ARTICLES

STRUCTURAL CONDITIONS FOR PERTURBATION ANALYSIS DERIVATIVE ESTIMATION: FINITE-TIME PERFORMANCE INDICES

PAUL GLASSERMAN

Columbia University, New York, New York

(Received July 1989; revisions received October 1990, March 1991; accepted April 1991)

In recent years, there has been a surge of research into methods for estimating derivatives of performance measures from sample paths of stochastic systems. In the case of queueing systems, typical performance measures are mean queue lengths, throughputs, etc., and the derivatives estimated are with respect to system parameters, such as parameters of service and interarrival time distributions. Derivative estimates potentially offer a general means of optimizing performance, and are useful in sensitivity analysis. This paper concerns one approach to derivative estimation, known as *infinitesimal perturbation analysis.* We first develop a general framework for these types of estimates, then give simple sufficient conditions for them to be unbiased. The key to our results is identifying conditions under which certain finite-horizon performance measures are almost surely continuous functions of the parameter of differentiation throughout an interval. The sufficient conditions we introduce are formulated in the setting of generalized semi-Markov processes, but translate into readily verifiable conditions for queueing systems. These results substantially extend the domain of problems in which infinitesimal perturbation analysis is provably applicable.

Most real-world queueing systems violate the rather restrictive conditions necessary to obtain exact analytic results, so networks of queues are often studied through discrete-event simulation. Simulation has the advantage of allowing complete model generality, but has the drawback of being computationally intensive. Hence, there is much to be gained from methods that make more efficient use of simulation by, for example, extracting more information from each run. Particularly valuable are methods that offer the possibility of optimization and sensitivity analysis, since these are the ultimate goals of most performance analysis.

One way to use simulation for optimization is through a stochastic version of a gradient search method—stochastic approximation—driven by gradient estimates obtained through simulation; see, e.g., Kushner and Clark (1978) for background on stochastic approximation. The viability of such an approach depends heavily on the ease with which good gradient estimates can be obtained. The past few years have witnessed significant advances in the development of efficient techniques for Monte Carlo gradient estimation for discrete-event systems, such as queueing networks. In this paper, we give verifiable sufficient conditions for the use of one of the more efficient techniques, *infinitesimal perturbation analysis* (IPA). Another approach, not touched on here, is considered in Glynn (1986), Reiman and Weiss (1989), and Rubinstein (1989).

Briefly, the idea of perturbation analysis is as follows: Let $\{Z(t, \theta), 0 \leq t \leq T, \theta \in \Theta\}$ be a parametric family of stochastic processes, and let $L(\theta)$ be a performance measure evaluated on each sample path $\{Z(t, \theta), 0 \leq t \leq T\}$. Suppose that the processes $\{Z(t, \theta), 0 \leq t \leq T, \theta \in \Theta\}$ are defined on a common probability space in such a way that L is, with probability one, differentiable in θ. Then the random variable $dL/d\theta$ is an IPA estimate of $dE[L(\theta)]/d\theta$. An IPA algorithm calculates the *exact* value of $dL/d\theta$ evaluated at, say θ_0, from a sample path $\{Z(t, \theta_0), 0 \leq t \leq T\}$ at θ_0 only, i.e., without ever actually perturbing θ. The IPA estimate $dL/d\theta$ is unbiased if

$$E\left[\frac{dL}{d\theta}\right] = \frac{dE[L]}{d\theta}. \tag{1}$$

The left side is what is obtained, in the limit, by averaging independent replications of $dL/d\theta$, but the right side is what is needed for optimization of $E[L]$.

Subject classifications: Queues: sample path analysis and optimization. Simulation: gradient estimation techniques.

0030-364X/91/3905-0724 $01.25

It has been understood for some time that (1) holds in some contexts and not others—that perturbation analysis is not universally applicable. In this paper, we introduce very general and surprisingly simple conditions for the consistency of a class of finite-horizon perturbation analysis estimates. At the heart of our results are conditions on Z, L and θ that ensure that L is, with probability one, a continuous function of θ. For the kinds of functions that commonly arise as performance measures in discrete-event simulations, ensuring continuity is the most important step in ensuring that (1) holds.

Along the way, we address some foundational issues and unify, in a general framework, many of the special cases of perturbation analysis estimates previously considered in the literature. Our formulation of IPA is similar to that developed in Suri (1987) in a different setting. Although we are mainly interested in applications to queueing systems, we find it convenient to work within the framework of *generalized semi-Markov processes* (especially, the formulation in Whitt 1980). This framework allows considerable generality, and, more importantly, permits us to separate the structural aspects of a discrete-event system from the distributions that drive it. Our main condition is, in fact, purely structural.

The essential feature of a generalized semi-Markov process is that it moves from state to state through the occurrence of "events." In a queueing context, a state might describe the arrangement of customers in queues; examples of events are service completions and arrivals of customers. With this rough description, our main condition can be paraphrased as requiring that the state reached from another state through the occurrence of two events be independent of their order. It has been observed widely that when IPA fails, it is typically because changes in a parameter change the order of events in such a way as to introduce discontinuities in the sample performance L. Our conditions guarantee the continuity of a class of performance measures even across event order changes.

Few other *general* results on the unbiasedness and consistency of IPA estimates are available. Throughput in Jackson networks is considered in Cao (1988); waiting time in the M/G/1 queue is discussed in Suri and Zazanis (1988). Necessary conditions for a class of throughput derivatives, and necessary and sufficient conditions for derivatives based on regenerative cycles are given in Heidelberger et al. (1988). These papers propose nothing like our main condition, which grew out of an argument in Glasserman (1988), Section 4, for the special case of a birth–death process. A related generalization of this argument, arrived at independently, is reported in Li and Ho (1989), but the conditions there are not purely dependent on system structure. There is, moreover, no overlap between our results and those of Heidelberger et al. This is partly because we consider a different class of performance measures, but, more importantly, because their conditions are stated in terms of possible equalities between unknown quantities and can only be checked in special cases. Our conditions are easy to check. The conditions in Heidelberger et al. are probably best suited to identifying cases where IPA is unlikely to work, whereas our emphasis is on understanding those cases where it does.

To prepare the way for considering sample path derivative estimates, in Section 1 we define and construct generalized semi-Markov processes. Working at this level of generality requires introducing a bit of notation, but this is necessary for a concise statement of our main condition. In Section 2, we derive IPA estimates for a broad class of finite-horizon performance measures. In Section 3, we introduce sufficient conditions for these estimates to be unbiased. Section 4 considers an example; Section 5 discusses an extension of the results of Section 3. Section 6 contains some concluding remarks.

1. THE GENERALIZED SEMI-MARKOV PROCESS FRAMEWORK

1.1. Basic Description

Generalized semi-Markov processes—GSMPs, for short—provide a broad framework ideally suited for the consideration of IPA estimates. Originally introduced to study the phenomenon of *insensitivity* (as in Schassberger 1978), GSMPs have turned out to be a powerful tool for analyzing discrete-event simulation because their dynamics mimic the evolution of such simulations. Even when the applications of interest are networks of queues, the generality of the GSMP model is useful; see Glynn and Iglehart (1988) for an overview of simulation methods for queues using the GSMP framework. In the case of perturbation analysis derivative estimates, it would be difficult to state general and succinct conditions for consistency without something like a GSMP. In particular, the notion of *event* seems essential to an understanding of when perturbation analysis works.

A brief description of a GSMP goes as follows: The states of a GSMP represent possible "physical" configurations of a system, which need not be states in the

Markovian sense. In a queueing context, the state may be simply a vector of queue lengths, perhaps supplemented by information about the classes of customers in queue, which servers are blocked, etc. The process jumps from state to state upon the occurrence of "events;" for us, the most important events will be departures from and external arrivals to queues. The state to which the process moves when an event occurs is governed by a set of transition probabilities. In a queueing network, these determine the routing of customers. Just when events occur is determined by random *clocks* associated with the possible events in a state. Each clock represents the time remaining until the associated event occurs, so the event with the shortest remaining clock time is the next to occur. When, for example, the events are arrivals to and departures from a queue, the initial settings of the respective clocks are simply interarrival and service times. After being set, all clocks are run down at unit rate. When a clock runs out, the corresponding event occurs, the process changes state, and new clocks may be set for new events possible in the new state. (Below we will require that all clocks from the previous state continue to run down; in the more general settings of, e.g., Whitt 1980 and Glynn and Iglehart 1988, the occurrence of one event may interrupt clocks for other events.)

To characterize a GSMP we need the following elements:

S = a state space (finite or countably infinite) representing the set of physical states of a system;
A = a finite subset of the integers enumerating the events; typical events will be denoted by α and β;
$\mathscr{E}(s)$ = the set of possible events (the *event list*) is state s; for example, departure from a queue is only a possible event in those states in which the queue is busy; we do not allow $\mathscr{E}(s)$ to be empty;
$p(s'; s, \alpha)$ = the probability of jumping to s' from s when event α occurs;
$F_\alpha(\cdot)$ = the distribution of new clock samples for events of type α; if α is an external arrival to a queue, then F_α is the interarrival time distribution; if α is the departure from a server, then F_α is the service time distribution.

We now show how to use the GSMP framework to model some simple systems; these examples will be useful later.

Example 1. GI/G/1 Queue. Take **S** to be the nonnegative integers (the set of possible queue lengths); let α denote arrival and β denote departure. Then F_α and F_β are the interarrival and service time distributions; $\mathscr{E}(s) = \{\alpha, \beta\}$ if $s > 0$, and $\mathscr{E}(0) = \{\alpha\}$. Also, $p(s+1; s, \alpha) = 1$ and $p(s-1; s, \beta) = 1$ (for $s > 0$) while all other transition probabilities are zero.

Example 2. Closed Jackson-Like Networks. Let **S** be the set of possible queue-length vectors $s = (n_1, \ldots, n_M)$, where n_i is the number at queue i and M is the number of servers. Let β_i denote departure from server i. Then F_{β_i} is the service time distribution at server i, and $\beta_i \in \mathscr{E}(s)$ if and only if $n_i > 0$ in s. Suppose that the routing in the network is Markovian in the sense that with probability P_{ij} customers leaving server i join queue j (independent of everything else). Then p is given by $p(s - e_i + e_j; s, \beta_i) = P_{ij}$, where e_i is the ith unit vector.

1.2. Construction of a GSMP

In order to consider sample path derivatives associated with a GSMP we need an explicit construction of the sequences of states, events and jump epochs that characterize a sample path. The construction, though seemingly intricate, amounts to little more than a generic algorithm for a discrete-event simulation. Presenting the construction explicitly will allow us to investigate the effect of small changes in the clock samples on the timing of events. The construction is greatly simplified if we impose the following from the outset (it would, in any case, be needed for our main results):

C1. (Noninterruptive Condition) *For every* $s, s' \in \boldsymbol{S}$ *and* $\alpha \in \boldsymbol{A}$, *if* $\alpha \in \mathscr{E}(s)$ *and* $p(s'; s, \alpha) > 0$, *then* $\mathscr{E}(s) - \{\alpha\} \subseteq \mathscr{E}(s')$.

This condition requires that the occurrence of one event not interrupt clocks for other events. (It is also used, for entirely unrelated reasons, in Schassberger 1976.) This condition excludes preemptive mechanisms in which the occurrence of one event deactivates another pending event.

We denote the GSMP itself by $Z(t)$. We need additional notation for various sample path characteristics. For easy reference, we provide the following informal descriptions; precise definitions are given via the recursions below.

τ_n = the epoch of the nth state transition;
a_n = the nth event;
Y_n = the nth state visited by the process: $Y_n = Z(\tau_n^+)$;

c_n = the vector of clock readings at τ_n^+;
$c_n(\alpha)$ = at τ_n, the time remaining until α occurs, provided $\alpha \in \mathscr{E}(Y_n)$;
$N(\alpha, n)$ = the number of instances of α among $a_1, \ldots, a_n$.

We now construct Z. Our basic inputs are two doubly-indexed sequences of independent random variables $\{X(\alpha, k), \alpha \in \mathbf{A}, k = 1, 2, \ldots\}$ and $\{U(\alpha, k), \alpha \in \mathbf{A}, k = 1, 2, \ldots\}$. For each k, $X(\alpha, k)$ is distributed according to F_α and represents the kth clock sample for α (e.g., the kth service or interarrival time). Every *routing indicator* $U(\alpha, k)$ is uniformly distributed on [0, 1] and will be used to determine the state transition at the kth occurrence of α. Think of each $U(\alpha, k)$ as a random number used to sample from a set of transition probabilities. Transitions are then determined by a mapping ϕ: $\mathbf{S} \times \mathbf{A} \times [0, 1] \to \mathbf{S}$: If α occurs in state s with routing indicator u, the process jumps to $s' = \phi(s, \alpha, u)$. The only condition we require of ϕ is that for all s, α, s':

$$P(\phi(s, \alpha, U) = s') = p(s'; s, \alpha)$$

whenever U is uniformly distributed on [0, 1]. For now, we assume that ϕ is given. Later, when we need it (in Sections 3 and 4) we will define a particular ϕ.

Choose an initial state Y_0 and initialize by setting $\tau_0 = 0$ and every $N(\alpha, 0) = 0$. Set clocks for the possible events: If $\alpha \in \mathscr{E}(Y_0)$, then set $c_0(\alpha) = X(\alpha, 1)$. Now repeat the following recursions:

$$\tau_{n+1} = \tau_n + \min\{c_n(\alpha): \alpha \in \mathscr{E}(Y_n)\} \tag{2}$$

$$a_{n+1} = \min\{\alpha \in \mathscr{E}(Y_n): c_n(\alpha) = \min\{c_n(\alpha'): \alpha' \in \mathscr{E}(Y_n)\}\} \tag{3}$$

(this definition picks out a unique $n + 1$st event a_{n+1} even when multiple events occur simultaneously);

$$N(\alpha, n+1) = \begin{cases} N(\alpha, n) + 1, & \alpha = a_{n+1} \\ N(\alpha, n), & \text{otherwise} \end{cases} \tag{4}$$

$$Y_{n+1} = \phi(Y_n, a_{n+1}, U(a_{n+1}, N(a_{n+1}, n+1))). \tag{5}$$

At each state transition, the clock readings are adjusted by setting clocks for any "new" events and reducing the time left on any "old" clocks by the time since the last transition. Thus, if $\alpha \in \mathscr{E}(Y_n)$ and $\alpha \neq a_{n+1}$, then under (C1), $\alpha \in \mathscr{E}(Y_{n+1})$ and

$$c_{n+1}(\alpha) = c_n(\alpha) - (\tau_{n+1} - \tau_n). \tag{6}$$

If $\alpha \in \mathscr{E}(Y_{n+1})$ and either $\alpha \notin \mathscr{E}(Y_n)$ or $\alpha = a_{n+1}$, then

$$c_{n+1}(\alpha) = X(\alpha, N(\alpha, n+1) + 1). \tag{7}$$

From these recursions we define Z by setting $Z(t) = Y_n$ on $[\tau_n, \tau_{n+1})$.

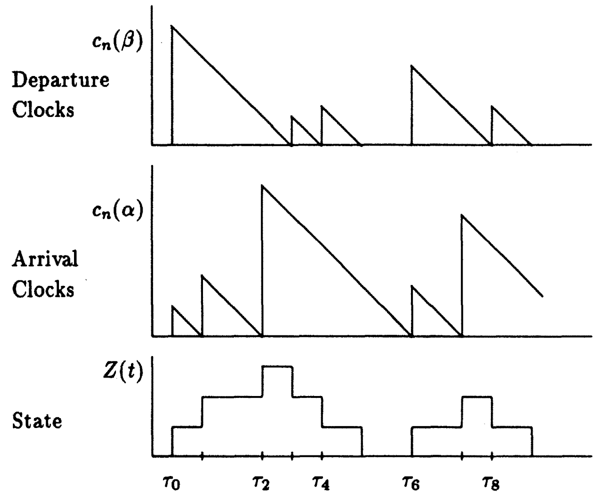

Figure 1. GSMP view of the GI/G/1 queue.

Figure 1 illustrates how clocks for arrivals and departures drive the state of a single server queue. The vertical jumps of the clock processes correspond to the setting of new clocks; thus, the kth jump of $c_n(\alpha)$ would have height $X(\alpha, k)$. When a departure clock runs out, the queue length, Z, jumps down one unit (corresponding to a departure); and when an arrival clock runs out, it jumps up one unit. A new clock is set only when another runs out. In particular, the occurrence of an arrival always causes a new arrival clock to be set; but a departure starts a new departure clock only if it leaves behind a nonempty queue. Thus, at τ_5, no new departure clock is set; it must wait until the next arrival, which occurs at τ_6.

To further specify characteristics of the sample paths of Z, we define $T(\alpha, k)$ to be the epoch of the kth occurrence of event α. That is, $T(\alpha, k)$ is equal to τ_{n^*} where

$$n^* = \min\{n \geq 0: N(\alpha, n) = k\}.$$

If the event α does not occur k times, then

$$T(\alpha, k) = \infty.$$

In comparing different sample paths, it is useful to be able to identify corresponding events on the two paths. We do this by identifying the kth occurrence of α on one path with the kth occurrence on the other, and so on. Call such an (α, k) an *event-order pair*, and if $r = (\alpha, k)$ write $T(r)$ for $T(\alpha, k)$. Now consider the nth event on some path. By definition, this event is a_n and this is its kth occurrence with $k = N(a_n, n)$. Hence, if we define $r_n = (a_n, N(a_n, n))$, then the first component of r_n is the type of the nth event and the second component is the number of times this event has occurred. Call r_n the nth event-order pair. Since $T(r_n)$ is just the epoch of the nth event, we always have $T(r_n) = \tau_n$.

We will find it convenient to work with a more explicit representation of the jump epochs τ_n and $T(\alpha, k)$. To obtain such a representation, the key observation is that every τ_n and $T(\alpha, k)$ is the sum of a subset of the $\{X(\beta, j), \beta \in \mathbf{A}, j = 1, 2, \ldots\}$. In (2), the clock that runs out and determines the jump at τ_{n+1} was set at some earlier transition τ_m, $m \leq n$ to some clock sample $X(\beta, j)$ (in particular, with $\beta = a_{n+1}$ and $j = N(\alpha_{n+1}, n+1)$) so that $\tau_{n+1} = \tau_m + X(\beta, j)$. In general, working backward in this way from any τ_n we can find a sequence of events and indices $(\beta_1, j_1), \ldots, (\beta_{m_n}, j_{m_n})$ such that

$$\tau_n = X(\beta_1, j_1) + \ldots + X(\beta_{m_n}, j_{m_n})$$

and such that the j_ith occurrence of event β_i triggers the setting of the j_{i+1}th clock of event β_{i+1}. If τ_n is the epoch of the transition caused by the kth occurrence of α (i.e., $a_n = k$ and $N(\alpha, n) = k$), call this $(\beta_1, j_1), \ldots, (\beta_{m_n}, j_{m_n})$ the *triggering sequence* for (α, k). To pick out which event-order pairs are in the triggering sequence for (α, k) we use indicators that take only the values zero and one.

Definition 1. Suppose that **C1** holds. The *triggering indicators* $\eta(\cdot, \cdot; \cdot, \cdot)$ are equal to zero except as follows:

1. for every α and k, $\eta(\alpha, k; \alpha, k) = 1$;
2. if the kth clock for α is set at the jth occurrence of β, then $\eta(\alpha, k; \beta, j) = 1$;
3. if $\eta(\alpha, k; \beta, j) = 1$ and $\eta(\beta, j; \beta', j') = 1$, then $\eta(\alpha, k; \beta', j') = 1$.

Intuitively, $\eta(\alpha, k; \beta, j) = 1$ indicates that a small delay in the jth occurrence of β delays the kth occurrence of α by the same amount. The following is an immediate consequence of the definition of η.

Lemma 1. *Suppose that* **C1** *holds. With $T(\alpha, k)$ the epoch of the kth occurrence of α and τ_n the epoch of the nth state transition, we have the following. For every $\alpha \in \mathbf{A}$ and every $k > 0$, if $T(\alpha, k) < \infty$, then*

$$T(\alpha, k) = \sum_{\beta, j} X(\beta, j)\eta(\alpha, k; \beta\, j). \tag{8}$$

For every $n \geq 0$, if $\tau_n < \infty$, then

$$\tau_n = \sum_{i=1}^{n} X(r_i)\eta(r_n; r_i). \tag{9}$$

In Figure 1 we see that every arrival (α) triggers the setting of the next arrival clock; hence, the only $\eta(\alpha, k; \cdot, \cdot)$ equal to one are of the form $\eta(\alpha, k; \alpha, j)$ with $j \leq k$. But departures (β) may have both arrivals and departures in their triggering sequences. For example, the last event is the departure at τ_9. The service time that ends at τ_9 was initiated at τ_8 when the previous customer departed. Hence, the departure at τ_8 is in the triggering sequence for the departure at τ_9. Continuing backward, the service time that ends at τ_8 is initiated by the arrival at τ_6; hence, that arrival is also in the triggering sequence. Since each interarrival clock is set by the previous arrival, all arrivals prior to τ_6 are in the triggering sequence. In this way we get

$$\tau_9 = X(\alpha, 1) + X(\alpha, 2) + X(\alpha, 3) + X(\beta, 4) + X(\beta, 5).$$

This is checked in the figure by adding the corresponding interarrival and service times along the time axis.

An important observation is that the triggering sequences and indicators are determined by the order in which events occur, but do not depend on the particular epochs of their occurrence.

Triggering here corresponds to *scheduling* in Suri.

2. DERIVATIVE ESTIMATES FOR PARAMETRIC GSMPs

With the construction of the previous section, we can calculate derivatives of sample performance measures for GSMPs that depend on a parameter. The derivative expressions we derive generalize and unify those in, for example, Ho and Cao (1983) and Cao (1988). They are similar to those formulated in a different framework in Suri. We consider the case where some or all of the clock setting distributions depend on a scalar parameter θ in a finite interval $\Theta = (\theta_a, \theta_b)$. Vector parameters are handled by considering each component separately. *We do not allow the transition probabilities $p(\cdot\,; \cdot, \cdot)$ to depend on θ.* For emphasis, we sometimes write $F_\alpha(x, \theta)$ for the α clock-setting distribution; we also write $X_\theta(\alpha, k)$ as a reminder that the clock samples themselves depend on θ. For each α and k, we view $X_\theta(\alpha, k)$ as a random function of θ satisfying

$$P(X_\theta(\alpha, k) \leq x) = F_\alpha(x, \theta) \quad \text{for all } \theta \in \Theta.$$

In practice—especially in simulation—this is often achieved by setting

$$X_\theta(\alpha, k) = F_\alpha^{-1}(U; \theta) \tag{10}$$

where U is uniformly distributed on the unit interval.

To consider derivatives, we need some conditions on the clock samples and their distributions:

A1. For each $\theta \in \Theta$ and $\alpha \in \mathbf{A}$, $F_\alpha(x, \theta)$ is continuous in x and zero at $x = 0$.

A2. For every α and k, $X_\theta(\alpha, k)$ is, with probability one, a continuously differentiable function of θ on Θ.

A3. There exists a constant $B > 0$ such that for all $\theta \in \Theta$, α and k

$$\left| \frac{dX_\theta(\alpha, k)}{d\theta} \right| \leqslant B(X_\theta(\alpha, k) + 1).$$

Condition **A1** guarantees that, for *each* θ, two or more events never occur at the same time. (We still need the generality of (3) because as θ varies the possibility of multiple events must be considered.) It also allows us to work only with strictly positive $X_\theta(\alpha, k)$. Condition **A2** requires that the construction of the parametric family $\{X_\theta(\alpha, k), \theta \in \Theta\}$ be differentiable. If F_α^{-1} is differentiable, then (10) satisfies **A2**. For example, if X_θ is exponential with mean θ, then $F(x, \theta) = 1 - \exp(-x/\theta)$ and this method yields

$$X_\theta = -\theta \ln(1 - U)$$

and

$$\frac{dX_\theta}{d\theta} = -\ln(1 - U) = \frac{X_\theta}{\theta}.$$

Other examples are considered in Suri (1987) and Suri and Zazanis (1988). See also Glynn (1987) for related results.

Finally, **A3** regulates the dependence of the clock samples on the parameter, and is broadly applicable. In particular, it permits location parameters and scale parameters that are bounded away from zero. Condition **A3** will not be needed until Theorem 2.

Notational Convention. Whenever a sample path characteristic appears without a parameter argument, it is understood to be evaluated at a fixed, nominal value of θ; thus, $a_i = a_i(\theta)$ and $Y_i = Y_i(\theta)$. When we need to emphasize a small change in θ we write, for example, $\tau_n(\theta + h)$ and $T_{\theta+h}(\alpha, k)$.

2.1. Event Time Derivatives

We can now turn to expressions for derivatives of performance measures with respect to θ. The first step is to calculate $d\tau_n/d\theta$ for each $n > 0$, and $dT(\alpha, k)/d\theta$ for each $\alpha \in \mathbf{A}$ and $k > 0$. Once the clock samples depend on θ, so do all the sample path characteristics in (2)–(7). Under **A1**, for each θ we may assume events occur singly so that the (finitely many) inequalities that determine $\tau_1, \ldots, \tau_n$ and $a_1, \ldots, a_n$ (via 2–3) are strict. This implies that throughout a sufficiently small neighborhood of θ, these inequalities retain their sense and remain strict. Throughout such a neighborhood, the τ_i change continuously—and, under **A2**, differentiably—in θ; the a_i, and hence the Y_i, remain constant. A potential discontinuity in some τ_i, a_i or Y_i can only occur where the change in θ is large enough to change the argument of minimization in (2) or (3).

Observe, next, that so long as the a_i and Y_i, $i \leqslant n$, remain unchanged, so will the triggering sequence for each τ_i, $i \leqslant n$. Writing $\eta_\theta(r_i; r_j)$ to emphasize the dependence of the triggering indicators on θ, we conclude that, for all sufficiently small h, $\eta_{\theta+h}(r_i; r_j) = \eta_\theta(r_i; r_j)$ for all $i, j \leqslant n$. Using (9), we find that for all sufficiently small h:

$$\tau_n(\theta + h) - \tau_n(\theta) = \sum_{i=1}^{n} (X_{\theta+h}(r_i) - X_\theta(r_i))\eta_\theta(r_n; r_i).$$

Furthermore, if $T_\theta(\alpha, k) < \infty$, then so is $T_{\theta+h}(\alpha, k)$ for all sufficiently small h, and

$$T_{\theta+h}(\alpha, k) - T_\theta(\alpha, k) = \sum_{\beta, j} (X_{\theta+h}(\beta, j) - X_\theta(\beta, j))\eta_\theta(\alpha, k; \beta, j).$$

Combining these observations, we have the following lemma.

Lemma 2. *Suppose that* **C1** *holds. For each θ and n, with probability one, the following hold: a_n and Y_n are constant in a neighborhood of θ; τ_n is differentiable at θ with*

$$\frac{d\tau_n}{d\theta} = \sum_{i=1}^{n} \frac{dX_\theta(r_i)}{d\theta} \eta_\theta(r_n; r_i) \tag{11}$$

and if $T_\theta(\alpha, k) < \infty$, then $T_\theta(\alpha, k)$ is differentiable with

$$\frac{dT_\theta(\alpha, k)}{d\theta} = \sum_{\beta, j} \frac{dX_\theta(\beta, j)}{d\theta} \eta_\theta(\alpha, k; \beta, j). \tag{12}$$

In the example of Figure 1, we see that small increases in departure clocks introduce delays in the occurrence of subsequent departures within the same busy period. Small increases in arrival clocks delay all future arrivals, and all departures during the *next* busy period. In other words, the "perturbations" $dX(r_i)/d\theta$ propagate to the event epochs along the triggering sequences, which is what (11) says.

2.2. Derivatives of Performance Measures

From the derivatives of the state transition epochs we can build up expressions for derivatives of a general class of finite-horizon performance measures. Let f be a *bounded*, real-valued function on the state-space **S** of the GSMP $Z(t, \theta)$. For any (deterministic) real $T > 0$ and any integer $n > 0$ define

$$L_T(\theta) = \int_0^T f(Z(t, \theta)\, dt \tag{13}$$

and

$$L_n(\theta) = \int_0^{\tau_n} f(Z(t, \theta))\, dt. \tag{14}$$

If $T_\theta(\alpha, k) < \infty$, also define

$$L_{\alpha,k}(\theta) = \int_0^{T_\theta(\alpha,k)} f(Z(t, \theta))\, dt. \tag{15}$$

These are the types of performance measures we consider. Through choice of f, many quantities of interest can be obtained. For example, f could count the number of customers at a queue in a closed network, in which case L_T/T is the average queue length on the interval $[0, T]$. By taking f to be the indicator function of the set of busy states for a queue or the state-dependent departure rate, we obtain utilization and throughput measures. More generally, $f(s)$ could represent the rate at which a cost is incurred in state s, in which case the L's are total costs incurred over periods of operation. From the perspective of simulation, L_T, L_n and $L_{\alpha,k}$ differ in how they terminate a run. With L_T, the number of events is random and the time is fixed; with L_n, the number of events is fixed but the time is random; and with $L_{\alpha,k}$, both the number of events and the time are random. *Henceforth, whenever we refer to* $L_{\alpha,k}$ *we assume that* $T_\theta(\alpha, k) < \infty$, *with probability one, for all* $\theta \in \Theta$.

For the following, let $N(t)$ count the number of events in $[0, t]$; that is:

$$N(t) = \sup\{k \geqslant 0 : \tau_k \leqslant t\}.$$

Since **A** is finite, and the clock samples $X(\cdot\,;\,\cdot)$ are greater than zero and independent, $N(t)$ is, with probability one, finite for finite t (use p. 155 of Prabhu 1965, for example).

Lemma 3. *Under* **C1**, **A1** *and* **A2**, *for each* $\theta \in \Theta$, L_T, L_n *and* $L_{\alpha,k}$ *are, with probability one, differentiable at* θ, *with*

$$\frac{dL_T}{d\theta} = \sum_{i=1}^{N(T)} \frac{d\tau_i}{d\theta}\,[f(Y_{i-1}) - f(Y_i)] \tag{16}$$

$$\frac{dL_n}{d\theta} = \sum_{i=0}^{n-1} f(Y_i)\left[\frac{d\tau_{i+1}}{d\theta} - \frac{d\tau_i}{d\theta}\right] \tag{17}$$

and

$$\frac{dL_{\alpha,k}}{d\theta} = \sum_{i=o}^{N(T(\alpha,k))-1} f(Y_i)\left[\frac{d\tau_{i+1}}{d\theta} - \frac{d\tau_i}{d\theta}\right]. \tag{18}$$

Proof. Note, first, that because Z is constant (and equal to Y_i) on $[\tau_i, \tau_{i+1})$, we may write

$$L_T = \sum_{i=1}^{N(T)} f(Y_{i-1})[\tau_i - \tau_{i-1}] + (T - \tau_{N(T)})f(Y_N(T)) \tag{19}$$

$$L_n = \sum_{i=0}^{n-1} f(Y_i)[\tau_{i+1} - \tau_i] \tag{20}$$

and, since $T(\alpha, k) = \tau_{N(T(\alpha,k))}$

$$L_{\alpha,k} = \sum_{i=0}^{N(T(\alpha,k))-1} f(Y_i)[\tau_{i+1} - \tau_i]. \tag{21}$$

Starting with L_n, recall from Lemma 2 that for each θ, with probability one, there exists a neighborhood of θ throughout which $Y_0, \ldots, Y_n$ are constant (and $\tau_1, \ldots, \tau_n$ are differentiable). Thus, with probability one, for sufficiently small h:

$$\begin{aligned} L_n(\theta + h) - L_n(\theta) &= \sum_{i=0}^{n-1} f(Y_i(\theta))\{[\tau_{i+1}(\theta + h) - \tau_i(\theta + h)] - [\tau_{i+1}(\theta) - \tau_i(\theta)]\} \\ &= \sum_{i=0}^{n-1} f(Y_i(\theta))\{[\tau_{i+1}(\theta + h) - \tau_{i+1}(\theta)] - [\tau_i(\theta + \mathrm{h}) - \tau_i(\theta)]\}. \end{aligned}$$

Dividing by h and letting $h \to 0$ we obtain (17). The same argument yields (18) because $N(T(\alpha, k))$ is, with probability one, constant throughout a neighborhood of θ.

Differentiating (19), we get

$$\frac{dL_T}{d\theta} = \sum_{i=1}^{N(T)} f(Y_{i-1})\left[\frac{d\tau_i}{d\theta} - \frac{d\tau_{i-1}}{d\theta}\right] + \left(-\frac{d\tau_{N(T)}}{d\theta}\right) f(Y_N(T))$$

because T is constant, and $N(T)$ (hence, $Y_{N(T)}$) is, with probability one, constant throughout a neighborhood of θ. Rearranging the terms in the sum yields (16).

Remark. The derivatives (16), (17) and (18) admit a simple interpretation: Sufficiently small perturbations in the clock samples that drive Z introduce small changes in the state transition epochs without changing the sequence of states visited (at least over a finite horizon). As θ varies, the sample paths of Z deform by stretching or contracting the state holding times

$\tau_{i+1} - \tau_i$ without changing the basic shape, as given by the Y_i.

2.3. Implementation

The expression (11) for $d\tau_n/d\theta$, while convenient for analysis, is cumbersome to implement because the triggering indicators, η, are defined by working *backward* from τ_n. But it is not necessary to evaluate all η's to calculate $d\tau_n/d\theta$ or to calculate the performance derivatives (16), (17) and (18). We now describe a simpler scheme, which is the way IPA is usually implemented (see, in particular, the algorithm in Suri). We assume that a simulation of Z is available; the recursions (2)–(7) effectively prescribe a simulation algorithm.

With each event $\alpha \in \mathbf{A}$ associate an accumulator A_α. Initialize every A_α to zero. Add to (2)–(7) the following steps:

At τ_{n+1} set
  $k_{n+1} := N(a_{n+1}, n+1)$;
  $A_{a_{n+1}} := A_{a_{n+1}} + dX(a_{n+1}, k_{n+1})/d\theta$;
  for every $\beta \in \mathscr{E}(Y_{n+1}) \backslash \mathscr{E}(Y_n)$
    (i.e., a new clock is set for β at τ_{n+1})
   $A_\beta := A_{a_{n+1}}$.

With this scheme, $d\tau_i/d\theta$ is just the contents of A_{a_i} at τ_i (after updating). The derivatives of the L's are calculated *exactly* from a sample path of Z by substituting A_{a_i} for $d\tau_i/d\theta$ in (16), (17) and (18).

This scheme also shows that it is possible to compute IPA estimates from observation (as opposed to simulation) of Z under one additional condition: namely, that for each $\alpha \in \mathbf{A}$ there exists a function ψ_α such that, for all k

$$\frac{dX_\theta(\alpha, k)}{d\theta} = \psi_\alpha(X_\theta(\alpha, k), \theta).$$

Since every $X_\theta(a_n, k_n)$ can be "observed" from $\{Z(t, \theta), 0 \leq t \leq \tau_n\}$, the existence of ψ_α makes it possible to "observe" $dX_\theta(a_n, k_n)/d\theta$ as well.

Suri and Zazanis (1988) and Suri (1987) give examples where such a ψ_α can be found. Glynn (1987) contains closely related results. These papers show that, often, the inversion representation (10) leads to ψ_α given by

$$\psi_\alpha(x, \theta) = \frac{\partial F_\alpha(x, \theta)/\partial\theta}{\partial F_\alpha(x, \theta)/\partial x}.$$

3. CONTINUITY AND CONSISTENCY

Having shown the existence of, and derived expressions for, the derivatives of a class of performance measures, we can now turn to the key question of whether or not

$$E\left[\frac{dL}{d\theta}\right] = \frac{dE[L]}{d\theta} \tag{22}$$

for any of the L's. This fundamental issue is at the heart of understanding the domain of applicability of infinitesimal perturbation analysis.

It is worth describing why, in practice, (22) may sometimes fail to hold. In Lemmas 2 and 3, the size of the neighborhod throughout which the τ_j's and L's are continuous (and differentiable) depends on the particular realization of the process, i.e., depends on the outcome of the clock samples $X_\theta(\alpha, k)$ and routing indicators $U(\alpha, k)$. But at any θ, for any *fixed* $h > 0$, there may, in general, be a positive probability that some L (and any τ_j) has a discontinuity somewhere in $(\theta, \theta + h)$. Such a discontinuity will typically preclude (22). These potential discontinuities arise when changes in θ introduce changes in the clock samples large enough to change the event that triggers the transition out of a state (in 3).

3.1. Conditions on the Structure of a GSMP

With the above points in mind, we introduce conditions on a GSMP that will guarantee the continuity of the L's even at points where triggering events change. Our conditions ensure this by restricting the possible effect of order changes among events. As functions of θ, the L's may have "kinks" where a change in the parameter changes the order of events, but they will still be continuous—and this is most of what we need for (22). We first state the main condition in provisional form, then give a more general, if less intuitively clear, statement. Recall from Section 1.2 that $\phi(s, \alpha, u)$ is the state reached from s under routing indicator u when event α occurs.

C2′. (Provisional Commuting Condition) *The mapping ϕ can be chosen so that for every $s \in \mathbf{S}$, every α, $\beta \in \mathbf{A}$ and every $u_1, u_2 \in [0, 1]$, if $\{\alpha, \beta\} \subseteq \mathscr{E}(s)$, then*

$$\phi(\phi(s, \alpha, u_1), \beta, u_2) = \phi(\phi(s, \beta, u_2), \alpha, u_1).$$

This condition says that when two events are possible, the state reached through the occurrence of both is independent of their order—provided the same routing indicators are used. (Under our construction of a GSMP, routing indicators are assigned to event types, so that if events of different types change order they keep their routing indicators; cf. 5.) Under this condition, if a parameter perturbation causes a change in the order of two events, the perturbed sample path and the original one will return to the same state after both events have occurred. This limits the possible difference in performance due to parameter changes. Precise arguments are given in Section 3.2. As a simple illustration of **C2′**, we mention an example.

Example 1. (Continued, GI/G/1 Queue) Recall that s records the number in the system, α denotes arrival and β denotes departure. Since the state changes deterministically, the routing indicators actually play no role: an arrival changes s to $s+1$, a departure to $s-1$. Thus, for any $s > 0$, and any u_1, u_2:

$$\begin{aligned}\phi(\phi(s, \alpha, u_1), \beta, u_2) &= \phi(s+1, \beta, u_2)\\ &= s\\ &= \phi(s-1, \alpha, u_1)\\ &= \phi(\phi(s, \beta, u_2), \alpha, u_1).\end{aligned}$$

Condition **C2′** makes no requirement of $s = 0$ because $\mathscr{E}(0)$ contains only one event (α), whereas **C2′** is a condition on states with at least two possible events.

C2′ has the shortcoming that it is stated in terms of ϕ, an object we introduce for our construction, and not solely in terms of the basic GSMP data **S**, **A**, $\mathscr{E}$, p and $\{F_\alpha, \alpha \in \mathbf{A}\}$. Another point is that there may be several equally valid choices of ϕ for the same GSMP, and some may satisfy **C2′** while others do not. (It would be enough to find one ϕ that works.) We take account of these considerations by giving a condition in terms of the transition probabilities p. The proof of the proposition that follows shows how to define ϕ to satisfy **C2′** when the condition on p (plus another minor condition) is satisfied.

C2. (Commuting Condition) *For any s_1, if $\{\alpha, \beta\} \subseteq \mathscr{E}(s_1)$, and s_2 and s_3 satisfy $p(s_2; s_1, \alpha)\, p(s_3; s_2, \beta) > 0$, then there is an s_4 such that*

$$p(s_4; s_1, \beta) = p(s_3; s_2, \beta)$$

and

$$p(s_3; s_4, \alpha) = p(s_2; s_1, \alpha).$$

This condition says that if it is possible to go from s_1 to s_3 through the occurrence of α then β, it must also be possible through the occurrence of β then α, in such a way that each transition triggered by the same event has the same probability. This situation is depicted in Figure 2, where the transitions represented by opposite sides of the square must have the same probability.

To get ϕ from **C2** for arbitrary GSMPs, we will impose one additional condition on p, namely, that

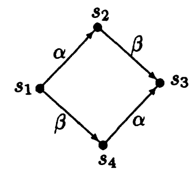

Figure 2. Under the commuting condition **C2**, changing the order of α and β does not change the state reached.

for any s, s', s'' and α if $p(s'; s, \alpha) > 0$, then

$$p(s'; s, \alpha) = p(s''; s, \alpha) \Rightarrow s'' = s'. \tag{23}$$

In words, no two possible transitions from the same state due to the occurrence of the same event can have exactly the same probability. In practice, this is not much of a restriction since the difference could be arbitrarily small. In the queueing example of Section 4.1, we will be able to define ϕ to satisfy **C2′** without recourse to (23).

Proposition 1. *If p satisfies* **C2** *and* (23), *ϕ can be chosen to satisfy* **C2′**.

Proof. The ϕ we choose samples from the mass function $p(\cdot\,; s, \alpha)$ by inverting the corresponding cumulative distribution. After defining this ϕ precisely, we verify the (intuitively clear) fact that it satisfies **C2′**.

Thus, fix a state s and an event α, and let $s_1, s_2, \ldots$ be the states for which $p(s_i; s, \alpha) > 0$, ordered so that for all i, $p(s_{i+1}; s, \alpha) < p(s_i; s, \alpha)$. For $u \in [0, 1]$ let

$$m^* \equiv m^*(u; s, \alpha) = \min\left\{m > 0 : u \leqslant \sum_{k=1}^{m} p(s_k; s, \alpha)\right\} \tag{24}$$

and define $\phi(s, \alpha, u) = s_{m^*}$. Clearly, the set of $u \in [0, 1]$ for which $\phi(s, \alpha, u) = s_k$ is an interval of length $p(s_k; s, \alpha)$, so if U is uniformly distributed on the unit interval:

$$P(\phi(s, \alpha, U) = s') = p(s'; s, \alpha).$$

Therefore, (24) defines a legitimate choice for ϕ. We now show that it satisfies **C2′**.

Suppose that $\{\alpha, \beta\} \subseteq \mathscr{E}(s)$ and $p(s'; s, \alpha) > 0$, then **C2** implies that for every s_i' with $p(s_i'; s', \beta) > 0$ there is some s_i with $p(s_i; s, \beta) = p(s_i'; s', \beta)$. Condition (23) implies that s_i is unique; s_i is completely determined by the value of $p(s_i'; s', \beta)$. Thus, the states reachable from s and s' via β are in one-to-one correspondence, with corresponding states determined by having the same transition probabilities under β. It follows, then, from (24) that

$$\phi(s', \beta, u_2) = s_i' \quad \text{if and only if } \phi(s, \beta, u_2) = s_i.$$

Reversing the roles of α and β, we get the analogous condition for states reachable via α from s and s'' whenever $p(s''; s, \beta) > 0$.

Suppose that $s_2 = \phi(s, \alpha, u_1)$ and $s_3 = \phi(s_2, \beta, u_2)$. Then $s_4 = \phi(s, \beta, u_2)$ and $s_5 = \phi(s_4, \alpha, u_1)$ satisfy $p(s_4; s, \beta) = p(s_3; s_2, \beta)$ and $p(s_5; s_4, \alpha) = p(s_2; s, \alpha)$. Comparing with **C2** and invoking the uniqueness in (23), we find that s_5 must, in fact, be s_3; that is:

$$\phi(\phi(s, \alpha, u_1), \beta, u_2) = \phi(\phi(s, \beta, u_2), \alpha, u_1)$$

which is **C2′**.

3.2. Continuity of the Performance Measures

The most important consequence of these conditions is as follows:

Lemma 4. *Suppose that* **C1** *and* **C2′** *hold (e.g.,* **C2** *and* (23) *hold) and* **A1** *and* **A2** *hold throughout* Θ. *Then*: i), *with probability one, every* τ_i *and finite* $T(\alpha, k)$ *is continuous in* θ *throughout* Θ; ii) *if* Y_i *is discontinuous at some* $\theta_0 \in \Theta$, *then* $\tau_{i+1}(\theta_0) = \tau_i(\theta_0)$; iii) *at a discontinuity of* $N(T)$, $\tau_{N(T)} = T$; iv) *at a discontinuity of* $N(T(\alpha, k))$, $T(\alpha, k) = \tau_{N(T(\alpha,k))-1}$.

Proof. See the Appendix.

Theorem 1. *Under the conditions of Lemma* 4, L_T *and* L_n *are, w.p.*1, *continuous functions of* θ *on* Θ. *The same is true of* $L_{\alpha,k}$ *if* $T(\alpha, k) < \infty$ *w.p.*1 *for all* $\theta \in \Theta$.

Proof. Consider first L_n. Let θ_v be any sequence in Θ converging to θ and consider $L_n(\theta_v) - L_n(\theta)$ *as* $v \to \infty$. Invoking (20), we may treat separately each term

$$f(Y_i(\theta_v))[\tau_{i+1}(\theta_v) - \tau_i(\theta_v)] - f(Y_i(\theta))[\tau_{i+1}(\theta) - \tau_i(\theta)], \quad i < n. \quad (25)$$

If Y_i is continuous at θ, then for all sufficiently large v, $Y_i(\theta_v) = Y_i(\theta)$ since Y_i takes on only discrete values. In this case, continuity of the τ_i implies that (25) goes to zero. Suppose, on the other hand, that Y_i has a discontinuity at θ. From ii of Lemma 4, this implies that $[\tau_{i+1}(\theta) - \tau_i(\theta)] = 0$; hence, part i implies that $[\tau_{i+1}(\theta_v) - \tau_i(\theta_v)] \to 0$ as $v \to \infty$. Boundedness of f now implies that (25) converges to zero as $v \to \infty$.

For L_T, start with (19). If $N(T)$ is continuous at θ, then the continuity of L_T at θ works the same way as that of L_n. Suppose, therefore, that $N(T)$ jumps at θ, say from m to $m - 1$. Suppose that for all sufficiently small $h > 0$, one less event occurs in $[0, T]$ at $\theta + h$ than at θ. Then, writing m for $N(T)$ evaluated at θ:

$$\begin{aligned} &L_T(\theta + h) - L_T(\theta) \\ &= \sum_{i=0}^{m-2} \{f(Y_i(\theta + h)[\tau_{i+1}(\theta + h) - \tau_i(\theta + h)] \\ &\qquad - f(Y_i(\theta))[\tau_{i+1}(\theta) - \tau_i(\theta)]\} \\ &\quad + [T - \tau_{m-1}(\theta + h)]f(Y_{m-1}(\theta + h)) \\ &\quad - [\tau_m(\theta) - \tau_{m-1}(\theta)]f(Y_{m-1}(\theta)) \\ &\quad - [T - \tau_m(\theta)]f(Y_m(\theta)). \end{aligned}$$

Each of the first $m - 1$ terms (the terms inside the summation) goes to zero by the argument used for L_n. For the remaining three terms, part iii of Lemma 4 implies that since $N(T)$ is discontinuous at θ, $\tau_m(\theta) = T$. Substituting $\tau_m(\theta)$ for T, the last three terms simplify to

$$[\tau_m(\theta) - \tau_{m-1}(\theta + h)]f(Y_{m-1}(\theta + h)) - [\tau_m(\theta) - \tau_{m-1}(\theta)]f(Y_{m-1}(\theta)).$$

The argument used for L_n shows that this, too, goes to zero. The case where $N(T)$ jumps *up* at θ works the same way. Finally, for $L_{\alpha,k}$ the argument is essentially the same, except that we must consider a possible jump in $N(T(\alpha, k))$ rather than $N(T)$ and use part iv of Lemma 4.

We can now prove our main result.

Theorem 2. *Consider a GSMP satisfying* **C1**. *Suppose that* **C2′** *holds (e.g.,* **C2** *and* (23) *hold), and* **A1–A3** *are satisfied throughout* Θ.

i. *If* $E[sup_{\theta\in\Theta} N(T)^2] < \infty$, *then* $dE[L_T]/d\theta = E[dL_T/d\theta]$ *on* Θ.
ii. *If* $E[sup_{\theta\in\Theta} \tau_n] < \infty$, *then* $dE[L_n]/d\theta = E[dL_n/d\theta]$ *on* Θ.
iii. *If* $E[sup_{\theta\in\Theta} T(\alpha, k)^2] < \infty$ *and* $E[sup_{\theta\in\Theta} N(T(\alpha, k))^2] < \infty$, *then* $dE[L_{\alpha,k}]/d\theta = E[dL_{l\alpha,k}/d\theta]$ *on* Θ.

Proof. Part of the content of the theorem is that the derivatives of the expectations actually exist. We justify the interchanges by bounding the $dL/d\theta$'s, and then use the dominated convergence theorem with a mean value theorem. To use the mean value theorem, we need to check that the L's are piecewise differentiable functions of θ. Pick any $\theta_0 \in (\theta_a, \theta_b)$ and consider the case of L_n. Since the event epochs are continuous, with probability one there is an $\epsilon > 0$ such that throughout each of the intervals $(\theta_0 - \epsilon, \theta_0)$ and

$(\theta_0, \theta_0 + \epsilon)$ the first n events $a_1, \ldots, a_n$ are unchanged; i.e., no events change order. Throughout each of these intervals, $Y_1, \ldots, Y_n$ are unchanged, and $\tau_1, \ldots, \tau_n$ are differentiable in θ. Hence, from (17) we see that L_n is differentiable throughout these intervals. Even if L_n fails to be differentiable right at θ_0, there are, therefore, intervals immediately to the left and right of θ_0 throughout which L_n is differentiable. Thus, the possible points of nondifferentiability are separated by open intervals, which is another way of saying that L_n is piecewise differentiable. The cases of L_T and $L_{\alpha,k}$ work similarly upon noting that $N(T)$ and $N(T(\alpha, k))$ are unchanged under sufficiently small parameter changes.

We now proceed to bound the $dL/d\theta$'s, beginning with bounds on the $d\tau_i/d\theta$'s. From (11) we find that

$$\left|\frac{d\tau_i}{d\theta}\right| \leq \sum_{j=1}^{i} \left|\frac{dX(r_j)}{d\theta}\right| \eta(r_i; r_j).$$

Applying **A3**, this is

$$\begin{aligned} &\leq \sum_{j=1}^{i} B(X(r_j) + 1)\eta(r_i; r_j) \\ &\leq B(\tau_i + i). \end{aligned} \tag{26}$$

Next, letting $f^* = \sup|f| < \infty$, from (16) we get

$$\begin{aligned} \left|\frac{dL_T}{d\theta}\right| &= \left|\sum_{i=1}^{N(T)} [f(Y_{i-1}) - f(Y_i)] \frac{d\tau_i}{d\theta}\right| \\ &\leq 2f^* \sum_{i=1}^{N(T)} \left|\frac{d\tau_i}{d\theta}\right| \\ &\leq 2f^*B \sum_{i=1}^{N(T)} [\tau_i + i]. \end{aligned}$$

And since $i \leq N(T)$ implies $\tau_i \leq T$, this is

$$\leq 2f^*B \cdot N(T)[T + N(T)]. \tag{27}$$

For L_n we find, similarly, that

$$\left|\frac{dL_n}{d\theta}\right| \leq 2f^*Bn[\tau_n + n] \tag{28}$$

and for $L_{\alpha,k}$ we get

$$\left|\frac{dL_{\alpha,k}}{d\theta}\right| \leq 2f^*B \cdot N(T(\alpha, k))[T(\alpha, k) + N(T(\alpha, k))]. \tag{29}$$

Next, for L_T, L_n and $L_{\alpha,k}$, let D_T, D_n and $D_{\alpha,k}$ be the (random) set of $\theta \in \Theta$ at which each L is differentiable. Since each L is, w.p.1, continuous and piecewise differentiable, we may apply a generalization of the mean value theorem (Dieudonné 1960, p. 160) to conclude that whenever θ and $\theta + h$ are in Θ

$$\left|\frac{L(\theta + h) - L(\theta)}{h}\right| \leq \sup_{\theta \in D} \left|\frac{dL}{d\theta}\right|. \tag{30}$$

Using (27), (28) and (29), it is immediate that under the hypothesis in i:

$$E\left[\sup_{\theta \in D_T} \left|\frac{dL_T}{d\theta}\right|\right] < \infty;$$

under the hypothesis in ii:

$$E\left[\sup_{\theta \in D_n} \left|\frac{dL_n}{d\theta}\right|\right] < \infty;$$

and under the hypothesis in iii:

$$E\left[\sup_{\theta \in D_{\alpha,k}} \left|\frac{dL_{\alpha,k}}{d\theta}\right|\right] < \infty.$$

Applying the dominated convergence theorem to (30), we find that for all $\theta \in \Theta$

$$\lim_{h \to 0} E\left[\frac{L(\theta + h) - L(\theta)}{h}\right] = E\left[\frac{dL}{d\theta}(\theta)\right].$$

That is, the limit on the left exists and is equal to the quantity on the right, which is what we needed to show.

It is useful to have conditions for Theorem 2 stated directly in terms of the clock samples. These are provided by the following result, which is proved in Glasserman (1990) using simple renewal theory arguments.

Corollary 1. *In Theorem* 2, *sufficient conditions for i and ii are, respectively,*

i. for every α, $P(\inf_{\theta \in \Theta} X_\theta(\alpha, k) = 0) < 1$;
ii. for every α, $E[\sup_{\theta \in \Theta} X_\theta(\alpha, k)] < \infty$.

Remark. In subsequent work, Glasserman, Hu and Strickland (1990) show how to extend results on unbiasedness of derivative estimates over finite horizons to *strong consistency* over an infinite horizon. The conditions given here, combined with results in Glasserman et al. validate the use of perturbation analysis in estimating derivatives of *steady-steady* performance measures for certain classes of regenerative systems.

In other work, Glasserman and Yao (1990) show that under conditions **C1** and **C2** it is possible to derive recursive expressions for the event epochs $T(\alpha, k)$ in terms of other event epochs and $X(\alpha, k)$, using only the operations +, min, and max. Using this representation, one could give a different proof of the continuity in Lemma 4i. From this representation it also follows that, under **C1** and **C2**, the event epochs $\{T(\alpha, k), \alpha \in \mathbf{A}, k \geq 1\}$ are monotone increasing functions of the clock samples $\{X(\alpha, k), \alpha \in \mathbf{A}, k \geq 1\}$.

4. AN EXAMPLE AND AN EXTENSION

The GSMP setting is useful for proving results that apply to many different kinds of systems, but it is not the most familiar or intuitive context in which to interpret results. In particular, it is not always immediately obvious what GSMP properties mean for queueing systems. Conversely, such distinctly queueing-related concepts as waiting times are hard to formulate in GSMP notation. To make the results here more vivid, in Glasserman (1990) we give detailed characterizations of queueing systems which do and do not satisfy conditions **C1** and **C2**. In each case—multiclass networks, networks with blocking, etc.—these conditions take on fairly simple interpretations for queues. To give an indication of how **C1** and **C2** are applied, we consider here only the simplest case of a Jackson-like network. In Section 4.2, we comment on extending our results to waiting times.

4.1. Jackson-Like Networks

By a *Jackson-like* network we mean one consisting entirely of first come, first served, infinite buffer, single-server nodes and a single class of customers whose transitions are governed by a Markovian routing matrix P. A Jackson-like network may be open or closed. For simplicity, we suppose that there is a single external arrival stream. Let P_{0i} be the probability that an external arrival joins queue i, and P_{i0} be the probability that a departure from i leaves the network. When $P_{00} = 1$ and every other P_{i0} is zero, the network is closed.

We take the GSMP state of such a network to be its population vector $s = (n_1, \ldots, n_M)$, where M is the number of nodes and n_i is the number of customers at node i. With node i associate the event β_i = departure from node i. Denote external arrivals by β_0. For every state s, β_0 is in the event list $\mathscr{E}(s)$, and $\beta_i \in \mathscr{E}(s)$, $i = 1, \ldots, M$ if and only if $n_i > 0$ in s.

We have seen that the GI/G/1 queue satisfies **C2′** (and **C2** as well). Jackson-like networks generalize this example. Intuitively, the order in which customers move between queues should not change the resulting queue lengths. We now verify this. Let e_i denote the ith standard M-dimensional unit vector. For any s

$$p(s + e_i; s, \beta_0) = P_{0i}.$$

For those s in which $n_i > 0$:

$$p(s - e_i + e_j; s, \beta_i) = P_{ij}$$

and

$$p(s - e_i; s, \beta_i) = P_{i0}$$

and these are the only nonzero transition probabilities. To check **C2** for pairs of events, start with a state s_1 in which there is at least one busy server, say, $n_i > 0$. There are at least two events β_i, β_j in the event list; β_j may be β_0. Consider any k and r for which $P_{ik} > 0$ and $P_{jr} > 0$. Let $s_2 = s_1 - e_i + e_k$ and $s_3 = s_2 - e_j + e_r$, taking e_0 to be the vector of all zeros. Observe that $p(s_2; s_1, \beta_i)\, p(s_3; s_2, \beta_j) = P_{ik}P_{jr} > 0$. If we take $s_4 = s_1 - e_j + e_r$, then **C2** is satisfied because

$$p(s_4; s_1, \beta_j) = P_{jr} = p(s_3; s_2, \beta_j)$$

and

$$p(s_3; s_4, \beta_i) = P_{ik} = p(s_2; s_1, \beta_i).$$

Equation 23 would require that no two nonzero entries of P in the same row have exactly the same value. To avoid this minor restriction, we define ϕ directly to satisfy **C2′**. For each node $i = 0, \ldots, M$, partition the unit interval into $M + 1$ intervals, the kth having length P_{ik}. Take $\phi(s, \beta_i, u) = s - e_i + e_k$ whenever u falls in the kth interval for i (and, of course, $\beta_i \in \mathscr{E}(s)$). In the situation considered before, suppose that u_1 falls in the kth interval for i and u_2 in the rth interval for j. Then

$$\phi(\phi(s_1, \beta_i, u_1), \beta_j, u_2) = s_3 = \phi(\phi(s_1, \beta_j, u_2), \beta_i, u_1).$$

Thus, Jackson-like networks satisfy **C2** and **C2′**. The implications of **C2** for many other kinds of queueing systems—including networks with blocking, with multiple customer classes, and with state-dependent routing—are considered in Glasserman (1990).

4.2. Waiting Times

We now use the example of Jackson-like networks to describe an extension of Theorem 2. This extension has no obvious counterpart for abstract GSMPs, but is important in queueing.

The waiting time (including time in service) of a customer in a queue is the difference between its departure and arrival times. In a GI/G/1 queue that is idle at time zero, the waiting time of the kth customer is $T(\beta, k) - T(\alpha, k)$, if β denotes departure and α denotes arrival. This is the difference between two performance measures $L_{\beta,k}$ and $L_{\alpha,k}$ (with $f \equiv 1$). But in a network, the epoch of the kth arrival to a queue will not, in general, correspond to $T(\alpha, k)$ for any α (though the kth departure will). Consequently, waiting times cannot be expressed in terms of the performance measures we have considered, even by taking differences. We will outline an extension that corrects this shortcoming in the case of Jackson-like networks.

Let $S_{i,k}$ be the epoch of the kth transition that sends a customer to node i; then $T(\beta_i, k) - S_{i,k}$ is the waiting time of the kth customer at node i. We could therefore obtain unbiased estimates of derivatives of expected waiting times if, more generally, derivative estimates for

$$L_{S_{i,k}} = \int_0^{S_{i,k}} f(Z(t))\, dt$$

are unbiased.

By construction, the routing of the kth departure from, say, node j is completely determined by the routing indicator $U(\beta_j, k)$ via ϕ. In particular, the routing decision does not depend on θ (because $U(\beta_j, k)$ does not), and furthermore, *it is independent of the order of events*. Consequently, similar proofs to those of Lemma 4 and Theorem 1 show that $S_{i,k}$ and $L_{S_{i,k}}$ are almost surely continuous in θ, whenever they are finite. Exactly the same proof used for Theorem 2 proves that if, in part iii, $T(\alpha, k)$ and $N(T(\alpha, k))$ are replaced with $S_{i,k}$ and $N(S_{i,k})$, then $dE[L_S]/d\theta = E[dL_S/d\theta]$.

5. RELAXING CONDITION C2: "RELEVANCE"

Our results thus far have not placed any restrictions on which event clock distributions $F_\alpha(x, \theta)$ may depend on θ. Our main condition can be relaxed if we know in advance that only some of the clock distributions depend on θ.

To motivate this idea, consider a single-server queue fed by two classes of arrivals, α_1 and α_2. The GSMP state of the queue must reflect the order of waiting customers of different classes. Consequently, a change in the order of occurrence of α_1 and α_2 changes the resulting state, violating **C2**. Let β_1 and β_2 denote departure of class 1 and class 2 customers. Changes in service times may cause some α_i and β_j to change order, but *cannot* change the order of α_1 and α_2. Departure events are not relevant to the timing of arrivals. Furthermore, a change in the order of an arrival and a departure would not violate **C2**. Thus, if the service time distributions F_{β_1} and F_{β_2} depend on θ, but the interarrival time distributions F_{α_1} and F_{α_2} do not, our previous continuity results should still hold, and it should be possible to use IPA. We now formalize this idea.

Definition 2. For any fixed $\alpha_0 \in \mathbf{A}$, define the set of α_0-*relevant* events recursively as follows:

a. α_0 is α_0-relevant;
b. if α_1 is α_0-relevant and there are states s and s' with $p(s'; s, \alpha_1) > 0$ and $\alpha_2 \in \mathscr{E}(s')\backslash\mathscr{E}(s)$, then α_2 is α_0-relevant.

Part a states that if a clock for α_2 is potentially set by the occurrence of an α_0-relevant event, then α_2 is α_0-relevant. In fact, we have the next lemma.

Lemma 5. *Unless α_1 is α_0-relevant, $\eta(\alpha_1, k_1; \alpha_0, k_0)$ must always be zero for all k_1 and k_0.*

Proof. From Definition 1, if $\eta(\alpha_1, k_1; \alpha_0, k_0) = 1$, then there is a sequence $(\beta_0, j_0), \ldots, (\beta_m, j_m)$ such that $(\beta_0, j_0) = (\alpha_0, k_0)$, $(\beta_m, j_m) = (\alpha_1, k_1)$, and the j_ith clock for β_i is set at the j_{i-1}th occurrence of β_{j-1}, $i = 1, \ldots, m$. Part b on Definition 2 applied repeatedly implies that every β_i (in particular, α_1) is α_0-relevant.

We now state, for fixed $\alpha_0 \in \mathbf{A}$, the next condition.

R. (Relevance Condition) *Condition* **C2′** *holds whenever either α or β is α_0-relevant.*

The multiclass queue with which we began this section satisfies **R** if α_0 is taken to be either β_1 or β_2. The next theorem verifies the applicability of IPA in this case. Other applications of relevance are considered in Glasserman (1990).

Theorem 3. *If only F_{α_0} depends on θ, then the condition of Theorem 2 holds with condition* **C2′** *replaced by* **R**.

Proof. Theorem 2 relies on **C2′** only via the continuity results of Theorem 1, which, in turn, rely on **C2′** only via Lemma 4; thus, we only need to verify that Lemma 4 still holds. The proof of Lemma 4 checks that no discontinuities are introduced when changes in θ cause two events α and β to occur at the same time. If either α or β is α_0-relevant then, by hypothesis, **C2′** still holds for $\{\alpha, \beta\}$ so the original argument remains valid. We now argue that if neither α nor β is α_0-relevant, changes in θ will not cause them to change order. For any k, we have from (12) that

$$\frac{dT_\theta(\alpha, k)}{d\theta} = \sum_{\alpha', j'} \frac{dX_\theta(\alpha', j')}{d\theta}\, \eta(\alpha, k; \alpha', j').$$

Since only α_0-clocks depend on θ, this is

$$\frac{dT_\theta(\alpha, k)}{d\theta} = \sum_j \frac{dX_\theta(\alpha_0, j)}{d\theta}\eta(\alpha, k; \alpha_0, j).$$

But Lemma 5 implies that every $\eta(\alpha, k; \alpha_0, j) = 0$ because α is not α_0-relevant. Thus, for every k, $dT(\alpha, k)/d\theta = 0$, and similarly for every j, $dT(\beta, j)/d\theta = 0$. Therefore, in the proof of Lemma 4, perturbations in θ cannot cause order changes among pairs of irrelevant events; at order changes involving a relevant event, the original argument of Lemma 4 applies.

6. CONCLUDING REMARKS

We have given a general formulation of infinitesimal perturbation analysis derivative estimates for a broad class of discrete-event systems, and verifiable sufficient conditions for these estimates to be unbiased. These conditions should facilitate the application of IPA, especially when translated to queueing systems as in Glasserman (1990).

There is a fundamental flexibility in the derivation of IPA estimators which has not been considered here, and for this reason our conditions cannot be taken to define the "limits" of IPA. There are generally many ways of constructing a parametric family of stochastic processes, and different constructions lead to different IPA algorithms. For instance, the empirical example in Glasserman (1988) shows that changing the way a family of processes is represented, and correspondingly changing the IPA algorithm used, can indeed make IPA work where it initially appears to fail. The construction used here for GSMPs is the most obvious one, and the resulting derivative estimator might well be called "standard" IPA. But other constructions potentially lead to IPA estimators that are unbiased, even when our conditions do not hold. The investigation of alternative constructions and their attendant sample path derivatives is an area of current research.

APPENDIX

Proof of Lemma 4

Part i. From (2)–(5) it is clear that for each i, $\tau_1, \ldots, \tau_i$ are continuous wherever $a_1, \ldots, a_i$ are. Suppose, then, that at some θ some a_j is discontinuous, and let j be, in fact, the smallest index of a discontinuous event. For a_j to be discontinuous, we see from (3) that clocks for two or more events in $\mathscr{E}(Y_{j-1})$ must run out simultaneously; and any of these clocks potentially determines a_j. But if, say, clocks for α and β run out simultaneously, then $c_{j-1}(\alpha) = c_{j-1}(\beta)$. That is, $c_{j-1}(a_j)$ is continuous at θ even if a_j is not. Thus, we see that $\tau_j \equiv \tau_{j-1} + c_{j-1}(a_j)$ is also continuous at θ.

Now suppose, for simplicity, that there are exactly two events α and β in $\mathscr{E}(Y_{j-1})$ whose clocks run out at the same time. (The case where many clocks run out together works the same way; considering only two makes the argument clearer.) If α occurs first then (from (6)) $c_j(\beta) = c_{j-1}(\beta) - c_{j-1}(\alpha) = 0$ and, similarly, if β occurs first, then $c_j(\alpha) = 0$. In the first case, the next event must therefore be β, while in the second case the next must be α. (Recall that under **A1**, no *new* clock is ever set to zero.) In either case, $\tau_{j+1} = \tau_j$. Observe, next, that regardless of the order in which α and β occur, under **C2′**:

$$\begin{aligned} Y_{j+1} &= \phi(\phi(Y_{j-1}, \alpha, U(\alpha, N(\alpha, j-1)+1)), \\ &\qquad \beta, U(\beta, N(\beta, j-1)+1))) \\ &= \phi(\phi(Y_{j-1}, \beta, U(\beta, N(\beta, j-1)+1)), \\ &\qquad \alpha, U(\alpha, N(\alpha, j-1)+1))). \end{aligned}$$

Furthermore, for any $\alpha' \in \mathscr{E}(Y_{j+1})$, either α' was in $\mathscr{E}(Y_{j-1})$ or α' was activated by the occurrence of α or β. Either way, $c_{j+1}(\alpha')$ is independent of the order of α and β: in the first case because (7) applies (which depends only on the number of occurrences of α'); in the second case because (6) applies and $\tau_{j+1} - \tau_j = 0$. For every $\alpha'' \in \mathbf{A}$, $N(\alpha'', j+1)$ is also independent of the order of α and β. But if Y_{j+1}, c_{j+1} and $N(\cdot, j+1)$ are all independent of the order of α and β (at θ), then so is the rest of the sample path, since it is determined by recursion from these quantities. In particular, every $\tau_i(\theta)$ with $i > j$ is independent of the order of α and β (at θ). Thus, if j' is the least index greater than $j+1$ for which $\alpha_{j'}$ is discontinuous at θ, then every τ_i, $i < j'$ is continuous at θ. At j', we may repeat the whole argument and proceed to the next discontinuous event (if any). Thus, we conclude that every τ_i is continuous.

The same argument shows that if $T(\alpha, k)$ is finite, it is continuous. Suppose that $(\alpha, k) = r_j$ so $T(\alpha, k) = \tau_j$ and $a_j = \alpha$. As argued, in order that a_j jump to, say, β, it is necessary that $c_{j-1}(\beta) = c_{j-1}(\alpha)$, in which case, just after a_j becomes β, $c_j(\alpha) = 0$. This implies that α is the next event, a_{j+1}, to occur, and it occurs just after a_j. This makes $(\alpha, k) = r_{j+1}$ and $T(\alpha, k) = \tau_{j+1} = \tau_j$. In short, changing the order of α and β does not change $T(\alpha, k)$.

Part ii. For Y_i to be discontinuous there must be (at least) two events α and β in $\mathscr{E}(Y_{i-1})$ with $c_{i-1}(\alpha) = c_{i-1}(\beta)$. As noted, this implies that $\tau_{i+1} = \tau_i$.

Part iii. Recall that $N(T)$ is defined by

$$N(T) = \sup\{n: \tau_n \leq T\}.$$

Since the τ_n are continuous in θ, $N(T)$ jumps only when a transition occurs right at time T; i.e., when there is an n such that $\tau_n = T$. By definition, then, $\tau_{N(T)} = T$.

Part iv. Since $T(\alpha, k)$ is continuous and every τ_n is continuous, a discontinuity of $N(T(\alpha, k))$ occurs only when two events occur at $T(\alpha, k)$. One of these is the $N(T(\alpha, k))$th event, the other is the $N(T(\alpha, k)) - 1$st. Since both occur at $T(\alpha, k)$, $T(\alpha, k) = \tau_{N(T(\alpha,k))-1}$.

ACKNOWLEDGMENT

This paper is based on part of the author's doctoral dissertation written in the Division of Applied Sciences at Harvard University under the guidance and support of Professor Y. C. Ho. The author and his advisor were supported by the NSF (ECS-85-15449, CDR-85-001-08), the ONR (N00014-84-K-D465, -86-K-0075) and the DAAL (03-86-K-0171). The author is grateful to Professor Ho and to these agencies for their assistance.

REFERENCES

CAO, X. R. 1988. A Sample Performance Function of Closed Jackson Queueing Networks, *Opns. Res.* **36,** 128–136.

DIEUDONNÉ, J. A. 1960. *Foundations of Modern Analysis.* Academic Press, New York.

GLASSERMAN, P. 1988. Infinitesimal Perturbation Analysis of a Birth and Death Process. *Opns. Res. Letts.* **7,** 43–49.

GLASSERMAN, P. 1990. Structural Conditions for Perturbation Analysis of Queueing Systems. *J. Assoc. Comput. Mach.* (to appear).

GLASSERMAN, P., J. Q. HU AND S. G. STRICKLAND. 1990. Strongly Consistent Steady-State Derivative Estimates. *Probability in the Engineering and Informational Sciences* (to appear).

GLASSERMAN, P., AND D. D. YAO. 1990. Generalized Semi-Markov Processes: Anti-Matroid Structure and Second-Order Properties. *Math. Opns. Res.* (to appear).

GLYNN, P. W. 1986. Stochastic Approximation for Monte Carlo Optimization. In *Proceedings of the Winter Simulation Conference*, 350–361.

GLYNN, P. W. 1987. Process Differentiable Representations of Parametric Families of Distributions. Technical Report, Mathematics Research Center, University of Wisconsin, Madison.

GLYNN, P. W., AND D. L. IGLEHART. 1988. Simulation Methods for Queues: An Overview. *Queueing Systems* **3,** 221–256.

HEIDELBERGER, P., X. R. CAO, M. A. ZAZANIS AND R. SURI. 1988. Convergence Properties of Infinitesimal Perturbation Analysis Estimates. *Mgmt. Sci.* **34,** 1281–1302.

HO, Y. C., AND X. R. CAO. 1983. Optimization and Perturbation Analysis of Queueing Networks. *J. Optim. Theory Appl.* **40,** 559–582.

KUSHNER, H. J., AND D. S. CLARK. 1978. *Stochastic Approximation Methods for Constrained and Unconstrained Systems.* Springer, New York.

LI, S., AND Y. C. HO. 1989. Sample Path and Performance Homogeneity of Discrete Event Dynamical Systems. *Automatica* **25,** 907–916.

PRABHU, N. U. 1965. *Stochastic Processes.* Macmillan, New York.

REIMAN, M. I., AND A. WEISS. 1989. Sensitivity Analysis for Simulation Via Likelihood Ratios. *Opns. Res.* **37,** 830–844.

RUBINSTEIN, R. 1989. Sensitivity Analysis and Performance Extrapolation for Computer Simulation Models. *Opns. Res.* **37,** 72–81.

SCHASSBERGER, R. 1976. On the Equilibrium Distribution of a Class of Finite-State Generalized Semi-Markov Processes. *Math. Opns. Res.* **1,** 395–406.

SCHASSBERGER, R. 1978. Insensitivity of Steady-State Distributions of Generalized Semi-Markov Processes With Speeds. *Adv. Appl. Prob.* **10,** 836–851.

SURI, R. 1987. Infinitesimal Perturbation Analysis for General Discrete Event Systems. *J. Assoc. Comput. Mach.* **34,** 686–717.

SURI, R., AND M. A. ZAZANIS. 1988. Perturbation Analysis Gives Strongly Consistent Sensitivity Estimates for the M/G/1 Queue. *Mgmt. Sci.* **34,** 39–64.

WHITT, W. 1980. Continuity of Generalized Semi-Markov Processes. *Math. Opns. Res.* **5,** 494–501.