Proceedings of the 1991 Winter Simulation Conference, pp. 207–217.

AN OVERVIEW OF DERIVATIVE ESTIMATION

Pierre L'Ecuyer

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

ABSTRACT

We explain the main techniques for estimating derivatives by simulation and survey the most recent developments in that area. In particular, we discuss perturbation analysis (PA), likelihood ratios (LR), weak derivatives (WD), finite differences (FD), and many of their variants. We also mention some other approaches. Our discussion emphasizes the relationships between the methods. For that purpose, all of them are presented in the same framework, which is based on L'Ecuyer (1990).

1. INTRODUCTION

Simulation is a popular tool for estimating the expected (average) performance measure of a complex stochastic system. Various statistical techniques have been develop ed in that context. Estimating the derivative or sensitivity of such an expectation certainly looks more difficult, but is nevertheless important for many practical applications. For example, let θ be a real-valued (continuous) parameter and suppose that the performance measure of interest depends on θ either directly, or indirectly through the probability law that governs the evolution of the system, or both.

If θ is a decision parameter, one might want to *optimize* the expected performance, say $\alpha(\theta)$, as a function of θ. It is a well known fact that if α is well behaved, much more efficient algorithms are available when derivative evaluations (or estimations), and not just function evaluations, can be obtained. See, e.g., Andradóttir (1991), Benveniste, Métivier, and Priouret (1987), Glynn (1986, 1989a), Kushner and Clark (1978), L'Ecuyer, Giroux, and Glynn (1991), Luenberger (1984), Meketon (1987), Métivier and Priouret (1984), Pflug (1990), and Rubinstein (1991). For that purpose, one would need a way of estimating $\alpha'(\theta)$ at any given point θ in the domain of α.

In other applications, θ is not a decision parameter, but a parameter of the model that has been estimated from statistical data. Then, one might be interested in the *sensitivity* of $\alpha(\theta)$ with respect to θ. Again, this is $\alpha'(\theta)$. Sensitivity analysis is useful for discovering which parameters in a model are important and which ones are not. It could also indicate that the model is excessively sensitive to some "critical" parameters. This may imply a questionable model or at least warn the decision maker not to be overly confident in the simulation results.

A third class of applications is *interpolation*. See Reiman and Weiss (1989) and the references given there.

In this paper, we look at different methods for estimating $\alpha'(\theta)$, for $\theta \in \mathbb{R}$. Here, $\alpha(\theta)$ can be either the expected performance measure (or "cost") over a finite (deterministic or random) horizon, or an infinite-horizon average cost per unit of time. For the case where θ is a vector of parameters, one can just apply the methods discussed here to obtain estimators of the derivative with respect to each component of θ. This yields the gradient of $\alpha(\theta)$ with respect to θ, which is the vector whose i-th component is the derivative of α with respect to the i-th component of θ.

We examine finite-differences (FD), perturbation analysis (PA), likelihood ratio (LR), and weak derivative (WD) methods. For more on these methods, the reader can look at the many recent references given at the end of this paper. We recommend in particular the following ones, which are more "general scope" or "survey style" : Glasserman (1991a), Glynn (1990), Ho and Cao (1991), L'Ecuyer (1990), Rubinstein (1991), and Suri (1989).

2. SAMPLE PERFORMANCE DERIVATIVE

2.1. Infinitesimal Perturbation Analysis (IPA)

Consider a simulation model defined over a probability space (Ω, Σ, P). Let $h(\theta, \omega)$ denote the sample value (cost), where the sample point $\omega \in \Omega$ obeys the probability law P, and $\theta \in \Theta$, where Θ is some open interval in $\mathbb{R}$. We assume that $h(\theta, \cdot)$ is measurable for each θ and (for the moment) that P does not depend on θ. For example, in a typical simulation, ω can be viewed as the sequence of underlying independent $U(0,1)$ variates that drive the simulation. The expected value (cost) is

$$\alpha(\theta) = \int_\Omega h(\theta, \omega) dP(\omega). \tag{1}$$

The basic idea of IPA is simply to take $h'(\theta, \omega)$, the derivative of $h(\theta, \omega)$ with respect to θ, for fixed ω, as an estimator of $\alpha'(\theta)$. This yields an unbiased estimator if one can differentiate α by differentiating inside the integral in (1), i.e. if the derivative and expectation can be interchanged. Sufficient conditions for that interchange to be valid are given in Glasserman (1988, 1991a), L'Ecuyer (1990, 1991b), and Pflug (1991). The conditions of L'Ecuyer (1991b) are recalled in the following theorem.

THEOREM 1. *Let $\theta_0 \in \Upsilon \subseteq \Theta$, where Υ is an open interval, and let $\Xi \subseteq \Omega$ be a measurable set such that $P(\Xi) = 1$. Suppose that for every $\omega \in \Xi$, there is a $D(\omega) \subseteq \Upsilon$, where $\Upsilon \setminus D(\omega)$ is at most a denumerable set, such that $h(\cdot, \omega)$ exists and is continuous everywhere in Υ, and is also differentiable everywhere in $D(\omega)$. Suppose also that there exists a P-integrable function $\Gamma : \Omega \to [0, \infty)$ such that*

$$\sup_{\theta \in D(\omega)} |h'(\theta, \omega)| \leq \Gamma(\omega)$$

for every ω in Ξ. Then, everywhere in Υ, α is differentiable and

$$\alpha'(\theta) = \int_\Omega h'(\theta, \omega) dP(\omega). \tag{2}$$

Further, if $h'(\cdot, \omega)$ is continuous all over Υ for each $\omega \in \Xi$, then α is continuously differentiable in Υ. ∎

Unfortunately, there are many practical applications where the conditions of Theorem 1 do not hold and where the interchange is not valid. Roughly speaking, in a discrete-event simulation, IPA assumes that an infinitesimal perturbation on θ does not affect the sequence of events, but only makes their occurence times "slide smoothly". (Note that the perturbation can also affect $h(\theta, \omega)$ *without* affecting the event times either). In some applications, varying θ affects the performance of the system only (or mainly) through a drastic change in the sequence of events. This is the case, for example, if θ is a routing probability in a multiclass queueing network, or a threshold in an production/inventory or repair/replacement model (L'Ecuyer 1990). In some cases, the model can be reformulated differently, i.e. the function $h(\theta, \omega)$ can be replaced by an alternative one which has the correct expectation but which is "smoother" in some sense, so that Theorem 1 applies. Sometimes, for a given interpretation of ω, there are many possible functions $h(\theta, \omega)$ which have the correct expectation and not all of them satisfy the conditions of Theorem 1. See Glasserman (1991c) for an example where IPA does not work when $h(\theta, \omega)$ is defined in a standard way, but does work if a more clever (and more complicated) definition is adopted. In most cases, the interpretation of ω is also changed to obtain a smoother function. Basically, all the methods that we will discuss later on under the names of LR, SPA, RPA, and so on, are variants of this "smoothing" idea.

Example 1 : Total Sojourn Time for the First t Customers in a GI/G/1 Queue. This is a classic example in the field. Consider a GI/G/1 queue with service time distribution B_θ and corresponding density b_θ. For simplicity, assume that θ is a scale parameter, that is $B_\theta(s) = B(s/\theta)$ for some distribution B with bounded density b. Let $\Theta = (\ell, u)$ for $0 < \ell < u$. For $i \geq 1$, let W_i, $S_i = \theta Z_i$, and $X_i = W_i + S_i$ be the *waiting* time, *service* time, and *sojourn* time for the i-th customer, and A_i be the time between arrivals of the i-th and $(i+1)$-th customer. Here, Z_i follows the distribution B and is assumed to have finite expectation. One has $W_1 = 0$, and for $i \geq 1$,

$$X_i := W_i + S_i = W_i + \theta Z_i \qquad \text{and} \qquad W_{i+1} := (X_i - A_i)^+ \tag{3}$$

where x^+ means $\max(x, 0)$. Suppose that we are interested in the expected total sojourn time in the system for the first t customers, and in the derivative of that. In other words, let

$$h(\theta, \omega) = \sum_{i=1}^{t} X_i. \tag{4}$$

One has

$$X_i' = \frac{\partial X_i}{\partial \theta} = Z_i + W_i' = Z_i + X_{i-1}' I(W_i > 0),$$

where I denotes the indicator function. From that,

$$X_i' = \sum_{j \in \Phi_i} Z_j$$

where Φ_i is the set containing customer i and all the customers that precede him in the same busy period (if any). The IPA estimator is then

$$h'(\theta, \omega) = \sum_{i=1}^{t} \sum_{j \in \Phi_i} Z_j. \quad (5)$$

In other words, an infinitesimal perturbation on θ becomes an infinitesimal perturbation on each S_j, which affects the system time of customer j and of all the customers (if any) that follow him in the same busy period.

The assumptions of Theorem 1 are easy to verify in this case. Indeed, from (3–4), $h(\theta, \omega)$ is clearly continuous in θ. It is also differentiable in θ everywhere except when two events (arrival or departure) occur simultaneously, which happens at most for a finite number of values of θ. Also, $|h'(\theta, \omega)| \leq t \sum_{i=1}^{t} Z_i$, which is integrable since each Z_i has finite expectation.

In this example, we have viewed the sample space in such a way that ω represents a sequence of $U(0,1)$ variates. Alternatively, one can view ω as representing $(Z_1, A_1, \ldots, Z_{t-1}, A_{t-1}, Z_t)$. The sample space and probability measure are then different, but this gives exactly the same IPA estimator. On the other hand, if ω is viewed as representing $(S_1, A_1, \ldots, S_{t-1}, A_{t-1}, S_t)$, the probability measure then depends on θ. We will see how to deal with that in Section 2.4.

Example 2 : Probability that the Sojourn Time Exceeds L. (This is taken from L'Ecuyer and Perron 1990 and Wardi et al. 1991a) Suppose that in Example 1, the objective function is replaced by $\alpha(\theta) = P[X_j > L]$, for given j and L. If ω is viewed as the underlying sequence of uniform variates, as in Example 1, and $h(\theta, \omega) = I(X_j > L)$, where I is the indicator function, then $h'(\theta, \omega) = 0$ except when $X_j = L$, where it is undefined. But since S_j has a density, $P[X_j = L] = 0$, so that $E[h'(\theta, \omega)] = 0 \neq \alpha'(\theta)$. In other words, this naïve IPA derivative estimator is completely useless. One can nevertheless use IPA in this case, as will be shown in the next subsection.

2.2. Smoothed Perturbation Analysis (SPA)

The basic idea of SPA is to replace the cost estimator $h(\theta, \omega)$ by its expectation conditional on some "part" of ω, before taking the derivative with respect to θ. See Glasserman and Gong (1990), Gong and Ho (1987), L'Ecuyer and Perron (1990), Vázquez-Abad and L'Ecuyer (1991), and Wardi et al. (1991a, 1991b). Equivalently, this corresponds to viewing ω as representing something different than the whole sequence of underlying uniform variates, and then using the same framework as for IPA. We illustrate that on Example 2.

Example 2 (continuation). Let us view ω as representing the sequence of uniform variates that have been used to generate $(S_1, A_1, \ldots, S_{j-1}, A_{j-1})$, and redefine

$$h(\theta, \omega) = P[X_j > L \mid \theta, \omega] = 1 - B_\theta(s),$$

where $s = L - (X_{j-1} - A_{j-1})^+$. The corresponding IPA estimator (L'Ecuyer and Perron 1990, and Wardi et al. 1991a) is then

$$\begin{aligned} h'(\theta, \omega) &= -\frac{\partial}{\partial \theta} B_\theta(s) + b_\theta(s) X'_{j-1} I(X_{j-1} \geq A_{j-1}) \\ &= b(s/\theta)[s/\theta^2 + X'_{j-1} I(X_{j-1} \geq A_{j-1})]. \end{aligned}$$

Since the service time density is bounded, $h(\theta, \omega)$ is continuous and piecewise differentiable for each ω, and $h'(\theta, \omega)$ is integrable. Therefore, Theorem 1 applies and we now have an unbiased IPA (or SPA) derivative estimator.

Note that strictly speaking, for the conditional probability $P[\cdot \mid \theta, \omega]$ to be well defined, it is assumed implicitly that there is a "lower-level" probability space in which the sample point $\tilde{\omega}$, say, can be interpreted as the whole sequence of uniform variates that is used throughout the simulation. In other words, that $\tilde{\omega}$ has the same meaning as ω in "standard" IPA. The reason why we rename it $\tilde{\omega}$ and use ω to denote what we condition on is to re-obtain cost and derivative estimators h and h' that are functions of ω. Then, the fact that SPA is just IPA applied over a different cost estimator is made more apparent, and the results of standard IPA can be applied directly with the same notation.

For a given application, different SPA derivative estimators can often be obtained by conditioning on different things, i.e. with different choices of what ω represents. In Example 2, for instance, one can view ω as representing the sequence of uniform variates that have been used to generate $(S_1, A_1, \ldots, A_{j-2}, S_{j-1})$. Since we condition on less, this will give and estimator with less variance. See L'Ecuyer and Perron (1990) for more details and numerical results.

2.3. Other IPA Variants

Another perturbation analysis variant which is somewhat related to SPA is called Rare Perturbation Analysis (RPA) (Brémaud and Vázquez-Abad 1991, Vázquez-Abad and Kushner 1991, and Vázquez-Abad and L'Ecuyer 1991) or Light Traffic Perturbation

(LTP) (Simon 1989). Instead of sliding events smoothly in time, RPA or LTP will cancel or add some events with very small (or infinitesimal) probabilities. Originally, RPA has been designed as a finite-difference (or FPA) method, but Brémaud and Vázquez-Abad (1991) have developped a way of taking RPA to the limit. This yields a derivative estimator which is an average of conditional expectations, i.e., roughly speaking, an average of SPA estimators. The conditioning can be, for example, on the number of events that are actually cancelled. An example is worked out in Vázquez-Abad and L'Ecuyer (1991), in these Proceedings. LPT is quite similar. Other variants of Perturbation analysis are also described in Ho and Strickland (1990), Ho and Cao (1991), and the references given there. Some of these are based on "cutting and pasting" sample paths (Ho and Li 1988), on Markov chain aggregation via LR (Zhang and Ho 1991), etc.

We also note that the term *perturbation analysis* often refers not only to the idea of computing $h'(\theta,\omega)$, but (perhaps above all) to the class of techniques that could be used to actually compute or approximate $h'(\theta,\omega)$ during a simulation, or during the operation of a real-time system, using a single sample path (Ho and Cao 1991).

2.4. Likelihood Ratios (LR)

In the previous subsections, the probability law that governs ω has been assumed independent of θ. In standard IPA, ω is usually viewed as an underlying sequence of independent $U(0,1)$ variates, or something equivalent to that. But for many applications, for ω viewed that way, valid IPA or SPA estimators are quite difficult or impossible to obtain. We will now consider the case where the probability law of ω depends on θ. The basic idea of the likelihood ratio (LR) approach is to transform the value function and probability law in such a way that the latter will not depend on θ anymore, and that IPA can be applied after the transformation. The method can be traced back to Aleksandrov, Sysoyev, and Shemeneva (1968). More recently, it has been studied in Asmussen (1991), Glynn (1986, 1987, 1990), Heidelberger and Towsley (1989), L'Ecuyer (1990, 1991b), Reiman and Weiss (1989), and Rubinstein (1986a, b, 1989, 1991).

Let $\{P_\theta,\ \theta \in \Theta\}$ be a family of probability measures, defined over the same measurable space (Ω, Σ), where Θ is some open interval in $\mathbb{R}$, as before. We now suppose that $\omega \in \Omega$ obeys the probability law P_θ. In that case, we cannot differentiate the expectation by differentiating directly the sample cost inside the integral, as in IPA, because the expectation itself is with respect to a probability measure P_θ that depends on θ. But if G is a given probability measure on (Ω, Σ) that dominates all the P_θ's, that is $G(B) = 0$ for a measurable set B implies $P_\theta(B) = 0$ for each θ, then the expected value (cost) can be written as a function of θ as

$$\alpha(\theta) = \int_\Omega h(\theta,\omega)dP_\theta(\omega) = \int_\Omega h(\theta,\omega)L(G,\theta,\omega)dG(\omega), \tag{6}$$

where $L(G,\theta,\omega) = (dP_\theta/dG)(\omega)$ is the Radon-Nikodym derivative of P_θ with respect to G, evaluated at ω. Define

$$H(\theta,\omega) = h(\theta,\omega)L(G,\theta,\omega). \tag{7}$$

We are now in the same framework as for IPA, with h and P replaced by H and G, respectively. Therefore, the IPA results can be used again here. In particular, if the sufficient conditions given in Theorem 1 are verified for H, one obtains:

$$\alpha'(\theta) = \int_\Omega H'(\theta,\omega)dG(\omega), \tag{8}$$

where

$$H'(\theta,\omega) = L(G,\theta,\omega)h'(\theta,\omega) + h(\theta,\omega)L'(G,\theta,\omega). \tag{9}$$

In typical models where P_θ depends on θ, ω can be viewed as a sequence $(\zeta_1,\ldots,\zeta_t)$, where the ζ_i's are independent, or where the distribution of ζ_i conditional on $(\zeta_1,\ldots,\zeta_{i-1})$ is known for each i so that the joint distribution (or likelihood) of $(\zeta_1,\ldots,\zeta_t)$ can be written down easily as a product. We assume for the moment that t is deterministic and finite and that ζ_i has density $f_{i,\theta}$. Our development extends trivially to the case where some or all of the ζ_i's have probability mass functions instead, or conditional laws as just explained. This will be illustrated in the next example. The case of a random t will be discussed later on. For $i = 1,\ldots,t$, let g_i be a density such that $\{\zeta \mid f_{i,\theta}(\zeta) > 0 \text{ for some } \theta \in \Theta\} \subseteq \{\zeta \mid g_i(\zeta) > 0\}$. The densities g_i define the probability measure G over (Ω, Σ). In that case, the Radon-Nikodym derivative becomes the likelihood ratio

$$L(G,\theta,\omega) = \prod_{i=1}^{t} \frac{f_{i,\theta}(\zeta_i)}{g_i(\zeta_i)}. \tag{10}$$

Its derivative is

$$L'(G,\theta,\omega) = L(G,\theta,\omega)S(\theta,\omega), \tag{11}$$

where

$$S(\theta,\omega) = \frac{\partial}{\partial\theta}\ln(L(G,\theta,\omega)) = \sum_{i=1}^{t} d_i \tag{12}$$

is called the *score function* and d_i is defined as

$$d_i = \frac{\partial}{\partial \theta} \ln(f_{i,\theta}(\zeta_i)). \qquad (13)$$

L'Ecuyer (1991b) gives specific conditions, in this context, under which Theorem 1 applies. In particular, one can take $g_i = f_{i,\theta_0}$, for some $\theta_0 \in \Theta$. In that case, the likelihood ratio at $\theta = \theta_0$ is 1 and the first term in (9) becomes a "direct IPA" part.

Note that (9) could be used to estimate the derivative α' *everywhere* in a given region by a *single* simulation. The idea is to compute an expression for $H'(\cdot, \omega)$ in that region (i.e. a whole function of θ, for fixed ω), and take that as an estimate of $\alpha'(\cdot)$. See Asmussen (1991) and Rubinstein (1991).

Example 3 : A GI/G/1 Queue with Rejections. Suppose that the i-th customer arriving to our GI/G/1 queue (Example 1) is rejected with probability $p_i = 1 - 1/(N_i\theta)$ when $N_i \geq 1$, $p_i = 0$ otherwise, where N_i is the number of customers already in the queue (waiting or in service) when customer i arrives. We assume that $\theta \in \Theta = (\ell, u)$ for $1 < \ell < u$. The Lindley equations (3) are still valid if the service times S_i of the rejected customers are replaced by $\tilde{S}_i = 0$. (The values of W_i and X_i corresponding to rejected customers are meaningless, but well defined.) Suppose that we want to estimate the total sojourn time for the customers that are not rejected, among the first t. Let $I_i = 1$ if customer i is rejected, $I_i = 0$ otherwise. A straightforward cost estimator is

$$h(\theta, \omega) = \sum_{i=1}^{t} X_i(1 - I_i). \qquad (14)$$

For simplicity, suppose that the service time distribution does not depend on θ.

If ω represents the underlying sequence of uniform variates, changing θ by a small amount for fixed ω may change N_i, so that p_i is discountinuous in θ. Then, the continuity condition in Theorem 1 fails to hold. In fact, one can easily check that $h'(\theta, \omega) = 0$ with probability one, so that this derivative estimator is useless. Further, a valid SPA estimator does not come to (my) mind easily.

Now, let us view ω as representing the sequence $(S_1, I_1, A_1, \ldots, A_{t-1}, S_t, I_t)$. From that, the N_i's, W_i's, and X_i's can be deduced easily, independently of θ. Therefore, for fixed ω, $h(\theta, \omega)$ does not depend anymore on θ, so that $h'(\theta, \omega) = 0$. Now, only the probability law of ω depends on θ, via the p_i's. For each i, let $p_i(\theta) = 1 - 1/(N_i\theta) = P_\theta[I_i = 1 \mid N_i]$ if $N_i > 0$, and $p_i(\theta) = 0$ otherwise. Here, the I_i's are not independent, but the joint probability mass of $(I_1, \ldots, I_t)$, for given $(S_1, A_1, \ldots, A_{t-1}, S_t)$, is the product $\prod_{i=1}^{t} p_i(\theta)^{I_i}(1 - p_i(\theta))^{1-I_i}$, with the convention that $0^0 = 1$. Let $G = P_{\theta_0}$ for some $\theta_0 \in \Theta$. The likelihood ratio is then

$$L(G, \theta, \omega) = \prod_{i=1}^{t} \frac{p_i(\theta)^{I_i}(1 - p_i(\theta))^{1-I_i}}{p_i(\theta_0)^{I_i}(1 - p_i(\theta_0))^{1-I_i}},$$

and the score function is

$$\begin{aligned} S(\theta, \omega) &= \frac{\partial}{\partial \theta} \ln L(G, \theta, \omega) \\ &= \sum_{i=1,\, N_i>0}^{t} \frac{\partial}{\partial \theta} \Bigg[I_i \ln \left(1 - \frac{1}{N_i\theta}\right) \\ &\qquad + (1 - I_i) \ln(N_i\theta) \Bigg] \\ &= \frac{1}{\theta} \sum_{i=1,\, N_i>0}^{t} \left[\frac{I_i}{N_i\theta - 1} + 1 - I_i \right]. \end{aligned}$$

Also, $H(\theta, \omega)$ is continuously differentiable in θ. Assume that $E_{\theta_0}[S_i] < K$ for each i, for some constant $K < \infty$. Then,

$$\begin{aligned} |H'(\theta, \omega)| &= \left[\frac{1}{t} \sum_{i=1}^{t} (W_i + S_i)(1 - I_i) \right] L(G, \theta, \omega) |S(\theta, \omega)| \\ &\leq \left[\sum_{i=1}^{t} S_i \right] \left(\frac{u}{\ell - 1} \right)^t \left(\frac{t}{\ell - 1} \right) (t-1)!, \end{aligned}$$

which is P_{θ_0}-integrable. Therefore, from Theorem 1, $E_{\theta_0}[H'(\theta, \omega)] = \alpha'(\theta)$ for each $\theta \in \Theta$.

2.5. The Weak Derivative (WD)

In LR, we obtained a derivative estimator by changing both the value function and the probability measure, and differentiating inside the integral. In IPA, the value function was differentiated directly and the probability measure kept unchanged. Now, we will look at the opposite idea: keep the value function unchanged and try to differentiate directly the probability measure. We will in fact replace it by its *weak derivative*. This idea was introduced by Pflug (1989, 1991).

Under appropriate conditions (Pflug 1991), one can write

$$\alpha'(\theta) = \int_\Omega h'(\theta, \omega) dP_\theta(\omega) + \int_\Omega h(\theta, \omega) dP'_\theta(\omega), \qquad (15)$$

where P'_θ is a finite signed measure. The first term in (15) is an IPA term that can be dealt with as usual: generate ω according to P_θ and compute $h'(\theta, \omega)$. In the second term, P'_θ is not necessarily a probability measure, but from the Jordan-Hahn decomposition, it can be written as

$$P'_\theta = c(\theta)(\dot{P}_\theta - \ddot{P}_\theta), \tag{16}$$

where $\dot{P}_\theta$ and $\ddot{P}_\theta$ are probability measures on (Ω, Σ). Then, if $\dot{\omega}$ and $\ddot{\omega}$ are sample paths generated from $\dot{P}_\theta$ and $\ddot{P}_\theta$, respectively, a weak derivative (WD) estimator for the second term in (15) is given by

$$c(\theta)(h(\theta, \dot{\omega}) - h(\theta, \ddot{\omega})). \tag{17}$$

(Note that in Equation (11) in Pflug (1989), the $c(x)$ should be in the numerator.)

For example, let $\Omega = \mathbb{R}$, P_θ have a density f_θ, and $h'(\theta, \omega) = 0$. Then,

$$\alpha(\theta) = \int_\Omega h(\theta, \omega) f_\theta(\omega) d\omega,$$

and, under appropriate regularity conditions,

$$\alpha'(\theta) = \int_\Omega h(\theta, \omega) \left(\frac{\partial}{\partial \theta} f_\theta(\omega) \right) d\omega.$$

One can decompose

$$\frac{\partial}{\partial \theta} f_\theta = c(\dot{f}_\theta - \ddot{f}_\theta)$$

where

$$\begin{aligned} c &= \int_{-\infty}^{\infty} \left(\frac{\partial}{\partial \theta} f_\theta(\omega) \right)^+ d\omega \\ &= \int_{-\infty}^{\infty} \left(\frac{\partial}{\partial \theta} f_\theta(\omega) \right)^- d\omega, \\ \dot{f}_\theta &= \frac{1}{c} \left(\frac{\partial}{\partial \theta} f_\theta \right)^+, \\ \ddot{f}_\theta &= \frac{1}{c} \left(\frac{\partial}{\partial \theta} f_\theta \right)^-. \end{aligned}$$

A WD estimator of $\alpha'(\theta)$ is then

$$c(h(\theta, \dot{\omega}) - h(\theta, \ddot{\omega}))$$

where $\dot{\omega}$ and $\ddot{\omega}$ have respective densities $\dot{f}_\theta$ and $\ddot{f}_\theta$. One disadvantage of this method is that *two* simulations must be performed. Numerical illustrations are given in Pflug (1989).

2.6. Frequency Domain Experimentation (FDE)

Jacobson (1991a, b) and Jacobson and Schruben (1991) study the use of frequency domain experimentation (FDE) for estimating derivatives. They point out the relationship between FD and FDE. The basic idea of FDE is to oscillate the different parameters at (different) given frequencies during the simulation and analyze the oscillations in the response through, e.g., harmonic analysis. Whether or not (and in which situations) FDE would be competitive in practice for derivative estimation is not clear at this point.

2.7. Higher Order Derivatives

IPA and LR estimation can be generalized to higher order derivatives. See, e.g., L'Ecuyer (1990) or Rubinstein (1989, 1991). For example, by taking the derivative of (9), under appropriate regularity conditions given by an easy adaptation of Theorem 1, one obtains the following second derivative estimator:

$$\begin{aligned} H''(\theta, \omega) &= L(G, \theta, \omega) h''(\theta, \omega) + h(\theta, \omega) L''(G, \theta, \omega) \\ &\quad + 2h'(\theta, \omega) L'(G, \theta, \omega). \end{aligned}$$

There are also generalizations for mixted partial derivatives of higher order when there are many parameters. For example, one can estimate a Hessian.

3. FINITE DIFFERENCE METHODS

Finite differences (FD) have been used for a long time to estimate derivatives. See Glynn (1989a), Kushner and Clark (1978), L'Ecuyer and Perron (1990), Pflug (1989), Rubinstein (1986a), and Zazanis and Suri (1988). FD can be used when none of the methods presented in the previous section would apply, or when they are judged too complicated to implement.

3.1. FD and FDC

Suppose that the sample point ω represents the underlying sequence of uniform variates, as in IPA. Let ω^- and ω^+ be two independent sample points generated under P, and let $\epsilon > 0$. The *forward* FD estimator is

$$\frac{h(\theta + \epsilon, \omega^+) - h(\theta, \omega^-)}{\epsilon}, \tag{18}$$

while the *central* FD estimator is

$$\frac{h(\theta + \epsilon, \omega^+) - h(\theta - \epsilon, \omega^-)}{2\epsilon}. \tag{19}$$

These estimators are biased. If α is three times continuously differentiable, the bias is in the order of ϵ for the forward version and in the order of ϵ^2 for the central version. The major problem here is that as ϵ decreases to zero, the variance of these FD estimators goes to infinity.

This variance problem can be addressed by using common random numbers. The idea is simply to generate only one ω using P and to take $\omega^- = \omega^+ = \omega$ in the above formulas. This yields the forward and central FDC estimators (FD with common random numbers). Using $\omega^- = \omega^+$ corresponds to comparing very similar systems, under the same conditions. For small ϵ one should expect $h(\theta+\epsilon, \omega^+)$ and $h(\theta-\epsilon, \omega^-)$ to be highly correlated, so that a considerable variance reduction could be obtained. Conditions that *guarantee* variance reductions are given in Glasserman and Yao (1990) and Rubinstein (1986a).

3.2. Convergence Rates

In practice, to estimate a derivative, one would usually compute n independent replications of a given derivative estimator, take the average, and perhaps compute a confidence interval. The mean square error (MSE) of such an average should converge to zero. But at which rate ? For IPA, SPA, or LR, if the estimator is unbiased and the variance is finite, the MSE is in $O(1/n)$, from the central-limit theorem. For FD or FDC, things are more complicated. Let ϵ be a function of n, say $\epsilon = c_n$. To get the MSE down to zero as $n \to \infty$, one must take $c_n \to 0$ to get rid of the bias, but c_n should not go to zero too fast, otherwise the variance would not go to zero. Glynn (1989a), Zazanis and Suri (1988), and L'Ecuyer and Perron (1990), among others, give optimal sequences c_n and optimal orders of convergence for MSE in various FD and FDC contexts and under different sets of assumptions. With independent random numbers, the best convergence rates for MSE are in $O(n^{-1/2})$ in the forward case and in $O(n^{-2/3})$ in the central case. For FDC and under a given set of assumptions, Glynn (1989a) has obtained respective convergence rates of $O(n^{-2/3})$ and $O(n^{-4/5})$. It turns out that his assumptions could hold only when IPA does not apply, but it is precisely in that situation that FDC has more chances of being used. L'Ecuyer and Perron (1990) have shown that under the conditions of Theorem 1, if the variance is bounded and α is twice continuously differentiable, the MSE of FDC is in $O(1/n)$ provided that c_n is in $O(n^{-1/2})$. In other words, when IPA works, FDC has the same convergence rate as IPA. This is interesting to know because in some situations, computing the IPA estimator can be complicated and/or tedious and FDC might be a reasonably efficient alternative. On the other hand, IPA estimators require just one simulation, whatever be the number of parameters, while if there are d parameters, FDC estimators require $2d$ simulations for the central case and $d+1$ simulations for the forward case. IPA can be applied to real-life systems (not just simulations), but not FDC (see Suri 1989). FDC can also have numerical problems when the intervals are too small. Finally, implementing FDC with the proper synchronisation is not always easy in practice, especially for complex systems.

3.3. Finite Perturbation Analysis

When ϵ is small, there is sometimes little change between the two sample paths in FDC. Sometimes, it is even possible to perform only *one* simulation and trace the few changes. This is called *finite perturbation analysis* (FPA). There are applications where IPA does not apply easily, but where a cleverly designed FPA approach does. See Ho and Cao (1991). Note that PA was originally introduced in its FPA version (Ho, Eyler, and Chien 1979).

4. A DISCRETE-TIME MARKOV CHAIN

Most discrete-event simulation models can be viewed as discrete-time Markov chains over general state spaces. Consider a Markov chain $\{X_i,\ i \geq 0\}$, with Borel state space S, defined as follows. Let $X_0 = s_0$ for some fixed initial state $s_0 \in S$. Let $\{Q_\theta(\cdot \mid s),\ \theta \in \Theta,\ s \in S\}$ be a family of probability measures on $\mathbb{R}$, with corresponding densities $\{q_\theta(\cdot \mid s),\ \theta \in \Theta,\ s \in S\}$. Let $c : \Theta \times S \times \mathbb{R} \to \mathbb{R}$ and $\varphi : \Theta \times S \times \mathbb{R} \to S$ be measurable functions called the *cost* and *transition* functions. The chain evolves as follows. At step i, a real-valued random variable ζ_i is "generated" according to $Q_\theta(\cdot \mid X_{i-1})$. A cost $C_i = c(\theta, X_{i-1}, \zeta_i)$ is incurred at that stage and the next state is $X_i = \varphi(\theta, X_{i-1}, \zeta_i)$.

4.1. Deterministic Horizon

Let

$$h_t(\theta, \omega) = \frac{1}{t} \sum_{i=1}^{t} C_i, \tag{20}$$

the average cost for the first t steps, where t is fixed. Here, one can view the sample point as $\omega = (\zeta_1, \ldots, \zeta_t)$. This ω obeys some probability law P_θ implied by the above framework.

If Q_θ is independent of θ and if c and φ are "smooth enough", IPA can be applied. The IPA derivative estimator is then

$$h'_t(\theta, \omega) = \frac{1}{t} \sum_{i=1}^{t} C'_i, \tag{21}$$

where the prime denotes the derivative with respect to θ, for fixed ω. Note that in general, computing C'_i involves X_{i-1}, X'_{i-1}, ζ_i, and ζ'_i. Therefore, for derivative estimation purposes, one can consider an "extended" Markov chain for which the state at step

i is defined as (X_{i-1}, X'_{i-1}). If this Markov chain is regenerative, then standard renewal theory can be used to study the behavior of $h_t(\theta,\omega)$, $h'_t(\theta,\omega)$, and of their expectations, as t goes to infinity. In particular, when $h_t(\theta,\omega)$ has the form (20), the variances of $h_t(\theta,\omega)$ and $h'_t(\theta,\omega)$ decrease linearly in t.

If Q_θ really depends on θ, one can use LR. Let $\{g(\cdot \mid s),\ s \in S\}$ be a family of densities such that for each s, the support of $g(\cdot \mid s)$ contains the support of each $q_\theta(\cdot \mid s)$. We will simulate using the densities g instead of q_θ, i.e. these densities g define the probability measure G. The likelihood ratio is then

$$L_t(G,\theta,\omega) = \prod_{i=1}^{t} \frac{q_\theta(\zeta_i \mid X_{i-1})}{g(\zeta_i \mid X_{i-1})}. \tag{22}$$

and the score function is

$$S_t(\theta,\omega) = \sum_{i=1}^{t} \frac{\partial}{\partial\theta} \ln(q_\theta(\zeta_i \mid X_{i-1})). \tag{23}$$

When $t \to \infty$, a major problem is that the variance of (23) increases linearly with t. Suppose that $h_t(\theta,\omega)$ converges to a constant K, as should be expected, that the likelihood ratio is one, i.e. that we want a derivative estimate at θ_0 and we take $g = q_{\theta_0}$, and that $h'(\theta,\omega) = 0$. Then, under mild conditions, the variance of the LR derivative estimator $H'_t(\theta,\omega) = h_t(\theta,\omega)S_t(\theta,\omega)$ increases linearly in t. This is really bad. However, if $K = 0$, things are slightly better. Indeed, L'Ecuyer and Glynn (1991) show that the variance is then in the order of 1 (with respect to t). This can be exploited as follows: just replace C_i by $C_i - K$ in (20). If K is unknown (the usual case), replace it by an estimation. Appropriate limit theorems are given in L'Ecuyer and Glynn (1990).

Another (less dramatic) way of reducing the variance is to use a *triangular* LR derivative estimator. The idea is to estimate the derivative of $E_\theta[C_i]$ separately for each i, and then take the average. Since C_i depends only on $\zeta_1, \ldots, \zeta_i$, the appropriate score function that should multiply C_i is the sum up to i instead of up to t. The triangular LR estimator is then

$$\frac{1}{t}\sum_{i=1}^{t} C_i \sum_{j=1}^{i} \frac{\partial}{\partial\theta} \ln(q_\theta(\zeta_j \mid X_{j-1})).$$

An alternative way of defining the likelihood ratio in this context, assuming that C_i can be expressed as a function of X_i, is to base it directly on the transition probabilities (or densities) between the successive visited states. In other words, one can use the likelihood of $(X_1, \ldots, X_t)$ instead of the likelihood of $(\zeta_1, \ldots, \zeta_t)$. This is what is done, for example, in Glynn (1987, 1990) and Pflug (1991). With X_0 known, $(X_1, \ldots, X_t)$ never contains more information than $(\zeta_1, \ldots, \zeta_t)$ and often contains less, so that the score function should be expected to have less variance when based on the former. However, the latter is often easier to deal with in actual implementations.

4.2. An IPA vs LR Paradox

In L'Ecuyer (1990), IPA was presented as a special case of LR: the case where P_θ does not depend on θ. In Section 2.4 of this paper, we have presented LR as IPA applied to a modified mathematical expectation. So, broadly speaking, IPA and LR represent the same approach. But we saw in Section 4.1 that for the model described there, the variance of the IPA estimator *decreases* linearly in t, while that of the LR estimator (which is IPA applied to $H_t(\theta,\omega)$) *increases* linearly in t. This apparent paradox can be explained by noting that the function $H_t(\theta,\omega) = h_t(\theta,\omega)L_t(G,\theta,\omega)$ does not have the form (20). When we said that the variance of $h'_t(\theta,\omega)$ decreases linearly in t, we assumed that $h_t(\theta,\omega)$ had the form (20). This is a special case. In general, when $h_t(\theta,\omega)$ does not have this form, anything can happen to the variance of the IPA estimator.

4.3. Random Horizon

In the previous subsection, for each t, let $\Im_t$ denote the sigma field generated by $\{\zeta_1, \ldots, \zeta_t\}$. Now, suppose that $h(\theta,\omega)$ is $\Im_\tau$-measurable for some (random) stopping time τ. We assume that for ω fixed, τ does not depend on θ. In the development of the previous subsection, t must now be replaced by τ. In the LR context, verifying the conditions of Theorem 1 then becomes more difficult in general, because the likelihood ratio is now the product of a random (generally unbounded) number of terms. L'Ecuyer (1991b) gives a set of sufficient conditions for that context. Other such conditions are given by Reiman and Weiss (1989) and Glynn (1986), in his Theorem 4.9.

4.4. GSMP and Other Frameworks

The Markov chain model of this section could be generalized to semi-markovian or Markov renewal models, in which the times between transitions are also random. The cost per step would be replaced by a cost per unit of time. Glynn (1989b) and Glasserman (1991a, b, c) suggest using a generalized semi-markov process (GSMP), with denumerable state space. This offers a convenient framework for analysis and is used in most of Glasserman's work.

Note that a GSMP can be viewed as a Markov chain where the state of the chain is comprised of the state of the GSMP (with denumerable state space) and of the event list (with the "planned" time of occurence of each event in the list). Glasserman (1991a, b, c) gives specific conditions for IPA to work under the GSMP framework.

5. DERIVATIVE OF THE STEADY-STATE AVERAGE COST

Consider again the Markov chain model of the previous section and suppose that it is regenerative, with finite second moment of the cycle length. For a given regenerative cycle, let τ be the length of the cycle (number of steps) and h_τ be the total cost during that cycle. Let $u(\theta) = E_\theta[h_\tau]$, $\ell(\theta) = E_\theta[\tau]$, and $\alpha(\theta) = u(\theta)/\ell(\theta)$. Then, from the renewal-reward theorem

$$\lim_{t\to\infty} \frac{1}{t}\sum_{i=1}^{t} C_i \overset{\text{a.s.}}{=} \alpha(\theta) = \lim_{t\to\infty} \frac{1}{t}\sum_{i=1}^{t} E_\theta[C_i]. \tag{24}$$

5.1. Estimators Based on the Regenerative Method

Suppose that one simulates n regenerative cycles of that process. Let Y_j and T_j denote the respective values of h_τ and τ for the j-th cycle. A consistent estimator of $\alpha(\theta)$ is then

$$\hat{\alpha}_n(\theta) = \frac{\sum_{j=1}^{n} Y_j}{\sum_{j=1}^{n} T_j}. \tag{25}$$

The derivative of α, assuming that it exists, is

$$\alpha'(\theta) = \frac{u'(\theta)\ell(\theta) - \ell'(\theta)u(\theta)}{\ell^2(\theta)} \tag{26}$$

$$= \frac{u'(\theta) - \alpha(\theta)\ell'(\theta)}{\ell(\theta)}. \tag{27}$$

This could be estimated by

$$\hat{\alpha}'_n(\theta) = \frac{\left(\sum_{j=1}^{n} Y'_j\right) - \hat{\alpha}_n(\theta)\sum_{j=1}^{n} T'_j}{\sum_{j=1}^{n} T_j}, \tag{28}$$

where Y'_j and T'_j are estimators of $u'(\theta)$ and $\ell'(\theta)$, respectively, based on the j-th cycle. Such estimators can be obtained, for example, by the LR method, as discussed in subsection 4.3.

Observe that with standard IPA, the derivative of τ will always be zero and will not be a valid estimator of $\ell'(\theta)$. However, Heidelberger et al. (1988) give conditions under which (28) is nevertheless a consistent estimator of $\alpha'(\theta)$. This looks surprising at first sight, but makes sense if the IPA estimator (21) is unbiased for the derivative of the expected average cost for the first t steps, for each fixed t, because of (24) and because (28) with IPA derivative estimators becomes the same as (21) with t replaced by the total number of steps in the n simulated cycles. For more on consistency of IPA, see also Glasserman (1991), Glasserman, Hu, and Strickland (1990), and Hu and Strickland (1991).

Heidelberger and Towsley (1989) have extended LR, in a regenerative setting, to the case where θ is unknown and its value in the derivative estimator is replaced by an estimate. They prove almost sure convergence and discuss the computation of confidence intervals.

5.2. Growing Horizon Estimators

Instead of exploiting the regenerative structure, one can use an estimator based on a fixed (but long) horizon t to estimate the steady-state derivative. As t goes to infinity, the bias goes to zero. Suppose n replications (of length t) are performed. Clearly, the optimal t should be a function of n, say t_n. L'Ecuyer (1991a) has obtained the optimal rate of increase of t_n for FD, FDC, IPA, and LR, under different sets of assumptions. As an example, in the context of the model of Section 4, for LR when the variance increases linearly in t, the optimal rate is t_n in the exact order of $n^{1/3}$ and the MSE for the steady-state gradient is then in $O(C^{-1/2})$, where $C = nt_n$ corresponds to the total computer budget. For IPA, in contrast, t_n in the exact order of n^p for any $p \geq 1$ is optimal and the MSE is then in $O(C^{-1})$. Since there is no upper bound on the optimal p, it is optimal, in particular, to allocate the total budget to just one very long run. Such an IPA derivative estimator is as efficient (in terms of order of convergence) as a standard *cost* estimator.

6. PERFORMANCE ASPECTS

In our Markov chain setup, we saw that IPA gives the best convergence rate for growing horizons. Even for fixed and short horizons, empirical evidence shows that IPA, when it applies, typically gives a lower variance than LR. However, there are applications where IPA does not seem to apply directly and LR works well. Such applications usually exploit short regenerative cycles. When IPA applies, FDC is essentially as good, but requires two runs instead of one (and still more runs when there are more parameters). Assessing the competitiveness of WD and FDE would require further investigation.

REFERENCES

Aleksandrov, V. M., Sysoyev, V. I., and Shemeneva, V. V. 1968. Stochastic Optimization. *Engineering Cybernetics*, **5**, 11–16.

Andradóttir, S. 1991. Optimization of the Steady-State Behavior of Discrete Event Systems. Department of Industrial Engineering, University of Wisconsin, Madison.

Asmussen, S. 1991. Performance Evaluation for the Score Function Method in Sensitivity Analysis and Stochastic Optimization. Manuscript, Chalmers University of Technology, Göteborg, Sweden.

Bacelli, F. and Brémaud, P. 1990. Virtual Customers in Sensitivity and Light Traffic Analysis via Campbell's Formula for Point Processes. To appear in *Advances in Applied Probabilities*.

Benveniste, A., Métivier, M., and Priouret, P. 1987. *Algorithmes Adaptatifs et Approximations Stochastiques*, Masson, Paris.

Brémaud, P. 1991. Maximal Coupling and Rare Perturbation Analysis. Manuscript.

Brémaud, P. and Vázquez-Abad, F. 1991. On the Pathwise Computation of Derivatives with respect to the Rate of a Point Process: the Phantom RPA method. *QUESTA*. To appear.

Chong, E. K. P. and Ramadge, P. J. 1990. Convergence of Recursive Optimization Algorithms using Infinitesimal Perturbation Analysis Estimates. Draft paper, Dept. of Electrical Engineering, Princeton University.

Glasserman, P. 1988. Performance Continuity and Differentiability in Monte Carlo Optimization. *Proceedings of the Winter Simulation Conference 1988*, IEEE Press, 518–524.

Glasserman, P. 1990. Stochastic Monotonicity, Total Positivity, and Conditional Monte Carlo for Likelihood Ratios. AT&T Bell Laboratories, Holmdel, New Jersey.

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

Glasserman, P. 1991b. Structural Conditions for Perturbation Analysis Derivative Estimation: Finite-Time Performance Indices. To appear in *Operations Research*.

Glasserman, P. 1991c. Derivative Estimates from Simulation of Continuous-Time Markov Chains. To appear in *Operations Research*.

Glasserman, P. and Gong, W. B. 1990. Smoothed Perturbation Analysis for a Class of Discrete Event Systems. *IEEE Trans. on Automatic Control*, **AC-35**, 11, 1218-1230.

Glasserman, P. and Yao, D. D. 1990. Some Guidelines and Guarantees for Common Random Numbers. AT&T Bell Laboratories, Holmdel, New Jersey.

Glasserman, P., Hu, J.-Q., and Strickland, S. G. 1990. Strongly Consistent Steady-State Derivative Estimates. To appear in *Probability in the Engineering and Information Sciences*.

Glynn, P. W. 1986. Stochastic Approximation for Monte-Carlo Optimization. *Proceedings of the Winter Simulation Conference 1986*, IEEE Press, 356–364.

Glynn, P. W. 1987. Likelihood Ratio Gradient Estimation: an Overview. *Proceedings of the Winter Simulation Conference 1987*, IEEE Press, 366–375.

Glynn, P. W. 1989a. Optimization of Stochastic Systems Via Simulation. *Proceedings of the Winter Simulation Conference 1989*, IEEE Press, 90–105.

Glynn, P. W. 1989b. A GSMP Formalism for Discrete Event Systems. *Proceedings of the IEEE*, **77**, 14–23.

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

Gong, W.-B. 1988. Smoothed Perturbation Analysis Algorithm for a $GI/G/1$ Routing Problem. *Proceedings of the 1988 Winter Simulation Comference*, IEEE Press, 525–531.

Gong, W. B. and Ho, Y. C. 1987. Smoothed (Conditional) Perturbation Analysis of Discrete Event Dynamical Systems. *IEEE Trans. on Automatic Control*, **AC-32**, 10, 858–866.

Heidelberger, P., X.-R. Cao, M. A. Zazanis, and R. Suri. 1988. Convergence Properties of Infinitesimal Perturbation Analysis Estimates. *Management Science*, **34**, 11, 1281–1302.

Heidelberger, P. and Towsley, D. 1989. Sensitivity Analysis from Sample Paths using Likelihoods. *Management Science*, **35**, 12, 1475–1488.

Ho, Y.-C. 1987. Performance Evaluation and Perturbation Analysis of Discrete Event Dynamic Systems. *IEEE Transactions on Automatic Control*, **AC-32**, 7, 563–572.

Ho, Y.-C. and Cao, X.-R. 1991. Discrete-Event Dynamic Systems and Perturbation Analysis. Kluwer Academic.

Ho, Y.-C., Eyler, A., and Chien, T. T. 1979. A Gradient Technique for General Buffer Storage Design in a Serial Production Line. *International Journal of Production Research*, **17**, 6, 557–580.

Ho, Y.-C. and Li, S. 1988. Extensions to the Perturbation Analysis Techniques for Discrete Event Dynamic Systems. *IEEE Transactions on Automatic Control*, **33**, 5, 427–438.

Ho, Y.-C. and Strickland, S. 1990. A Taxonomy of Perturbation Analysis Techniques. Manuscript, Harvard University.

Hu, J. Q. and Strickland, S. G. 1991. General Conditions for Strong Consistency of Sample Path Derivative Estimates. To appear in *Applied Mathematics Letters*.

Jacobson, S. H. 1991a. Convergence Results for Frequency Domain Gradient Estimators. Manuscript, Dept. Oper. Res., Case Western Reserve University, Cleveland.

Jacobson, S. H. 1991b. Variance and Bias Reduction Techniques for the Frequency Domain Gradient Estimators. Manuscript, Dept. Oper. Res., Case Western Reserve University, Cleveland.

Jacobson, S. H. and Schruben, L. W. 1991. A Simulation Optimization Procedure Using Harmonic Analysis. Manuscript, Dept. Oper. Res., Case Western Reserve University, Cleveland.

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

L'Ecuyer, P. 1990. A Unified Version of the IPA, SF, and LR Gradient Estimation Techniques. *Management Sciences*, **vol. 36**, No. 11, pp. 1364-1383.

L'Ecuyer, P. 1991a. Convergence Rates for Steady-State Derivative Estimators. To appear in the *Annals of Operations Research*.

L'Ecuyer, P. 1991b. On the Interchange of Derivative and Expectation for Likelihood Ratio Derivative Estimators. Submitted for publication.

L'Ecuyer, P. and Perron, G. 1990. On the Convergence Rates of IPA and FDC Derivative Estimators for Finite-Horizon Stochastic Systems. Submitted for publication.

L'Ecuyer, P. and Glynn, P. W. 1991. A Control Variate Scheme for Likelihood Ratio Gradient Estimation. In preparation.

L'Ecuyer, P., Giroux, N., and Glynn, P. W. 1991. Stochastic Optimization by Simulation: Convergence Proofs and Experimental Results for the GI/G/1 Queue in Steady-State. In preparation.

Luenberger, D. G. 1984. *Linear and Nonlinear Programming*, Addison-Wesley, second edition.

Meketon, M. S. 1987. Optimization in Simulation: a Survey of Recent Results. *Proceedings of the Winter Simulation Conference 1987*, IEEE Press, 58–67.

Métivier, M. and Priouret, P. 1984. Application of a Kushner and Clark Lemma to General Classes of Stochastic Algorithms. *IEEE Trans. on Information Theory*, **IT-30**, 2, 140–151.

Pflug, G. Ch. 1989. Sampling Derivatives of Probabilities. *Computing*, **42**, 315–328.

Pflug, G. Ch. 1990. On-line Optimization of Simulated Markovian Processes. *Math. of Operations Research*, **15**, 3, 381–395.

Pflug, G. Ch. 1991. *Simulation and Optimization: The Interface*. In preparation.

Reiman, M. I. and Weiss, A. 1989) Sensitivity Analysis for Simulations via Likelihood Ratios. *Op. Res.*, **vol. 37**, No. 5, pp. 830-844.

Rubinstein, R. Y. 1986a. *Monte-Carlo Optimization, Simulation and Sensitivity of Queueing Networks*, Wiley.

Rubinstein, R. Y. 1986b. The Score Function Approach for Sensitivity Analysis of Computer Simulation Models. *Math. and Computers in Simulation*, **28**, 351–379.

Rubinstein, R. Y. 1989. Sensitivity Analysis and Performance Extrapolation for Computer Simulation Models. *Operations Research*, **37**, 1, 72–81.

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 Shapiro, A. 1991. *Discrete-Event Systems: Sensitivity Analysis and Stochastic Optimization via the Score Function Method*, Wiley, To appear.

Simon, B. 1989. A New Estimator of Sensitivity Measures for Simulations Based on Light Traffic Theory. *ORSA Journal on Computing*, **1**, 3, 172–180.

Suri, R. 1987. Infinitesimal Perturbation Analysis of General Discrete Event Dynamic Systems. *J. of the ACM*, **34**, 3, 686–717.

Suri, R. 1989. Perturbation Analysis: The State of the Art and Research Issues Explained via the $GI/G/1$ Queue. *Proceedings of the IEEE*, **77**, 114–137.

Vázquez-Abad, F. J. and Kushner, H. 1991. A Surrogate Estimation Approach for Adaptive Routing in Communication Networks. Submitted for publication.

Vázquez-Abad, F. J. and L'Ecuyer, P. 1991. Comparing Alternative Methods for Derivative Estimation when IPA does not Apply Directly. In these Proceedings.

Wardi, Y., Gong, W.-B., Cassandras, C. G., and Kallmes, M. H. 1991a. A New Class of Perturbation Analysis Algorithms for Piecewise Continuous Sample Performance Functions. Submitted for publication.

Wardi, Y., Kallmes, M. H., Cassandras, C. G., and Gong, W.-B. 1991b. Smoothed Perturbation

Analysis Algorithms for Estimating the Derivatives of Occupancy-Related Functions in Serial Queueing Networks. Submitted for publication.

Wardi, Y., McKinnon, M. W., and Schuckle, R. 1991c. On Perturbation Analysis of Queueing Networks with Finitely Supported Service Time Distributions. To appear in *IEEE Transactions on Automatic Control*.

Zazanis, M. A. and Suri, R. 1988. Comparison of Perturbation Analysis with Conventional Sensitivity Estimates for Stochastic Systems. Manuscript.

Zhang, B. and Ho, Y.-C. 1991. Performance Gradient Estimation for Very Large Markov Chains. *IEEE Transactions on Automatic Control*, To appear.

AUTHOR BIOGRAPHY

PIERRE L'ECUYER is an associate 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 a professor in the computer science department, at Laval University, Ste-Foy, 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 an Associate Editor for *Management Science* and for the *ACM Transactions on Modeling and Computer Simulation*. He is also a member of ACM, IEEE, ORSA and SCS.