{i=1}^{t} S_i,$$

respectively, and the SF-LR estimator of $\alpha'(\theta)$ is (3). To estimate the derivative in a functional form, one simulates at θ_0 and memorizes $h(\omega)$ and $\sum_{i=1}^{t} S_i$. From that, $\psi(\theta,\omega)$ can then be evaluated at any other value of θ of interest. To do that, if N replications are performed, then the values of $h(\omega)$ and $\sum_{i=1}^{t} S_i$ for *each replication* must be memorized. It is important to note that the computational work involved in evaluating $\bar{\psi}_N(\theta,\omega)$ at each value of θ of interest is *not* negligible, especially if N is large. However, it is less than re-running the simulation at all those different values of θ.

One might wonder about the variability of such an estimation scheme, especially when θ is far away from θ_0. For that particular example, it can be shown (see L'Ecuyer 1993) that the kth moment of the derivative estimator $\psi(\theta,\omega)$ is finite if and only if $(k-1)\theta < k\theta_0$. In particular, for a given $\theta_0 > 0$, $\psi(\theta,\omega)$ has finite expectation for all $\theta > 0$, but finite variance only for $\theta < 2\theta_0$. One can easily derive the following more precise expression for the variance from a development similar to that outlined in Section 2.3 of Rubinstein and Shapiro (1993):

$$\begin{aligned}&\mathrm{Var}\,[\psi(\theta,\omega)]\\ &= \left(\frac{\theta_0^2}{\theta(2\theta_0-\theta)}\right)^t E_{\tilde{\theta}}\left[h^2(\omega)\left(\sum_{i=1}^{t}\frac{S_i-\theta}{\theta^2}\right)^2\right]\\ &\quad - [\alpha'(\theta)]^2, \end{aligned} \quad (7)$$

where $\tilde{\theta} = \theta\theta_0/(2\theta_0 - \theta)$. When $\theta = \theta_0$, one has $\tilde{\theta} = \theta$, the expectation in (7) increases linearly in t, and the factor multiplying the expectation is equal to 1; therefore the variance is in $O(t)$. For $\theta \neq \theta_0$, the expectation is in $O(t^2)$, and so the variance increases as $O\left(\left[\theta_0^2/(\theta(2\theta_0-\theta))\right]^t t^2\right)$ as a function of t, that is, exponentially fast. Furthermore, for fixed t, the variance increases to infinity when θ approaches either 0 or $2\theta_0$.

2.2. The IPA-LR Estimator

To derive the IPA estimator of $\alpha'(\theta)$, we now reinterpret ω as representing the sequence of i.i.d. uniforms that are used to drive the simulation. If the service times are generated by inversion, one has $S_j = \theta Z_j = -\theta \ln(1-U_j)$, where U_j is the uniform variate used to generate S_j, and Z_j is an exponential random variable with mean 1. Then,

$$h'(\theta,\omega) = \frac{1}{t}\sum_{i=1}^{t}\sum_{j=\nu_i}^{i} Z_j \quad (8)$$

$$= \frac{1}{t}\sum_{i=1}^{t}\sum_{j=\nu_i}^{i} S_j/\theta \quad (9)$$

is an unbiased estimator of $\alpha'(\theta)$, where ν_i is the number (or index) of the first customer who is in the same busy period as customer i.

Now, to estimate the derivative everywhere by simulating at θ_0, we will *reinterpret* ω as representing the sequence of interarrival and service times. For that, $h'(\theta,\omega)$ must be expressed as a function of this "new" ω; that is, we must use the expression (9), not (8). With this new interpretation of ω, the likelihood ratio becomes (6), the same as for the SF-LR method. Then, the IPA-LR derivative estimator at θ when the simulation is performed at θ_0 is

$$\begin{aligned}&\psi(\theta,\omega)\\ &= \left(\frac{1}{\theta t}\sum_{i=1}^{t}\sum_{j=\nu_i}^{i} S_j\right)\left(\frac{\theta_0}{\theta}\right)^t \exp\left(\left(\frac{1}{\theta_0}-\frac{1}{\theta}\right)\sum_{i=1}^{t} S_i\right).\end{aligned}$$

If N replications are performed, then the values of $\sum_{i=1}^{t} S_i$ and $\sum_{i=1}^{t}\sum_{j=\nu_i}^{i} S_j$ for each replication must be memorized in order to make possible the subsequent evaluation of $\bar{\psi}_N(\cdot,\omega)$.

In this case, the variance of the derivative estimator is given by

$$\begin{aligned}&\mathrm{Var}\,[\psi(\theta,\omega)]\\ &= \left(\frac{\theta_0^2}{\theta(2\theta_0-\theta)}\right)^t E_{\tilde{\theta}}\left[\left(\frac{1}{\theta t}\sum_{i=1}^{t}\sum_{j=\nu_i}^{i} S_j\right)^2\right]\end{aligned}$$

$$- [\alpha'(\theta)]^2, \tag{10}$$

where $\tilde{\theta} = \theta\theta_0/(2\theta_0 - \theta)$. One can see that this expression is in $O\left((\theta_0^2/(\theta(2\theta_0 - \theta)))^t t^{-1}\right)$ as a function of t. So, the variance increases exponentially fast w.r.t. t for $\theta \neq \theta_0$, and *decreases* linearly in t for $\theta = \theta_0$. As a result, for θ close to θ_0, this IPA-LR estimator should be much better than its SF-LR counterpart for moderate or large t, since its variance decreases w.r.t. t instead of increasing. The variance also increases to infinity when θ approaches either 0 or $2\theta_0$, as was the case with SF-LR.

2.3. Numerical illustrations

For numerical illustrations, consider an $M/M/1$ queue with arrival rate $\lambda = 1$ and mean service time θ. We have performed experiments with different values of t and θ_0, each time doing $N = 10,000$ replications to estimate α' all over the interval $\Theta = [0.1, 0.9]$, using both the SF-LR and IPA-LR methods described in the previous subsections. Figures 1–6 show some of the results.

In the upper part of each figure, the solid line gives the exact value of $\alpha'(\theta)$ as a function of θ, while the two dotted lines show the boundaries of a 95% confidence interval for $\alpha'(\theta)$ computed from our N simulation replications. The dashed line (the center of the confidence interval) indicates the value of the derivative estimator $\bar{\psi}_N(\cdot, \omega)$.

In the lower part of each figure, the solid line is an estimation of the variance of $\psi(\theta, \omega)$ computed using (7) and (10), i.e., the expectation $E_{\tilde{\theta}}[\cdot]$ was estimated by the simulation and the other terms were computed exactly. The dashed line, on the other hand, represents the standard empirical variance, given by

$$\frac{1}{N-1} \sum_{j=1}^{N} \left(\psi_j(\theta, \omega) - \bar{\psi}_N(\theta, \omega)\right)^2.$$

The latter is in fact much less reliable as a variance estimator; indeed, it can be shown that its variance is finite only for $\theta < 4\theta_0/3$ and increases to infinity as θ approaches either 0 or $4\theta_0/3$ (the variance of the variance estimator has to do with the fourth moment of $\psi(\theta, \omega)$).

Figures 1–4 show the behavior of the two types of estimators for $t = 10$, and $\theta_0 = 0.5$ and 0.8. It can be seen that the IPA-LR method has much lower variance, and that both methods are reliable only in a restricted region around (and mostly to the left of) θ_0. Asmussen and Rubinstein (1991), Rubinstein (1991), and Rubinstein and Shapiro (1993) have already pointed out the variance explosion with the SF-LR method when θ becomes too large, and have advocated selecting θ_0 equal (or slightly larger than) the largest value of θ at which one is interested in estimating the gradient. It turns out that the same kind of variance explosion occurs with the IPA-LR estimators. The figures also illustrate the fact that for both methods, there is a (perhaps more sneaky) problem when θ gets too small: the empirical variance is close to zero with very high probability, despite the fact

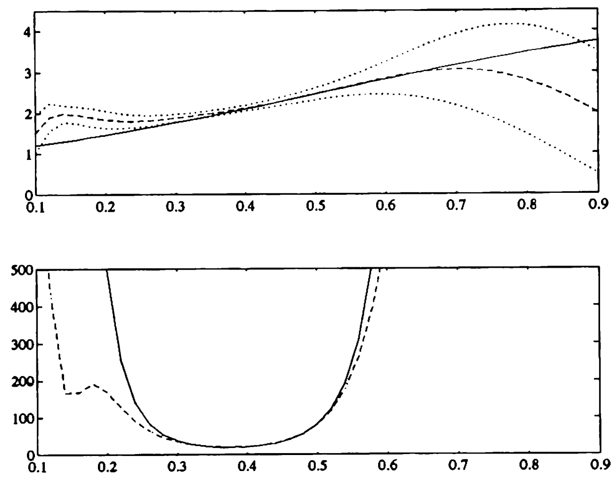

Figure 1: SF-LR Estimator, $t = 10$, $\theta_0 = 0.5$

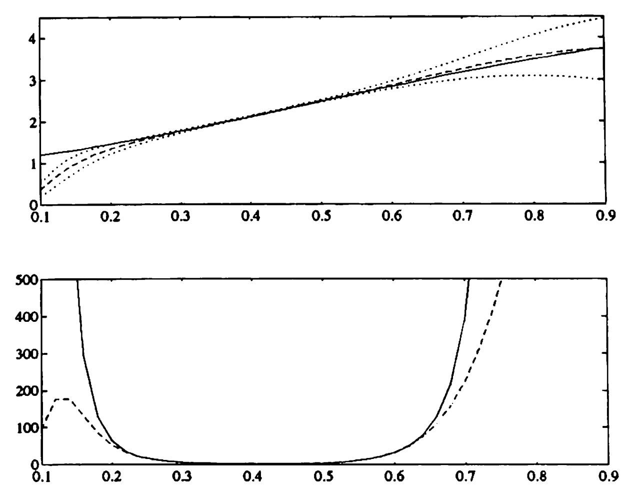

Figure 2: IPA-LR Estimator, $t = 10$, $\theta_0 = 0.5$

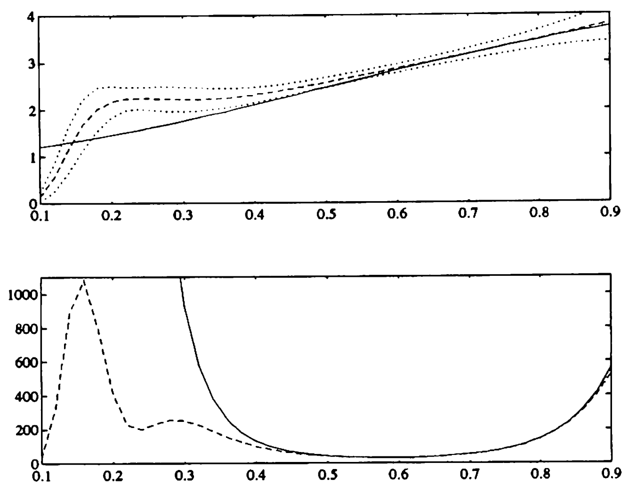

Figure 3: SF-LR Estimator, $t = 10$, $\theta_0 = 0.8$

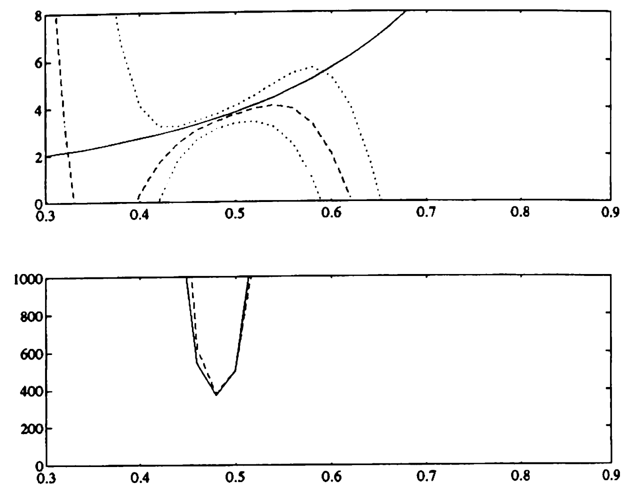

Figure 5: SF-LR Estimator, $t = 100$, $\theta_0 = 0.5$

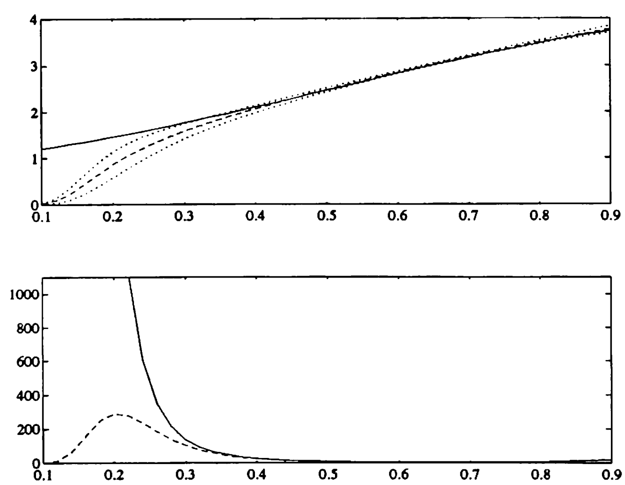

Figure 4: IPA-LR Estimator, $t = 10$, $\theta_0 = 0.8$

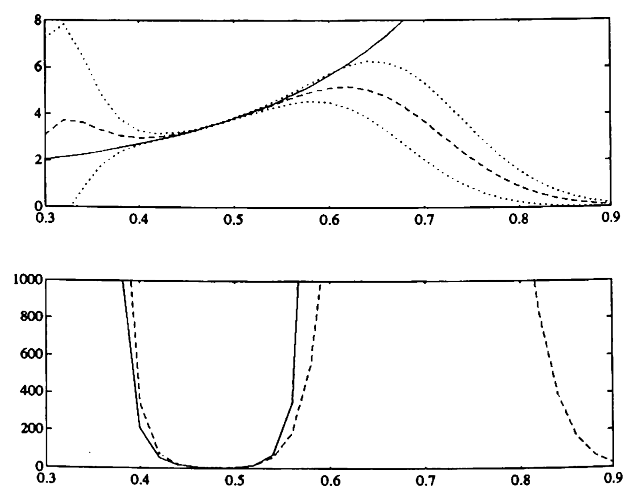

Figure 6: IPA-LR Estimator, $t = 100$, $\theta_0 = 0.5$

that the true variance is huge. So, using the empirical variance to compute confidence intervals becomes misleading. This is what happens, for example, when $\theta_0 = 0.8$ and $\theta < 0.3$.

Figures 5–6 show what happens with $\theta_0 = 0.5$ and $t = 100$. They show that as t increases, the estimators can be useful only in a very narrow area. Roughly, in that case, the SF-LR estimators are useful only for $\theta \in (0.45, 0.55)$, while the IPA-LR estimators are useful only for $\theta \in (0.4, 0.6)$. This, despite the fact that 10,000 replications are performed.

3. OTHER EXAMPLES AND APPLICATIONS

3.1. Fraction of Zero-Wait Customers in a $GI/M/1$ Queue

We consider again the $GI/M/1$ example, but we will now redefine h as the fraction of customers whose waiting time is zero, among the first t customers. That is,

$$\begin{aligned} h(\theta,\omega) &= \frac{1}{t}\sum_{i=1}^{t} I[W_i = 0] \\ &= \frac{1}{t}\sum_{i=1}^{t} I[A_i \geq X_{i-1}], \end{aligned} \quad (11)$$

where I denotes the indicator function. Let F be the interarrival-times distribution, with density f. If we differentiate the estimator (11) with respect to θ, we obtain zero with probability one, and so the sample derivative is not a useful estimator of $\alpha'(\theta)$. However, by conditioning on X_{i-1} for each term in the sum, we obtain the conditional Monte Carlo estimator

$$\begin{aligned} h(\theta,\omega) &= \frac{1}{t}\sum_{i=1}^{t} P[A_i \geq X_{i-1} \mid X_{i-1}] \\ &= \frac{1}{t}\sum_{i=1}^{t}(1 - F(X_{i-1})), \end{aligned} \quad (12)$$

whose expectation is also $\alpha(\theta)$. Then,

$$\begin{aligned} h'(\theta,\omega) &= -\frac{1}{t}\sum_{i=1}^{t} f(X_{i-1})X'_{i-1} \\ &= -\frac{1}{t}\sum_{i=1}^{t} f(X_{i-1}) \sum_{k=\nu_{i-1}}^{i-1} S'_k \end{aligned}$$

which can be multiplied by the likelihood ratio as in the previous example to get the IPA-LR derivative estimator. The technique that we just used to obtain the latter IPA estimator is called *smoothed perturbation analysis* (see Glasserman 1991). The SF-LR estimator is the same as in the previous example, with the only exception that h must be replaced by one of the two expressions (11) and (12).

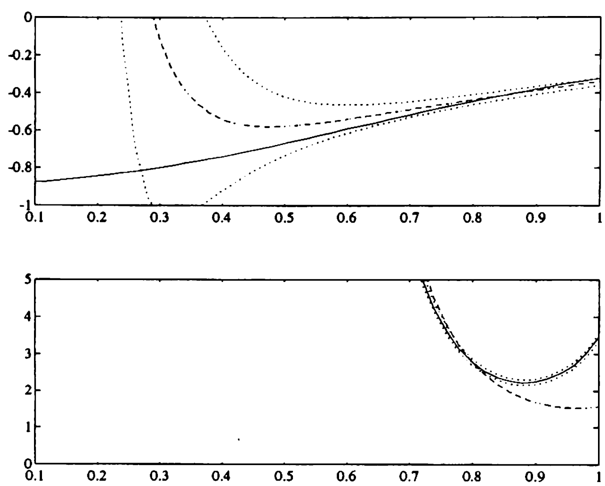

Figure 7: SF-LR Estimator, $t = 10$, $\theta_0 = 0.8$

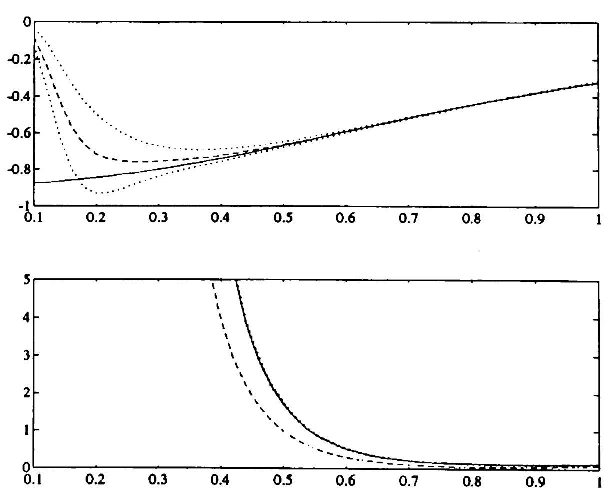

Figure 8: IPA-LR Estimator, $t = 10$, $\theta_0 = 0.8$

Figures 7–8 show the results of some numerical experiments made with this example., with $t = 10$ and $\theta_0 = 0.8$. For the results shown for SF-LR, we used the expression (11) for h, but there was very little difference in the results when we used (12) instead. It could be observed that for this example, the most favorable variance properties are obtained when θ is slightly larger than θ_0. In the previous example, the opposite was true. This means that to estimate $\alpha'(\theta)$ in some interval $[\theta_1, \theta_2]$, it was best in the previous example to take θ_0 near the upper bound of the interval. For this example, however, the opposite is true, and the reason is that the performance measure $h(\theta, \omega)$ here is *decreasing* w.r.t. θ; therefore, importance sampling considerations (Glynn and Iglehart 1989) tell us that it is better to take $\theta_0 < \theta$ (simulating at a smaller θ_0 stochastically increases the number of zero-waits).

3.2. A Reliability System with Age Limits

Consider a Q-component system, where the components evolve independently of each other. For each $q = 1, \ldots, Q$, component q has a lifetime distribution F_q, with age-dependent failure rate λ_q (failure ends the life of the component), and replacement-time distribution G_q. Component q also has an age limit L_q: whenever its age reaches L_q, it is replaced immediately by a new one. This is called a preventive replacement. So, for each q, component q lives for a while, then either fails or is replaced preventively, then the repair time must elapse, the new component starts working, lives for a while, and so on. Let $t_0 = 0$ and $t_1, t_2, t_3, \ldots$ be the successive event times in the evolution of this system, that is, the times at which there is either a failure, or reaching of an age-limit, or the end of a repair.

We assume that costs are incurred continuously, and that the cost rate $c(t)$ at time t is a continuous function of the component's states and ages at t. Then, the total cost incurred over a time-horizon T is given by $\int_0^T c(t)dt$. For example, the system could be a coherent reliability system (with redundancies) and $c(t) = 1$ when the system is operational, $c(t) = 0$ otherwise. If there is a failure cost K_q incurred each time component q fails, then this failure cost can be integrated into $c(t)$ by replacing it by a cost rate of $\lambda_q(x_q(t))K_q$, where $x_q(t)$ is the age of component q at time t.

Suppose that F_q or G_q depends on a parameter θ, and denote them by $F_{q,\theta}$ and $G_{q,\theta}$. Suppose that $F_{q,\theta}^{-1}(u)$ and $F_{q,\theta}^{-1}(u)$ are differentiable w.r.t. θ for each u. Now, a infinitesimal change in θ will provoke infinitesimal changes in the t_i's, but will not significantly change the sequence of events. In fact, two successive events can change order only if they occur almost simultaneously, and such a change of order will not affect the order of the subsequent events. Therefore, these ordering changes are not making the sample cost discontinuous w.r.t. θ and (under minor additional conditions), IPA yields an unbiased derivative estimator for the expected total cost. That estimator is given by:

$$h'(\theta, \omega) = \sum_{\{i:0<t_i<T\}} [c(t_i^-) - c(t_i)]t_i', \qquad (13)$$

where t_i' is the derivative of t_i w.r.t. θ. It is easily seen that the t_i''s can be expressed as random variables whose distributions depend on θ, as for the previous examples, and that a likelihood ratio based on the lifetime or repair-time random variables can then be defined to construct the IPA-LR gradient estimator.

Now, suppose that the parameter θ is an age-limit; i.e., $L_q = \theta$ for some q. Again, h is continuous w.r.t. an infinitesimal change in L_q, and the IPA estimator (13) still works nicely in that case. Indeed, even if a slight increase in L_q now makes a failure to occur before a preventive replacement, whereas the opposite was true with the original value of L_q (that is, an event changes type, from "preventive replacement" to "failure"), that changes the cost rate only during an infinitesimal duration. In (13), if t_i is the time of an event involving component q, then t_i' will be equal to the number of times the age limit L_q was reached for component q until time t_i. But what is the likelihood ratio in this case ? The only "random variable" that depends on θ here is L_q, and it is a degenerate one, whose *support* depends on θ. Therefore, the likelihood ratio does not exist and so the IPA-LR estimator cannot be constructed. The SF and SF-LR estimators do not exist either in this case.

4. CONCLUSION

We have shown how a change of measure can be combined with IPA to obtain a gradient estimator in functional form, which we called the IPA-LR estimator. We compared it with the previously introduced SF-LR estimator for simple examples. The IPA-LR estimator can be constructed in a similar way for other applications like stochastic PERT networks, queueing networks, some inventory systems, and so on. Typically, it will work only if *both* the IPA and SF gradient estimation methods apply to the problem at hand. One requirement is that the likelihood ratio must exist. There are applications where IPA applies and where the parameter θ is not a parameter of some probability distribution, but a *threshold*. Examples include some classes of inventory systems where θ is an inventory level, maintenance models where θ is an age-limit (as in §3.2), or production systems where θ an hedging-point (see Haurie, L'Ecuyer, and van Delft 1993). For those applications, the IPA-LR and SF-LR methods do not apply. There are other applications (see L'Ecuyer 1990 or Rubinstein and Shapiro 1993) where SF-LR does apply, but where IPA, and therefore IPA-LR, does not.

In the queueing examples that we have implemented, in turned out that θ was easy to factor out of $h'(\theta, \omega)$; see Equation (9). This is typical, e.g., for distributions from the exponential family, or for location and scale parameters, if h does not have a too complicated structure, but may not necessarily always be the case in general. If θ cannot be factored out, then the actual computation of $\psi(\theta, \omega)$ could become more costly, and making that computation possible for all values of θ could require storing much more information.

In the $GI/M/1$ example that we have examined, the variance of both the SF-LR and IPA-LR estimators was increasing dramatically fast (with vertical asymptotes) as θ was getting away from θ_0, and was also increasing exponentially fast as a function of the simulation horizon t. This is typical. As a result, those functional estimators are typically useful only in a small neighborhood of θ_0 and for short horizons (i.e., for cases where the likelihood ratio is a product of only a few terms). Steady-state models can nevertheless be treated if they have a regenerative structure that can be exploited and the regenerative cycles are short; see Rubinstein (1991) and Rubinstein and Shapiro (1993).

Rubinstein and Shapiro (1993) study the use of SF-LR estimators for stochastic optimization through the stochastic counterpart (SC) method. It turns out that IPA-LR estimators can also be used in SC. Another method for stochastic optimization is the well-known stochastic approximation (SA) method (Kushner and Clark 1978; Polyak and Juditsky 1992). In contrast to SC, SA is an iterative method which does not require a gradient estimator in functional form, but only a point (gradient) estimator at each iteration. The latter estimator can be obtained, for instance, by either SF or IPA. Comparison between SC and SA in terms of efficiency, ease of implementation, etc., in the asymptotic sense and in the practical sense (e.g., for a small computer budget), has not been much investigated at this point. Such comparative studies should be the subject of further investigation. If SC and SA are performed in the "optimal" way, then both have the same convergence rates, namely $O(N^{-1/2})$. However, performing them in the optimal way is typically hard to achieve in practice.

ACKNOWLEDGMENTS

Myriam Antaki helped producing the figures. This work has been supported by NSERC-Canada grant # OGP0110050 and FCAR-Québec grant # 93-ER-1654.

REFERENCES

Asmussen, S. and R. Y. Rubinstein. 1991. "The Performance of Likelihood Ratio Estimators Using the Score Function", Manuscript, Chalmers University of Technology, Göteborg, Sweden.

Glasserman, P., *Gradient Estimation via Perturbation Analysis*, Kluwer Academic, 1991.

Glynn, P. W. "Likelihood Ratio Gradient Estimation for Stochastic Systems", *Communications of the ACM*, **33**, 10 (1990), 75–84.

Glynn and Iglehart. 1989. Importance Sampling for Stochastic Simulations. *Management Science* **35**, 1367–1392.

Haurie, A., P. L'Ecuyer, and Ch. Van Delft. 1993. Monte-Carlo Optimization of Parametrized Policies in a Class of Piecewise Deterministic Control Systems Arising in Manufacturing Flow Control. To appear in *Discrete Event Dynamic Systems.*

L'Ecuyer, P., "A Unified View of the IPA, SF, and LR Gradient Estimation Techniques", *Management Science*, **36**, 11 (1990), 1364–1383.

Kushner, H. J. and D. S. Clark, *Stochastic Approximation Methods for Constrained and Unconstrained Systems*, Springer-Verlag, Applied Math. Sciences, vol. 26, 1978.

L'Ecuyer, P. 1993. On the Interchange of Derivative and Expectation.... *Management Science*, to appear.

Polyak, B. T. and A. B. Juditsky. 1992. Acceleration of Stochastic Approximation by Averaging. *SIAM J. on Control and Optimization* **30**, 838–855.

Reiman, M and A. Weiss. 1989. Sensitivity Analysis for Simulations via Likelihood Ratios. *Operations Research* **37**, 830–844.

Rubinstein, R. Y. 1991. How to Optimize Discrete-Event Systems from a Single Sample Path by the Score Function method. *Annals of Operations Research* **27**, 175–212.

Rubinstein, R. Y. and A. Shapiro, *Discrete Event Systems: Sensitivity Analysis and Stochastic Optimization by the Score Function Method*, Wiley, 1993.

AUTHOR BIOGRAPHY

PIERRE L'ECUYER is a professor in the department of "Informatique et Recherche Opérationnelle" (IRO), at the University of Montreal. He received a Ph.D. in operations research in 1983, from the University of Montreal. From 1983 to 1990, he was with the computer science department, at Laval University, Québec. His research interests are in Markov renewal decision processes, sensitivity analysis and optimization of discrete-event stochastic systems, random number generation, and discrete-event simulation in general. He is the Departmental Editor for the Simulation Department of *Management Science* and an Associate Editor for the *ACM Transactions on Modeling and Computer Simulation*.