

Performance continuity and differentiability in Monte Carlo optimization

Paul Glasserman
Division of Applied Sciences
Harvard University
Cambridge, MA 02138

ABSTRACT

This paper describes a class of Monte Carlo optimization problems for which unbiased derivative estimators of the infinitesimal perturbation analysis (IPA) type can be derived; and also a simple framework within which to establish unbiasedness. Of central importance are systems with continuous, piecewise differentiable sample performance functions. Experience suggests that continuity is, in practice, almost necessary for IPA to work. "Piecewise" differentiable is a concession to the discrete nature of many applied probability models. We discuss a variety of examples, including both static and dynamic systems.

1. INTRODUCTION

The subject of Monte Carlo optimization has received considerable attention in recent years, primarily through advances in derivative estimation techniques for discrete event systems. A typical setting for such optimization problems is the following: It is desired to maximize $g(\theta) = E_\theta(\gamma)$ over $\theta \in \Theta$ where γ is interpreted as the performance associated with each realization of a stochastic system which depends on a parameter θ . Algorithms of the Robbins-Monro type search for a zero of the gradient by iterating according to

$$\theta_{n+1} = \theta_n + a_n \nabla \widehat{g}(\theta_n) \quad (1)$$

where $\{a_n\}$ is a sequence of step sizes and $\nabla \widehat{g}(\theta_n)$ is an estimate of $\nabla E_\theta(\gamma)$. Specific characteristics of $\widehat{\nabla g}$ vary; but one issue of particular importance is whether or not it is unbiased in the sense that

$$E[\widehat{\nabla g}(\theta)] = \nabla g(\theta) \quad \forall \theta \in \Theta. \quad (2)$$

Recent advances consist of finding easily computed estimators satisfying (2). See Meketon [9] for a survey and references.

In this paper, we consider only derivative estimators of the *infinitesimal perturbation analysis* type: Suppose \underline{X} is a random variable, vector or sequence whose distribution depends on θ – for now, a scalar – and

$$g(\theta) = E_\theta[\gamma]$$

where

$$\gamma = \Gamma(\underline{X}).$$

The idea is to (try to) construct a family of random elements $\underline{X}(\theta)$ on a single probability space in such a way that $\gamma(\theta) = \Gamma(\underline{X}(\theta))$ is almost surely differentiable with respect to θ . Then $g'(\theta)$ is estimated using $\gamma'(\theta)$. (It may be more descriptive to refer to these as "common random number" or "common probability space" derivative estimators (as in Glynn [5]) and to reserve the name "perturbation analysis" for the particular implementation associated with networks of queues.)

Our goal is to use some simple ideas from calculus to provide easily checked conditions for the unbiasedness of these derivative estimators. We do not attempt to formulate the most general conditions; necessary and sufficient conditions can be stated succinctly in terms of uniform integrability, but these are rarely immediately applicable. Instead, we seek conditions that are simple enough to be easily checked and cover enough cases to be interesting. Furthermore, by illustrating the use of this framework in a variety of examples, we suggest new areas of application.

2. PRELIMINARIES

The most important ingredient in the unbiasedness of IPA-like estimators is usually the almost sure continuity of γ as a function of θ . This was first noted in Cao [1] and is also the motivation for the approach in Gong and Ho [6]. It may seem strange that continuity should be an issue once differentiability is assumed; the point is that even if γ is a.s. differentiable for each θ in the parameter space Θ , it may happen that γ has discontinuities in Θ with probability one. An analogy can be made here with the sample paths of a Poisson process as functions of t : a Poisson process has discontinuous sample paths which are a.s. differentiable (with respect to t) at each $t > 0$.

Our approach is to focus on continuous, piecewise differentiable γ ; we believe this to be the most interesting class of functions with which IPA can be used. We also advocate a particular approach to proving unbiasedness in such cases.

The basic result is the following lemma. Let the parameter space Θ be a finite interval (θ_a, θ_b) and let γ be a random function of θ on a probability space (Ω, \mathcal{F}, P) . Let $D_\gamma(\omega)$ be the subset of Θ on which γ' exists.

Lemma 1. If γ is a.s. a continuous, piecewise differentiable function of Θ and

$$\sup_{\theta \in D_\gamma} |\gamma'(\theta)| \text{ is integrable,} \quad (3)$$

then

$$E[\gamma'(\theta)] = \frac{dE[\gamma(\theta)]}{d\theta}. \quad (4)$$

Proof. Part of the content of the lemma is that the derivative on the right in (4) exists. From a generalization of the mean value theorem (see, e.g., Dieudonne [2] p. 160).

$$\left| \frac{\gamma(\theta_0 + h) - \gamma(\theta_0)}{h} \right| \leq \sup_{\theta \in D_\gamma} |\gamma'(\theta)| \quad (5)$$

whenever θ_0 and $\theta_0 + h$ are in Θ . Equation (4) follows from (3) and the dominated convergence theorem.

Some remarks are in order about what is, so far, a fairly trivial idea. First, in virtually all interesting cases, it is easy to define $\gamma(\theta, \omega)$ so that γ is a.s. piecewise differentiable; thus, the more important condition is the continuity of γ across points where it fails to be differentiable. Second, many proofs of results like (4) rely on constructing a dominating function for $|h^{-1}(\gamma(\theta + h) - \gamma(\theta))|$. We maintain that the choice of (3) is particularly convenient. In many applications, γ turns out to be a complicated function of θ - though the essence of IPA is to observe that the local dependence of γ on θ may be quite simple. Thus, finding expressions for γ' or at least bounds on γ' is often straightforward. On the other hand, for fixed h , there may be points between θ and $\theta + h$ at which γ' does not exist. Comparing $\gamma(\theta + h)$ and $\gamma(\theta)$ across such points is usually difficult, making bounding γ' easier than directly bounding $|\gamma(\theta + h) - \gamma(\theta)|$.

We should mention that there is a simple extension of Lemma 1 for vector parameters and derivatives in the Frechet sense (using, for example, the mean value theorem on page 176 of Luenberger [8]). Gradients may be analyzed by considering each partial derivative separately, so from now on we consider only scalar θ .

3. FUNCTIONS OF FINITELY MANY VARIABLES

We now tailor Lemma 1 to functions of the form

$$\gamma(\theta) = \Gamma(X_1(\theta), \dots, X_n(\theta)), \quad (6)$$

where $\Gamma : \mathbf{R}^n \rightarrow \mathbf{R}$ and the X_i are independent real-valued random functions of θ . Usually, only the marginal distributions (with respect to θ) are specified, say via

$$P(X_i(\theta) \leq x) = F_i(x; \theta), \quad \forall \theta \in \Theta, \quad (7)$$

where the F_i are given parametric families of c.d.f.s.

We take $\Omega = [0, 1]^n$, \mathcal{F} the Borel sets on Ω , and P equal to product Lebesgue measure. Writing $\omega = (\omega_1, \dots, \omega_n)$, each X_i depends on ω only through ω_i . There are usually several ways of constructing X_i on $\Theta \times \Omega$ to satisfy (7). An especially interesting construction is available whenever F_i is strictly increasing, namely the inversion representation

$$X_i(\theta, \omega) = F_i^{-1}(\omega_i; \theta). \quad (8)$$

Under suitable conditions on F_i , (8) implies ([11],[5])

$$\frac{dX_i(\theta, \omega)}{d\theta} = -\frac{\partial F_i / \partial \theta}{\partial F_i / \partial x}(X_i(\theta, \omega); \theta). \quad (9)$$

For example, when

$$F_i(x, \theta) = 1 - e^{-x/\theta},$$

(8) becomes

$$X_i(\theta, \omega) = -\theta \ln(1 - \omega_i)$$

and (9) becomes

$$\frac{dX_i}{d\theta} = \frac{X_i}{\theta}.$$

We will not be directly concerned with (8) and (9) but instead assume simply

(A1). Every X_i is a.s. a continuously differentiable function of θ on Θ .

In addition, we sometimes impose the regularity condition

(A2). $|dX_i/d\theta| \leq B(|X_i| + 1)$ for some constant $B > 0$.

Finally, we mention that it is occasionally desirable to require the existence of a function \mathcal{D}_i such that

$$\frac{dX_i}{d\theta}(\theta, \omega) = \mathcal{D}_i(X_i(\theta, \omega), \theta)$$

since in this case it is possible to calculate $dX_i/d\theta$ (and γ') from observation of the X_i . (Note that this holds in (9); see also Suri and Zazanis [12], eqn.(4.4) and Glynn [5], Theorem 2.14.) However, we will not distinguish this case here.

We now return to (6). We suppose that

(A3). Γ is continuous,

and that for every $\theta \in \Theta$, Γ is a.s. differentiable at $\underline{X}(\theta) = (X_1(\theta), \dots, X_n(\theta))$. That is, letting D_Γ be the domain of differentiability of Γ , we assume

(A4). $\forall \theta \in \Theta$ $P(\underline{X}(\theta) \in D_\Gamma) = 1$.

Lemma 2. Under (A1), (A3) and (A4), γ is a.s. continuous and piecewise differentiable on Θ .

Proof. Continuity is immediate. Let $d = |\theta_b - \theta_a|$ (where

$(\theta_a, \theta_b) = \Theta$) and for $m = 1, 2, \dots$ and $k = 1, 2, \dots, m$ let A_{km} be the event

$$\{\gamma \text{ non-differentiable throughout } (\theta_a + \frac{(k-1)d}{m}, \theta_a + \frac{kd}{m})\}.$$

Note that $\gamma(\cdot, \omega)$ fails to be piecewise differentiable if and only if ω is in some A_{km} . But from (A1) and (A4), γ is a.s. differentiable at each $\theta \in (\theta_a, \theta_b)$ so each A_{km} has probability zero. Hence, $\cup_m \cup_k A_{km}$ has probability zero, and this is just the set of ω for which $\gamma(\cdot, \omega)$ fails to be piecewise differentiable.

What makes Lemma 2 interesting is that it relates the a.s. differentiability of γ separately for each θ to a property that holds simultaneously for all θ .

Piecewise differentiability is only half of what is needed to apply Lemma 1. Integrability of the supremum of $|\gamma'|$ can be assured through a variety of hypotheses. As an example, we consider

(A5). $|\partial\Gamma/\partial x_i|, i = 1, \dots, n$ bounded on D_Γ ;
and

(A6). $\sup_{\theta \in \Theta} |X_i(\theta)|$ integrable, $i = 1, \dots, n$.

We then have

Lemma 3. Under (A1),(A2),(A4),(A5) and (A6)

$$E[\sup_{\theta \in D_\Gamma} |\gamma'(\theta)|] < \infty.$$

Proof. From the first three hypotheses, on D_γ ,

$$\begin{aligned} |\gamma'(\theta)| &= \left| \sum_{i=1}^n \frac{\partial \Gamma}{\partial x_i} \frac{dX_i}{d\theta} \right| \\ &\leq B \sum_{i=1}^n \left| \frac{\partial \Gamma}{\partial x_i} \right| (|X_i| + 1). \end{aligned}$$

Now apply (A5) and (A6).

Ordinarily, taking a supremum over uncountably many random variables as in (A6) might make integrability questionable. But it is important to remember that the $X_i(\theta)$ at different values of θ are highly dependent. Indeed, in applications it commonly happens that each X_i is monotonic in θ . This is the case, for example, in the construction of the exponential distributions given above. If X_i is increasing in θ , then $\sup_{\Theta} X_i(\theta)$ is just $X_i(\theta_b)$. Moreover, it is interesting to note that the inversion representation (8) stochastically minimizes $\sup_{\Theta} X(\theta)$ over all joint distributions with marginals $F(\cdot; \theta)$. To see this, observe that, in general,

$$P(\sup_{\Theta} X(\theta) > x) \geq \sup_{\theta \in \Theta} P(X(\theta) > x) = \sup_{\theta \in \Theta} (1 - F(x; \theta)),$$

and that equality is obtained using (8).

Combining Lemmas 1-3, we have

Theorem 1. Under (A1)-(A6), γ as defined in (6) satisfies (4).

Assumptions (A2), (A5) and (A6) could be replaced by others. For example, the boundedness of $\partial\Gamma/\partial x_i$ could be relaxed by placing more stringent conditions on $dX_i/d\theta$. In addition, it is possible and often convenient to permit explicit dependence of Γ on θ .

We now turn to some examples to illustrate the use of these ideas.

Example 1. The "min" operation serves as a building block in many models (for example, determining the timing of events in generalized semi-Markov processes) and also serves to illustrate the role of continuous piecewise differentiable functions. Thus, let

$$\Gamma(x_1, x_2) = \min(x_1, x_2);$$

then Γ is continuous ((A3)) and

$$D_\Gamma = \{(x_1, x_2) : x_1 \neq x_2\}.$$

Let $X_1(\theta)$ and $X_2(\theta)$ be samples from $F_1(\cdot; \theta)$ and $F_2(\cdot; \theta)$. Suppose that

(A7). $F_i(\cdot; \theta)$ is continuous and $F_i(0; \theta) = 0$ for all $\theta \in \Theta$.

then $P((X_1(\theta), X_2(\theta)) \in D_\Gamma) = 1$ ((A4)). Also note that the partials of Γ take values in $\{0, 1\}$ so (A5) is in effect. If we assume the X_i satisfy (A1),(A2) and (A6), we conclude from Theorem 1 that

$$E\left[\frac{d}{d\theta} \min(X_1(\theta), X_2(\theta))\right] = \frac{d}{d\theta} E[\min(X_1(\theta), X_2(\theta))].$$

Example 2. Notice that in Example 1, $\gamma(\theta) = \Gamma(X_1(\theta), X_2(\theta))$ could be interpreted as the time to failure of a structure consisting of two components in series with lifetimes X_1 and X_2 . This example extends to more general structures: Let C_1, \dots, C_m be subsets of $\{1, \dots, m\}$ and let

$$\Gamma(x) = \max_{1 \leq j \leq m} \min_{i \in C_j} x_i$$

or

$$\Gamma(x) = \min_j \sum_{i \in C_j} x_i$$

and $\gamma(\theta) = \Gamma(X(\theta))$. Then Γ is continuous and has bounded partials on D_Γ . By adopting some of the hypotheses above, it is a simple matter to guarantee

$$E[\gamma'(\theta)] = \frac{d}{d\theta} E[\gamma(\theta)].$$

Derivative estimates for this type of system via likelihood ratios were considered in Rubinstein [10].

A closely related class of problems helps to illustrate the importance of continuity. Consider the probability that a system is functioning at a fixed time t . For convenience, consider the

simple case of two components in series:

$$\Gamma(x_1, x_2) = 1_{\{\min(x_1, x_2) > t\}},$$

$$\gamma(\theta) = \Gamma(X_1(\theta), X_2(\theta)).$$

Under (A1) and (A7), γ is a.s. differentiable at every θ . But since γ only takes values in $\{0, 1\}$, it is only continuous if it is constant. Furthermore, $\gamma'(\theta) = 0$ wherever it exists. Thus, except in the uninteresting case where $dE[\gamma]/d\theta \equiv 0$,

$$E[\gamma'(\theta)] \neq \frac{d}{d\theta} E[\gamma(\theta)].$$

The sample derivative γ' completely misses the effect of discontinuities in γ . (It should be pointed out that in this and similar examples, $dE[\gamma]/d\theta$ can usually be estimated using likelihood ratios.)

Example 3. Let $\Gamma(x_1, x_2) = |x_1 - x_2|$; Γ is differentiable except where $x_1 = x_2$. Suppose the c.d.f.s $F(\cdot; \theta)$ and $G(\cdot)$ (with $G(\cdot) \notin \{F(\cdot; \theta)\}$) satisfy (A7) and admit inversion representations $X(\theta)$ and Y satisfying (A1) and (A2). Suppose also that X satisfies (A6). Let $\gamma(\theta) = \Gamma(X(\theta), Y)$, with $X(\theta)$ and Y constructed simultaneously on $[0, 1]$. Then γ is a.s. continuous and piecewise differentiable, and on D_γ

$$|\gamma'(\theta)| \leq \left| \frac{dX}{d\theta}(\theta) \right| \leq B(|X(\theta)| + 1).$$

So under the stated hypotheses, from Lemma 1,

$$E[\gamma'(\theta)] = \frac{d}{d\theta} E[\gamma(\theta)].$$

Notice that

$$E[\gamma(\theta)] = \int_0^1 |F^{-1}(t; \theta) - G^{-1}(t)| dt. \quad (10)$$

This has the interpretation of a distance between $F(\cdot; \theta)$ and G . In this context, the problem of minimizing $E[\gamma]$ becomes the problem of fitting a member of $F(\cdot; \theta)$ to G .

Example 4. Consider the average waiting time of the first n customers in a GI/G/1 queue that starts empty. Let W_i be the waiting time of the i th customer and $\bar{W}_n = \frac{1}{n} \sum_{i=1}^n W_i$. If U_i and X_i are the interarrival and service times of customer i , then

$$W_{i+1} = \max\{0, W_i + X_i - U_i\}. \quad (11)$$

Let Γ_j , $j = 1, \dots, n$ be such that

$$W_j = \Gamma_j(U_1, \dots, U_j, X_1, \dots, X_j).$$

It is clear from (11) that every Γ_j is continuous.

Suppose now that the X_i are functions of a service time parameter θ satisfying (A1) with a c.d.f. satisfying (A7). Then the functions

$$\gamma_j(\theta) = \Gamma_j(U_1, \dots, U_j, X_1(\theta), \dots, X_j(\theta))$$

are a.s. continuous and piecewise differentiable. From (11), it is immediate that

$$|\gamma'_{j+1}(\theta)| \leq |\gamma'_j(\theta)| + \left| \frac{dX_j}{d\theta} \right|$$

and $\gamma'_1(\theta) \equiv 0$, since the first customer never waits. Thus,

$$|\gamma'_n(\theta)| \leq \sum_{i=1}^{n-1} (n-i) \left| \frac{dX_i}{d\theta} \right|.$$

Under (A2),

$$|\gamma'_n(\theta)| \leq Bn \sum_{i=1}^{n-1} \left| \frac{dX_i}{d\theta} \right|;$$

so from (A6) and Lemma 1, (4) holds for γ and also for \bar{W}_n . (See Zazanis [14] for a different approach to this example as well as convergence to steady state.)

4. FUNCTIONS OF RANDOMLY MANY VARIABLES

We now consider the case where \underline{X} is a sequence of random variables (now with $\Omega = [0, 1]^\infty$) but Γ depends on only finitely many at a time. More precisely, suppose we have functions $\{\Gamma_n\}_{n=1}^\infty$, $\Gamma_n : \mathbb{R}^n \rightarrow \mathbb{R}$ and a positive integer valued function N on \mathbb{R}^∞ , and consider

$$\gamma(\theta) = \Gamma_{N(\underline{X}(\theta))}(X_1(\theta), \dots, X_N(\theta)). \quad (12)$$

We assume that every Γ_n is continuous and also, for simplicity, everywhere differentiable. We further assume that $N(\underline{X}(\cdot))$ is, for each θ , a.s. constant in a neighborhood of θ . Under these assumptions and (A1), γ is a.s. differentiable and

$$\gamma'(\theta) = \sum_{i=1}^{N(\underline{X}(\theta))} \frac{\partial \Gamma_{N(\underline{X}(\theta))}}{\partial x_i} \frac{dX_i}{d\theta}. \quad (13)$$

Equations (12) and (13) are typical of situations arising in complex discrete event systems, and also help to emphasize a point made earlier about the *local* and *global* dependence of γ on θ . Locally, γ is just some Γ_n ; thus it is a simple matter to write (13) and from (13) to find bounds on $|\gamma'|$. But for any fixed h , it may happen that $N(\underline{X}(\theta + h)) \neq N(\underline{X}(\theta))$, in which case it may be difficult to bound $|\gamma(\theta + h) - \gamma(\theta)|$.

Bounding $|\gamma'|$ is only useful if γ is continuous. To guarantee this, we impose the following condition on $\{\Gamma_n\}$ and N : If $\{\underline{x}^n\}$ and $\{\underline{y}^n\}$ are sequences in \mathbb{R}^∞ converging pointwise to \underline{x} , then

$$\lim_{n \rightarrow \infty} \Gamma_{N(\underline{x}^n)}(x_1^n, \dots, x_{N(\underline{x}^n)}^n) = \lim_{n \rightarrow \infty} \Gamma_{N(\underline{y}^n)}(y_1^n, \dots, y_{N(\underline{y}^n)}^n). \quad (14)$$

Note that it is not necessary that

$$\lim_{n \rightarrow \infty} N(\underline{x}^n) = \lim_{n \rightarrow \infty} N(\underline{y}^n).$$

Under (14), γ is continuous, even across discontinuities in $N(\underline{X}(\cdot))$ – though γ' may fail to exist at such points.

Example 5. Fix constants $T > 0$ and $\{f_i\}$ and let

$$N(x) = \max\{n : \sum_{i=1}^n x_i < T\},$$

$$\Gamma_n(x_1, \dots, x_n) = \sum_{i=1}^n f_i x_i + (T - \sum_{i=1}^n x_i) f_{n+1}. \quad (15)$$

If N is discontinuous at x , then

$$\sum_{i=1}^n x_i = T$$

for some n ; and

$$\Gamma_n(x_1, \dots, x_n) = \Gamma_{n-1}(x_1, \dots, x_{n-1}).$$

Therefore, if $\gamma(\theta) = \Gamma_N(X_1(\theta), \dots, X_N(\theta))$, γ is continuous. If the f_i are bounded, $\sup |\gamma'|$ is easily dominated using (13) and perhaps (A2).

This example easily extends to

$$\gamma(\theta) = \int_0^T f(Z_t(\theta)) dt \quad (16)$$

where Z_t is a semi-Markov process whose sojourn times X_i depend on θ but whose embedded transition probabilities do not. In this case, (16) takes the form (15) but with coefficients f_i determined by the embedded chain.

Example 6. The continuity condition (14) is, unfortunately, often violated in examples of interest. Consider again the single server queue with interarrival sequence $\{U_n\}$ and service times $\{X_n\}$. Let N be the number of arrivals in the first busy cycle and

$$\Gamma_n = \sum_{i=1}^n W_i$$

with W_i defined by (11). Discontinuities of N correspond to points where busy cycles merge or break apart due to changes in service or interarrival times. At such points, Γ_N is also discontinuous: the accumulated waiting time in one busy period may jump to that in two busy periods (see also [12]).

5. GENERALIZED SEMI-MARKOV PROCESSES

So far, we have considered only fairly simple examples. More interesting and more general systems can be handled through the framework of generalized semi-Markov processes (GSMPs). We describe GSMPs briefly but refer to, e.g., Whitt [13] for details.

The ingredients of a GSMP are a countable state space \mathbf{S} ; a set $\mathbf{I} = \{1, \dots, I\}$ of events; an event list mapping $\mathcal{E} : \mathbf{S} \rightarrow 2^{\mathbf{I}}$ with the interpretation that $\mathcal{E}(s)$ is the set of possible events in state s ; transition probability functions $p : \mathbf{S} \times \mathbf{S} \times \mathbf{I} \rightarrow [0, 1]$ with the interpretation that with probability $p(s'; s, \alpha)$, the process enters state s' when event $\alpha \in \mathcal{E}(s)$ occurs in state s . Associated with each event α is a clock reading c_α ; at any time, the event with the smallest clock reading is the next to occur. New clock

readings for event α are drawn from a c.d.f. F_α whenever α is activated.

Denote by $\tau_0 \equiv 0, \tau_1, \tau_2, \dots$ the times at which events occur; i.e., the state transition epochs of the GSMP Z_t . Let $\{Y_i\}$ be the embedded process $Y_i = Z_{\tau_i}$ and suppose Z is right-continuous. For each $\alpha \in \mathbf{I}$, let $\{X_{\alpha i}\}_{i=1}^\infty$ be an i.i.d. sequence of F_α -distributed random variables, interpreted as clock samples for α . We suppose that the c.d.f. F_{α_0} of one event α_0 depends on θ .

We are interested in

$$\gamma(\theta) = \int_0^{\tau_n} f(Z_t(\theta)) dt. \quad (17)$$

Since n events occur in $[0, \tau_n]$, γ depends on θ only through (at most) $X_{\alpha_0 1}, \dots, X_{\alpha_0 n}$.

The function (17) is superficially similar to (16), but there is an important difference. For (16) we assumed that the sojourns of Z may change with θ but that the state sequence cannot. But in (17), θ enters through the clock c.d.f.s, not the state sojourns.

Thus, changes in θ can, in general, introduce changes in $\{Y_i\}$. Such changes potentially introduce discontinuities in γ ; thus, we must place some restrictions on the possible changes in $\{Y_i\}$ that can occur.

We now present two conditions that guarantee the continuity of γ . To simplify the presentation, we restrict attention to GSMPs satisfying

$$\forall s \in \mathbf{S} \forall \alpha \in \mathbf{I}, p(s'; s, \alpha) = p(s''; s, \alpha) > 0 \Rightarrow s' = s''. \quad (18)$$

In words, no two transitions out of a state due to the same event can have exactly the same probabilities (though the difference could be arbitrarily small).

The conditions are

- (C1). If $\{\alpha, \beta\} \subseteq \mathcal{E}(s)$ and $p(s'; s, \alpha) > 0$ then $\beta \in \mathcal{E}(s)$.
- (C2). If $\{\alpha, \beta\} \subseteq \mathcal{E}(s)$ and s_1 and s_2 are such that $p(s_1; s, \alpha)p(s_2; s_1, \beta) > 0$ then there is a state s_3 such that

$$p(s_3; s, \beta) = p(s_2; s_1, \beta)$$

and

$$p(s_2; s_3, \alpha) = p(s_1; s, \alpha).$$

Condition (C1) can be paraphrased as saying that no event deactivates another, condition (C2) as saying that – in a probabilistic sense – the order in which two events occur does not affect the resulting state. In particular, (C2) implies that when every transition probability is in $\{0, 1\}$, the state reached from s by the occurrence of α then β is the same as that reached by the occurrence of β then α .

It is a simple matter in general to construct a family of GSMPs $Z_i(\theta)$ so that each event occurrence time is a.s. piecewise differentiable; differentiability fails where event times cross. Conditions (C1) and (C2) essentially ensure the continuity of event times across event order changes. The idea behind them might be called “the principle of quick return to the right state;” it is included as part of a heuristic discussion in Glasserman [3]. Here, we just sketch the implications of (C1) and (C2); details can be found in [4] and in a forthcoming paper. In addition, conditions in the same spirit as (C1)-(C2) can be found in Li and Ho [7].

Even for GSMPs violating (C1)-(C2), γ as in (17) is usually a.s. piecewise differentiable. This is made clearer by writing γ as

$$\sum_{i=0}^{n-1} f(Y_i(\theta))[\tau_{i+1}(\theta) - \tau_i(\theta)], \quad (19)$$

and noting that $f(Y_i(\cdot))$ is piecewise constant. Possible discontinuities of γ (and the τ_i) occur at points θ where some $\tau_j(\theta^-) = \tau_{j+1}(\theta^-)$ – where two events may change order. For simplicity, here we ignore the possibility that more than two event times merge. Then, an important consequence of (C1) and (C2) is

Lemma 4. Under (C1),(C2),(A1) and (A7), $Z_i(\theta)$ can be constructed so that every τ_j is a.s. a continuous function of θ , and so that if $\tau_j(\theta) = \tau_{j+1}(\theta)$, then for $i = 0, \dots, n-1, i \neq j$, $Y_i(\theta^-) = Y_i(\theta^+)$.

The point of the lemma is that when two events change order, the only change in the state sequence involves the state Y_j which has a sojourn of length $\tau_{j+1} - \tau_j = 0$.

With Lemma 4, we can easily prove

Theorem 2. Suppose that (C1) and (C2) hold, and also that (A1),(A2) and (A6) are satisfied by the α_0 clock samples $\{X_{\alpha_0 i}(\theta)\}$, $i = 1, \dots, n$, and (A7) by the $F_{\alpha_0}(\cdot; \theta)$ c.d.f.s. If, in addition, f is bounded, then

$$E[\gamma'(\theta)] = \frac{d}{d\theta} E[\gamma(\theta)],$$

with γ as in (17).

Proof. Under (C1),(C2),(A1) and (A7), γ is a.s. continuous. For even if $\tau_j(\theta) = \tau_{j+1}(\theta)$,

$$\begin{aligned} \gamma(\theta^+) - \gamma(\theta^-) &= \sum_{i=0}^{n-1} [f(Y_{i+1}(\theta^+)) - f(Y_i(\theta^-))][\tau_{i+1}(\theta) - \tau_i(\theta)] \\ &= [f(Y_{j+1}(\theta^+)) - f(Y_j(\theta^-))][\tau_{j+1}(\theta) - \tau_j(\theta)] \\ &= 0. \end{aligned}$$

Next, if K bounds $|f|$, then on D_γ ,

$$\begin{aligned} |\gamma'(\theta)| &= \sum_{i=0}^{n-1} f(Y_i(\theta)) \left| \frac{d\tau_{i+1}}{d\theta} - \frac{d\tau_i}{d\theta} \right| \\ &\leq 2K \sum_{i=0}^{n-1} \sum_{k=1}^{i+1} \left| \frac{dX_{\alpha_0 k}}{d\theta} \right| \\ &\leq 2KB \sum_{i=0}^{n-1} \sum_{k=1}^{i+1} |X_{\alpha_0 k} + 1|. \end{aligned}$$

Then, under (A6), $\sup_{D_\gamma} |\gamma'|$ is integrable. The result now follows from Lemma 1.

Generalizations of this theorem can be found in [4] and in a forthcoming paper. The extensions allow τ_n in (17) to be replaced by either a fixed time or the time of the n th occurrence of a fixed event. (Unfortunately, it is rarely possible to replace τ_n with a regeneration epoch and still maintain continuity.) Other extensions allow the introduction of “speeds” – state-dependent clock rates.

We close this section by mentioning some examples that satisfy (C1) and (C2). Consider a GI/G/1 queue where the state is the number in the system and the events are arrivals and departures. Condition (C1) is obviously satisfied; to see that (C2) is also, observe that an arrival followed by a departure leaves the same state as a departure followed by an arrival. Similarly, these conditions are satisfied by open and closed networks of FCFS queues with a single class of customers and state-independent routing.

On the other hand, (C2) is violated by the GI/G/1/K queue: in $s = K$, an arrival followed by a departure leaves $K - 1$ customers (since the arrival is lost), whereas the events taken in the opposite order leave K customers (since the departure makes room for the arrival). Many networks of queues with multiple classes of customers also violate (C2).

6. CONCLUSIONS

We have emphasized the role of sample performance continuity in the consistency of infinitesimal perturbation analysis derivative estimates, and presented a framework for establishing consistency when performance is in fact continuous. We maintain that dominating the supremum of the sample derivative is a fruitful approach, usually preferable to dominating finite differences. A variety of examples of the use of these ideas were given. Of particular interest are simple conditions for the consistency of a class of IPA estimates associated with generalized semi-Markov processes.

ACKNOWLEDGEMENTS

Research supported by NSF grants ECS-85-15449 and CDR-85-001-08, by ONR contracts N00014-84-K-D465 and -86-K-0075, and by DAAL contract 03-86-K-0171.

REFERENCES

- [1] Cao, X.R., 'Convergence of parameter sensitivity estimates in a stochastic experiment,' *IEEE Transactions on Automatic Control*, **AC-30**, pp.845-853, 1985.
- [2] Dieudonne, J.A., *Foundations of Modern Analysis*, Academic Press, New York, 1960.
- [3] Glasserman, P., 'Infinitesimal perturbation analysis of a birth and death process,' *Operations Research Letters*, **7**, no.1, 1988.
- [4] Glasserman, P., 'Equivalence methods in the perturbation analysis of queueing networks', Ph.D. Thesis, Harvard University, 1988.
- [5] Glynn, P.W., 'Process-differentiable representations of parametric families of random variables', Technical Report, Department of Industrial Engineering, University of Wisconsin - Madison, 1987.
- [6] Gong, W.B. and Ho, Y.C., 'Smoothed perturbation analysis of discrete-event dynamic systems', *IEEE Transactions on Automatic Control*, **AC-32**, pp.858-867, 1987.
- [7] Li, S. and Ho, Y.C., 'Sample path and performance homogeneity of discrete event dynamic systems', Manuscript, Systems and Industrial Engineering, University of Arizona, 1988.
- [8] Luenberger, D.G., *Optimization by Vector Space Methods*, Wiley, New York, 1969.
- [9] Meketon, M.S., 'Optimization in simulation: a survey of recent results', *Proc. of the 1987 Winter Simulation Conference*, A. Thesen, H. Grant and W.D. Kelton, ed.s.
- [10] Rubinstein, R., 'The score function approach for sensitivity analysis of computer simulation models', *Mathematics and Computers in Simulation*, 1986, pp.351-379.
- [11] Suri, R., 'Infinitesimal perturbation analysis for general discrete event systems', *Journal of the Association for Computing Machinery*, pp.686-717, 1987.
- [12] Suri, R. and Zazanis, M., 'Perturbation analysis gives strongly consistent sensitivity estimates for the M/G/1 queue', *Management Science*, **34**, pp.39-64, 1988.
- [13] Whitt, W., 'Continuity of generalized semi-Markov processes', *Mathematics of Operations Research*, **5**, pp.494-501, 1980.
- [14] Zazanis, M., 'Weak convergence of sample path derivatives for the waiting time in a single server queue', *Proceedings of the 25th Allerton Conference on Communication, Control and Computing*, 1987.

AUTHOR'S BIOGRAPHY

PAUL GLASSERMAN received an A.B. in Mathematics from Princeton University in 1984 and a Ph.D. in Applied Mathematics from the Division of Applied Sciences, Harvard University in 1988. In July of 1988 he joined the Operations Research Department at AT&T Bell Laboratories in Holmdel, New Jersey. His research centers on optimization through simulation.

Room 3K-324
AT&T Bell Laboratories
Holmdel, New Jersey 07733
(201) - 949 - 0669