

GRADIENT ESTIMATION FOR REGENERATIVE PROCESSES

Paul Glasserman

Graduate School of Business
 Columbia University
 New York, NY, U.S.A.

Peter W. Glynn

Department of Operations Research
 Stanford University
 Stanford, CA 94305, U.S.A.

ABSTRACT

This paper is concerned with gradient estimation techniques for steady-state performance measures associated with regenerative stochastic processes. The principal emphasis is on the discussion of conditions under which likelihood ratio methods and infinitesimal perturbation analysis techniques are valid.

1 INTRODUCTION

This paper is concerned with using simulation to estimate gradients of steady-state performance measures in the regenerative process setting. This turns out to be quite a rich class of stochastic processes from an applications viewpoint, encompassing all irreducible recurrent discrete state space Markov processes as well as a variety of more general discrete-event systems. In particular, a large class of generalized semi-Markov processes may be made regenerative by using "splitting" ideas from the theory of Harris recurrent Markov chains; see, for example, Glynn (1989).

Two different gradient estimation algorithms are explored in this paper; likelihood ratio gradient estimations and estimations based on infinitesimal perturbation analysis (IPA). The focus, in this paper, is on the tools needed to rigorously verify the mathematical validity of these techniques in the regenerative setting.

2 STEADY-STATE LIKELIHOOD RATIO GRADIENT ESTIMATION

Let $W = (W_n : n \geq 0)$ be a real-valued stochastic sequence. For each $\theta \in \Lambda = (a, b)$, let P_θ be a probability distribution on the path space of W . We assume that there exists a non-decreasing sequence

$T = T(n) : n \geq 0$ of random times such that $T(0) = 0$ and:

(A1) Under distribution P_θ , W is a non-delayed regenerative sequence with respect to T , for each $\theta \in \Lambda$.

Let $Y_n = \sum_{j=T(n-1)}^{T(n)-1} W_j$, $\tilde{Y}_n = \sum_{j=T(n-1)}^{T(n)-1} |W_j|$, and $\tau_n = T(n) - T(n-1)$. If $E_\theta(\cdot)$ is the expectation operator associated with P_θ , we require that $E_\theta \tilde{Y}_1 < \infty$ and $E_\theta \tau_1 < \infty$, in which case

$$\frac{1}{n} \sum_{j=0}^{n-1} W_j \rightarrow \alpha(\theta) \quad P_\theta \text{ a.s.} \quad (2.1)$$

as $n \rightarrow \infty$, where $\alpha(\theta) \triangleq u(\theta)/\ell(\theta)$ and $u(\theta) = E_\theta Y_1$, $\ell(\theta) = E_\theta \tau_1$.

Relation (2.1) implies that $\alpha(\theta)$ may be interpreted as the steady-state mean of W under P_θ . Our goal is to develop an estimation methodology for $\alpha'(\theta_0)$ for (fixed) $\theta_0 \in \Lambda$. Since $\alpha = u/\ell$, it is clearly sufficient to develop estimators for $u'(\theta_0)$, $\ell'(\theta_0)$, $u(\theta_0)$, and $\ell(\theta_0)$. Of course, $u(\theta_0)$ and $\ell(\theta_0)$ can easily be estimated via sample means formed from i.i.d. copies of Y_1 and τ_1 generated under P_{θ_0} . The greater challenge is to develop estimators for $u'(\theta_0)$ and $\ell'(\theta_0)$.

To accomplish this task, we assume that there exists a σ -field \mathcal{G} for which Y_1 and τ_1 are \mathcal{G} -measurable and such that P_θ is absolutely continuous with respect to P_{θ_0} on \mathcal{G} . Consequently, for each $\theta \in \Lambda$, we may represent $u(\theta)$ and $\ell(\theta)$ as

$$\begin{aligned} u(\theta) &= EY_1 L(\theta) \\ \ell(\theta) &= E\tau_1 L(\theta), \end{aligned}$$

where $E(\cdot) \triangleq E_{\theta_0}(\cdot)$ and $L(\theta)$ is the likelihood ratio (Radon-Nikodym derivative) of P_θ with respect to P_{θ_0} . Suppose that $L(\cdot)$ is P_{θ_0} a.s. differentiable at θ_0 . If

(A2) $(Y(L(\theta_0 + h) - L(\theta_0))/h : h > 0, \theta_0 + h \in \Lambda)$ and $(\tau(L(\theta_0 + L) - L(\theta_0))/h : h > 0, \theta_0 + h \in \Lambda)$ are uniformly integrable under P_{θ_0} ,

then it follows that

$$u'(\theta_0) = EY_1L'(\theta_0)$$

$$\ell'(\theta_0) = E\tau_1L'(\theta_0).$$

Hence, under (A2), $u'(\theta_0)$ and $\ell'(\theta_0)$ can be consistently estimated via sample means formed from i.i.d. copies of $Y_1L'(\theta_0)$ and $\tau_1L'(\theta_0)$ generated under P_{θ_0} . This then solves the problem of estimating $\alpha'(\theta_0)$ consistently.

The key, from a mathematical viewpoint, is therefore to verify (A2).

2.1 Harris Chains

A large class of regenerative systems are derived from discrete-time Markov chains $Z = (Z_n : n \geq 0)$ living on some general state space S . Assume that there exists a subset $A \subseteq S, \lambda > 0$, a positive integer $m \geq 1$, and a probability distribution φ such that:

- (A3) i) $P_\theta[Z_n \in A \text{ infinitely often} | Z_0 = z] = 1$ for $\theta \in \Lambda, z \in S$,
- ii) $P_\theta[Z_m \in dz | Z_0 = x] \geq \lambda\varphi(dz)$ for $x \in A, z \in S, \theta \in \Lambda$.

Condition (A3) guarantees that for Z is a Harris recurrent Markov chain under P_θ for each $\theta \in \Lambda$. It is well known that (A3) can be exploited to obtain regenerative structure for Z . Suppose that Z_0 has distribution φ . Now, observe that the minorization condition (A3) ii) guarantees that for $x \in A$ we can write

$$P_\theta[Z_m \in dz | Z_0 = x]$$

$$= \lambda\varphi(dz) + (1 - \lambda)Q(\theta, x, dz)$$

where $Q(\theta, x, \cdot)$ is a probability distribution on S for each $\theta \in \Lambda$ and $x \in A$. Hence, each time Z visits A , there is a probability $\lambda > 0$ that the chain will be distributed according to φ m time units later, thereby constituting a regeneration. More precisely, we can introduce a sequence $(\eta_n : n \geq 0)$ of "coin flip" r.v.'s such that $\eta_n = 0$ whenever $Z_n \notin A$, and $\eta_n = 1(0)$ whenever $Z_n \in A$ and a "successful" ("unsuccessful") coin flip occurs. Then, the regeneration time τ_1 can be defined as $\tau_1 = \inf\{n \geq m : \eta_{n-m} = 1\}$.

For a given function $f : S \rightarrow \mathbb{R}$, let $W_n = f(X_n)$. Then, it is evident that if $E_\theta\tau_1 < \infty$ and $E_\theta Y_1 < \infty$,

$$\frac{1}{n} \sum_{j=0}^{n-1} W_j \rightarrow \alpha(\theta) \triangleq E_\theta Y_1 / E_\theta \tau_1 \quad P_\theta \text{ a.s.} \quad (2.3)$$

as $n \rightarrow \infty$. To calculate the derivative $\alpha'(\theta_0)$, we have already shown that (A2) is the key.

To obtain a likelihood ratio estimator, we will assume that there exist "densities" $p(\theta, x, y), q(\theta, x, y)$ such that:

- (A4) i) For $x \notin A, y \in S, \theta \in \Lambda, P(\theta, x, dy) = p(\theta, x, y)P(\theta_0, x, dy)$
- ii) For $x \in A, y \in S, \theta \in \Lambda, Q(\theta, x, dy) = q(\theta, x, y)Q(\theta_0, x, dy)$

Assumptions (A3) ii) and (A4) together guarantee that there exist densities $p_n(\theta, x, y)$ such that $P_\theta[Z_n \in dz | Z_0 = x] = p_n(\theta, x, z)P_{\theta_0}[Z_n \in dz | Z_0 \in x]$ for $\theta \in \Lambda, z, z \in S$.

To obtain a likelihood ratio for the path of the sequence $X_n = (Z_n, \eta_n)$ up to time τ_1 , we decompose the path according to appropriately spaced visits to the set A . Specifically, let $S_0 = -m$ and set $S_k = \inf\{n > S_{k-1} + m : Z_n \in A\}$, and let $\beta = \inf\{k \geq 1 : \eta_{S_k} = 1\}$. Then, the likelihood ratio $L(\theta)$ on the σ -field $\mathcal{G} = \sigma(Z_n : 0 \leq n \leq \tau_1)$ of P_θ with respect to P_{θ_0} can be written in the form

$$L(\theta) = \prod_{k=1}^{\beta} L_k(\theta)$$

where

$$L_k(\theta)$$

$$= \prod_{j=S_{k-1}+m}^{S_k+m-1} p(\theta, Z_j, Z_{j+1}) \frac{q(\theta, Z_{S_k}, Z_{S_k+m})}{p_m(\theta, Z_{S_k}, Z_{S_k+m})}$$

for $k < \beta$ and

$$L_\beta(\theta)$$

$$= \prod_{j=S_{\beta-1}+m}^{S_\beta+m-1} p(\theta, Z_j, Z_{j+1}) \frac{1}{p_m(\theta, Z_{S_\beta}, Z_{S_\beta+m})}.$$

Assume that the functions $p_n(\cdot, x, y)$ and $q(\cdot, x, y)$ are continuously differentiable on Λ for $n \geq 1$ and $x, y \in S$. To verify (A2), note that $Y(L(\theta_0 + h) - L(\theta_0))/h = YL'(\xi)$ for some $\xi \in [\theta_0, \theta_0 + h)$. But

$$L'(\theta)$$

$$= L(\theta) \left(\sum_{j=0}^{\tau_1-1} p'(\theta, Z_j, Z_{j+1}) + \sum_{k=1}^{\beta} \ell'_k(\theta) \right)$$

where

$$\ell_k(\theta) = \frac{q(\theta, Z_{S_k}, Z_{S_k+m})}{p_m(\theta, Z_{S_k}, Z_{S_k+m})}$$

for $k < \beta$ and

$$\ell_\beta(\theta) = \frac{1}{p_m(\theta, Z_{S_\beta}, Z_{S_{\beta+m}})}.$$

We have used here the fact that $p_n(\theta, x, y) > 0$ ($q(\theta, x, y) > 0$) whenever $p'_n(\theta, x, y) \neq 0$ ($q'(\theta, x, y) \neq 0$). For $\varepsilon > 0$, set

$$\begin{aligned} \tilde{p}(x, y) &= \sup_{|\theta - \theta_0| \leq \varepsilon} |p(\theta, x, y)| \\ \tilde{p}'(x, y) &= \sup_{|\theta - \theta_0| \leq \varepsilon} |p'(\theta, x, y)| \\ \tilde{\ell}_k &= \sup_{|\theta - \theta_0| \leq \varepsilon} |\ell_k(\theta)| \\ \tilde{\ell}'_k &= \sup_{|\theta - \theta_0| \leq \varepsilon} |\ell'_k(\theta)|, \end{aligned}$$

and observe that by continuity $\tilde{p}(x, y)$ and $\tilde{\ell}_k$ will typically be close to 1. Then,

$$\begin{aligned} |YL'(\xi)| &\leq (\tau_1 \max_{0 \leq k < \tau_1} |f(Z_k)|) \\ &\cdot (\tau_1 \max_{0 \leq k < \tau_1} |\tilde{p}'(Z_k, Z_{k+1})| + \beta \max_{0 \leq k \leq \beta} |\tilde{\ell}'_k|) \\ &\cdot \sup_{|\theta - \theta_0| \leq \varepsilon} L(\theta). \end{aligned}$$

To establish (A2) requires proving that the dominating r.v. just defined is integrable. But its expectation is bounded by

$$\begin{aligned} &\| \sup_{|\theta - \theta_0| \leq \varepsilon} L(\theta) \|_{p_1} \cdot \| \tau_1 \|_{p_2} \cdot \| \max_{0 \leq k < \tau_1} |f(Z_k)| \|_{p_3} \\ &\cdot \| \tau_1 \|_{p_4} \cdot \| \max_{0 \leq k < \tau_1} |\tilde{p}'(Z_k, Z_{k+1})| \|_{p_5} \\ &+ \| \beta \|_{p_4} \cdot \| \max_{0 \leq k \leq \beta} |\ell'_k| \|_{p_5} \end{aligned}$$

when $p_1, p_2, \dots, p_5 > 0$ and $p_1^{-1} + p_2^{-1} + \dots + p_5^{-1} = 1$. But

$$\| \max_{0 \leq k < \tau_1} |f(Z_k)| \|_p \leq E^{1/p} \left(\sum_{j=0}^{\tau_1-1} |f(Z_j)|^p \right);$$

the other maximum terms can be similarly bounded by sums. These expectations, as well as $\| \tau_1 \|_p$, can be bounded by standard Lyapunov function methods; see Chapter 15 of Meyn and Tweedie (1992). (Note that β is geometric so $\| \beta \|_p$ is trivially finite.)

The greater difficulty is posed by the term $\| \sup_{|\theta - \theta_0| < \varepsilon} L(\theta) \|_r$. Let

$$\begin{aligned} a_1(z) &= E[\sup_{|\theta - \theta_0| \leq \varepsilon} L_1(\theta)^r | \beta > 1, Z_0 = z] \\ a_2(z) &= E[\sup_{|\theta - \theta_0| \leq \varepsilon} L_1(\theta)^r | \beta = 1, Z_0 = z]. \end{aligned}$$

By successively conditioning and using the strong Markov property at the times $S_0, S_1, \dots, S_{\beta-1}$, it follows that

$$\begin{aligned} &\| \sup_{|\theta - \theta_0| \leq \varepsilon} L(\theta) \|_r \\ &= E \prod_{j=0}^{\beta-1} q_1(Z_{S_j+m}) q_2(Z_{S_{\beta-1}+m}). \end{aligned}$$

Hence, establishing (A2) requires getting a handle on the functions q_1 and q_2 . In many applications, the set A is compact. Continuity arguments then permit one to control $p_m^{-1}(\theta, Z_{S_k}, Z_{S_{k+m}})$ and $q(\theta, Z_{S_k}, Z_{S_{k+m}})/p_m(\theta, Z_{S_k}, Z_{S_{k+m}})$ are bounded. Furthermore,

$$E \left[\prod_{j=0}^{m-1} \tilde{p}(Z_j, Z_{j+1})^r | Z_0 = z \right]$$

is typically bounded over $z \in A$. Consequently, the key to bounding q_1 and q_2 is to get a handle on

$$h(z) \triangleq E \left[\prod_{j=0}^{S_1-1} \tilde{p}(Z_j, Z_{j+1})^r | Z_0 = z \right]$$

for $z \in A^c$. We complete this discussion of the verification of (A2) by providing a Lyapunov function criterion for bounding h . Let $K(x, dz) = \tilde{p}(x, z)^r P_{\theta_0}[Z_1 \in dz | Z_0 = x]$ and set $g(x) = K(x, A)$. A standard argument establishes that if we can find a non-negative function \tilde{r} satisfying

$$\int_{A^c} K(x, dz) \tilde{r}(z) \leq \tilde{r}(x) - \varepsilon g(x)$$

for $x \in A^c$ and $\varepsilon > 0$, the bound $h(x) \leq \tilde{r}(x)/\varepsilon$ for $x \in A^c$ follows.

2.2 Stochastic Recursions

A large class of Markov chains Z satisfy stochastic recursions of the form

$$Z_{n+1} = \phi(Z_n, U_n)$$

where $\phi : S \times S' \rightarrow S$ and $U = (U_n : n \geq 0)$ is a sequence of S' -valued r.v.'s that is i.i.d. under P_θ for each $\theta \in \Lambda$. In order to guarantee that Z be regenerative, we require the existence of a subset B and a family $\varphi(\theta, \cdot)$ of probability distributions such that:

(A5) i) $P_\theta[Z_n \in B \text{ infinitely often} | Z_0 = z] = 1$ for $z \in S, \theta \in \Lambda$,

ii) $P_\theta[Z_1 \in dz | Z_0 = x] = \varphi(\theta, dz)$ for $x \in B, z \in S$.

Clearly, $\tau_1 = \inf\{n \geq 1 : Z_{n-1} \in B\}$ is a regeneration time for Z . (In fact, Z satisfies (A3) with $\lambda = 1$.) Suppose that:

$$(A6) P_\theta[U_0 \in du] = p(\theta, u)P_{\theta_0}[U_0 \in du]$$

for some density $p(\theta, u)$ that is continuously differentiable on Λ . Then, the likelihood ratio of P_θ with respect to P_{θ_0} on $\mathcal{F} = \sigma(U_n : 0 \leq n < \tau_1)$ is simply given by

$$\tilde{L}(\theta) = \prod_{j=0}^{\tau_1-1} p(\theta, U_j).$$

To verify (A2) for $\tilde{L}(\theta)$ in this setting is much simpler than the verification of the previous section for Harris chains. In particular, since (2.5) continues to hold, the key is to show that

$$E \sup_{|\theta - \theta_0| \leq \varepsilon} \tilde{L}(\theta)^r < \infty \tag{2.6}$$

for some $\varepsilon > 0$. But

$$\begin{aligned} & E \sup_{|\theta - \theta_0| \leq \varepsilon} \tilde{L}(\theta)^r \\ & \leq \sum_{n=0}^{\infty} E^{n/2} \sup_{|\theta - \theta_0| \leq \varepsilon} p(\theta, U_0)^{2r} \cdot P^{1/2}[\tau_1 = n]. \end{aligned}$$

Since $E \sup_{|\theta - \theta_0| \leq \varepsilon} p(\theta, U_0)^{2r}$ can typically be made arbitrarily close to 1 for ε small enough, it is evident that (2.6) holds if $E \exp(\lambda \tau_1) < \infty$ for $|\lambda|$ sufficiently small and positive. However, this can be verified easily by using suitable Lyapunov functions. (In fact, the Lyapunov function r of the previous section can be suitably specialized.) By verifying (2.6), this permits us to establish that $u'(\theta_0)$ and $l'(\theta_0)$ can be estimated via sample means of $Y_1 \tilde{L}'(\theta_0)$ and $\tau_1 \tilde{L}'(\theta_0)$ respectively.

However, if the density $p(\theta, x, y)$ of Z can easily be calculated, the derivative $L'(\theta_0)$ of the previous section is a competing estimator. Note that

$$E[\tilde{L}(\theta)|\mathcal{G}] = \prod_{j=0}^{\tau_1-2} p(\theta, Z_j, Z_{j+1}) \triangleq L(\theta)$$

where $\mathcal{G} = \sigma(Z_n : 0 \leq n < \tau_1)$. If $((\tilde{L}(\theta_0 + h) - \tilde{L}(\theta_0))/h : 0 < h < \varepsilon)$ is P_{θ_0} uniformly integrable, then $(\tilde{L}(\theta_0 + h) - \tilde{L}(\theta_0))/h \rightarrow \tilde{L}'(\theta_0)$ in $L^1(P_{\theta_0})$ so $(L'(P_{\theta_0+h}) - L(\theta_0))/h \rightarrow E[\tilde{L}'(\theta_0)|\mathcal{G}]$ in $L^1(P_{\theta_0})$. It follows that the difference quotients $(L(\theta_0+h) - L(\theta_0))/h$ are uniformly integrable as demanded by (A2) and

$$L'(\theta_0) = E[\tilde{L}'(\theta_0)|\mathcal{G}].$$

By the principle of conditional Monte Carlo, the r.v. $Y_1 L'(\theta_0)$ has smaller variance than $Y_1 \tilde{L}'(\theta_0)$. So, using $L'(\theta_0)$ is statistically desirable.

However, our discussion also shows that for systems satisfying stochastic recursions, the easiest way to establish (A2) for $L(\theta)$ may be to instead establish (A2) for $\tilde{L}(\theta)$. Hence, introducing $\tilde{L}(\theta)$ can be a useful theoretical tool.

3 STEADY-STATE IPA

We now turn to steady-state derivative estimation using IPA, focusing on discrete-time problems. Here, too, regenerative structure plays a key role in the convergence and consistency of the derivative estimates.

We begin with some background on IPA, then give conditions for the derivative estimates to be regenerative, and finally use regenerative structure to prove strong consistency.

3.1 IPA Estimates

Many discrete-time sequences studied through simulation satisfy recursions of the general form

$$W_{n+1} = \phi(W_n, U_n), \quad n \geq 0, \tag{1}$$

where $\{U_n, n \geq 0\}$ are inputs to the simulation and $\{W_n, n \geq 0\}$ are the outputs of interest. We allow the W_n 's to be d -vectors and the U_n 's to be l -vectors. With this generality, (1) is by no means restrictive; but to use IPA, we will need to put further conditions on ϕ .

A familiar example of (1) is the Lindley equation for the waiting times in a single-server queue:

$$W_{n+1} = [W_n + S_n - A_n]^+, \quad n \geq 0, \tag{2}$$

where S_n is the n -th service time and A_n is the time between the n -th and $(n+1)$ -st arrivals. In this example, $d = 1, l = 2$, and $U_n = (A_n, S_n)$; ϕ is defined by (2).

Suppose now that each U_n is a (random) function of a parameter θ ranging over an interval $[a, b]$; then each $W_n, n \geq 1$, depends on θ , and we may also assume that W_0 depends on θ . We want to compute derivatives with respect to this parameter. Let $U_{n,i}$ and $W_{n,i}$ denote the i -th components of U_n and W_n . Formally differentiating (1), we get

$$W'_{n+1,i}(\theta) = \sum_{j=1}^d \frac{\partial \phi_i}{\partial w_j} W'_{n,j}(\theta) + \sum_{j=1}^l \frac{\partial \phi_i}{\partial u_j} U'_{n,j}(\theta), \tag{3}$$

$i = 1, \dots, d$, where the partial derivatives of the i -th component of ϕ are with respect to the indicated components of its arguments and are all evaluated at $(W_n(\theta), U_n(\theta))$. This is another recursion, mapping (W_n, W'_n, U_n, U'_n) to W'_{n+1} . Combined with (1), it gives a mapping from (W_n, W'_n) to (W_{n+1}, W'_{n+1}) with input (U_n, U'_n) . Equation (3), when valid, defines an IPA algorithm for computing $\{W'_n, n \geq 1\}$.

Returning to the Lindley equation (2), we find that the ϕ implicitly defined there is not a differentiable function: differentiability fails where $W_n + S_n - A_n = 0$. For the queue, this event corresponds to one busy period ending at exactly the same time the next one begins. This phenomenon is typical of many discrete-event systems: differentiability may fail when two events occur simultaneously, and these are often the only points of non-differentiability. So, justification of (3) requires some care. Fortunately, this is more of a theoretical than a practical concern; differentiability is often assured by conditions implying that events occur singly with probability one.

We now give a set of conditions from Glasserman (1992ab) justifying (3) and further implying that $E[W'_n(\theta)] = E[W_n(\theta)]'$. These conditions will also be useful in our analysis of steady-state derivative estimation.

Recall that a function $f : S \rightarrow R^d$, $S \subseteq R^n$ is Lipschitz if there exists a constant k_f , called a modulus, such that

$$\|f(x) - f(y)\| \leq k_f \|x - y\|, \quad x, y \in S.$$

Lipschitz functions are differentiable almost everywhere. A Lipschitz function of a scalar is absolutely continuous and is therefore the indefinite integral of its (almost-everywhere defined) derivative. The class of Lipschitz functions is just broad enough to include min, max, and similar functions arising in discrete-event systems, and just smooth enough to be compatible with IPA.

We call a random function $X = \{X(\theta), \theta \in [a, b]\}$ almost-surely Lipschitz if its sample paths are Lipschitz with probability one. If X is Lipschitz, let K_X be a (random) modulus for X .

We now proceed with the conditions. Our first assumption puts minimal smoothness conditions on the inputs and on the initial state:

(A1) W_0 and $\{U_n, n \geq 0\}$ are a.s. Lipschitz functions on $[a, b]$. For each $\theta \in [a, b]$, W_0 and $\{U_n, n \geq 0\}$ are a.s. differentiable at θ , taking one-sided derivatives at the endpoints.

The first part of (A1) restricts dependence on θ for a fixed sample path; the second part fixes θ and varies the sample path. Both types of conditions are needed. We also assume

(A2) ϕ is Lipschitz.

By itself, (A2) implies that the partial derivatives of ϕ exist almost everywhere. But this is not quite enough for (3), since it is possible for $\{(W_n, U_n), n \geq 0\}$ to return infinitely often to the null set of non-differentiable points of ϕ , with positive probability. To rule this out, define

$$C_\phi = \{x \in R^d \times R^l : \phi \text{ is differentiable at } x\},$$

and require

(A3) $P((W_n(\theta), U_n(\theta)) \in C_\phi) = 1$, for all $n \geq 0$, for all $\theta \in [a, b]$.

This is not a primitive condition, in the sense that the distributions of $\{(W_n(\theta), U_n(\theta)), n \geq 0\}$ are generally unknown. Nevertheless, (A3) is often easy to verify in practice. For example, in the Lindley recursion (2), C_ϕ is the complement in $R \times R^2$ of the set $\{(w, s, a) : w + s - a = 0\}$. If, say, $\{S_n, n \geq 0\}$ and $\{A_n, n \geq 0\}$ are i.i.d. and mutually independent, and if either A_0 or S_0 has a density, then $\{(W_n, S_n, A_n), n \geq 0\}$ never leaves C_ϕ , a.s.

We now combine these conditions to validate IPA estimates for (1):

Lemma 3.1. Suppose (A1)-(A3) hold. Then each W_n , $n \geq 0$, is a.s. Lipschitz on $[a, b]$. If W_0 and $\{U_n, n \geq 0\}$ have integrable moduli K_{W_0} and $\{K_{U_n}, n \geq 0\}$, then at every $\theta \in (a, b)$ for which $W_n(\theta)$ is integrable, $E[W_n(\theta)]'$ exists and equals $E[W'_n(\theta)]$.

Proof. The Lipschitz property is preserved by composition, so under (A1) and (A2), W_n is a.s. Lipschitz. Under (A1) and (A3), W_n is also differentiable, a.s., at each θ . Let $k_{\phi,1}$ and $k_{\phi,2}$ be moduli for ϕ as a function of its first and second arguments, respectively, for all values of its other argument. (For example, take $k_{\phi,i} = k_\phi$, $i = 1, 2$, the modulus guaranteed by (A2).) Simple induction shows that

$$\|W'_n(\theta)\| \leq K_{W_n} \equiv k_{\phi,1}^n K_{W_0} + \sum_{j=0}^{n-1} k_{\phi,2}^{n-j} K_{U_j}, \quad (4)$$

and K_{W_n} is an integrable modulus for (each component of) W_n . Then, by dominated convergence,

$$\lim_{h \rightarrow 0} h^{-1} E[W_{n,i}(\theta + h) - W_{n,i}(\theta)]$$

exists and equals $E[W'_{n,i}(\theta)], i = 1, \dots, d$. \square

3.2 Regeneration

Suppose, now, that for each θ ,

$$n^{-1} \sum_{i=0}^{n-1} W_i(\theta) \rightarrow w(\theta), \text{ a.s.}, \tag{5}$$

for some deterministic function $w(\cdot)$. Lemma 3.1 motivates an examination of whether similar conditions imply

$$n^{-1} \sum_{i=0}^{n-1} W'_i(\theta) \rightarrow w'(\theta), \text{ a.s.} \tag{6}$$

There are two considerations in (6) — whether the limit exists and, if it does, whether it equals $w'(\theta)$. Regenerative structure is particularly useful in addressing the first question.

While convergence in (5) is also often based on (possibly implicit) regenerative properties, it turns out that a somewhat stronger regenerative structure is usually needed for the convergence in (6). Indeed, it is possible to have all the sequences $W(\theta) = \{W_n(\theta), n \geq 0\}, \theta \in [a, b]$ regenerate simultaneously infinitely often and yet for $\{W'_n(\theta), n \geq 0\}$ to fail to be regenerative for all θ , as the following example illustrates:

Example 3.1 For $i = 1, 2$, let $\tau^{(i)} = \{\tau_n^{(i)}, n \geq 1\}$ be the points of two independent, unit-rate Poisson processes. The sequence $\{\theta\tau_n^{(2)}, n \geq 0\}$ is the set of points of a rate- $(1/\theta)$ Poisson process if $\theta > 0$. Let $\tau(\theta)$ be the superposition of these points and $\tau^{(1)}$. Let X_n be 1 or 2 depending on whether the n -th point of $\tau(\theta)$ is from $\tau^{(1)}$ or $\theta\tau^{(2)}$. Then

$$\tau'_n(\theta) = \begin{cases} \tau_n(\theta), & X_n = 2; \\ 0, & X_n = 1. \end{cases}$$

Now let $W(\theta)$ be the sequence of spacings, $W_n(\theta) = \tau_n(\theta) - \tau_{n-1}(\theta), n \geq 1$. For each $\theta, W(\theta)$ is an i.i.d. sequence (of exponentially distributed random variables with mean $\theta/(\theta + 1)$). Thus, the sequences $\{W(\theta), \theta > 0\}$ trivially regenerate simultaneously because each regenerates at every n . However, the corresponding derivatives are given by

$$W'_n(\theta) = \begin{cases} W_n(\theta), & X_{n-1} = 2, Z_n = 2; \\ 0, & X_{n-1} = 1, Z_n = 1; \\ \tau_n(\theta), & X_{n-1} = 1, Z_n = 2; \\ -\tau_{n-1}(\theta), & X_{n-1} = 2, Z_n = 1. \end{cases}$$

The first two cases pose no problem, but all four cases occur infinitely often, and the last two show that there can be no N for which the distribution of $W'_N(\theta)$ is independent N .

When $\{W'_n(\theta), n \geq 0\}$ is in fact regenerative, it is often because of the structure present in the following example:

Example 3.2. Consider, again, the Lindley recursion (2). Suppose service requirements and interarrival times are each i.i.d. and mutually independent of each other. The server works at rate $1/\theta$, with $0 < \theta < E[A_1]/E[S_1]$. If the 0-th customer finds the queue empty, then $W_0(\theta) \equiv 0$ and subsequent waiting times obey

$$W_{n+1}(\theta) = [W_n(\theta) + \theta S_n - A_n]^+.$$

It follows that

$$W'_{n+1}(\theta) = \begin{cases} W'_n(\theta) + S_n, & \text{if } W_n(\theta) + \theta S_n > A_n; \\ 0, & \text{otherwise,} \end{cases}$$

with $W'_0(\theta) = 0$. Thus, $\{W'_n(\theta), n \geq 0\}$ returns to zero whenever $\{W_n(\theta), n \geq 0\}$ does; the state $(0, 0)$ is recurrent for the Markov chain $\{(W_n(\theta), W'_n(\theta)), n \geq 0\}$.

The regeneration in Example 3.2 can be explained in rough terms as follows. The process $\{W_n(\theta), n \geq 0\}$ returns to zero infinitely often because many states are mapped to zero by $\phi(\cdot, S_n, A_n)$, for given (S_n, A_n) . In particular, if $\phi(W_n(\theta), S_n, A_n) = 0$, then the same is true throughout a neighborhood of $W_n(\theta)$, a.s. Similarly, $W_{n+1}(\theta)$ remains zero under a sufficiently small change in θ . But if $W_{n+1}(\theta) = 0$ throughout a neighborhood of θ , then $W'_{n+1}(\theta) = 0$. Thus, the fact that the waiting times couple from different initial states forces the derivatives to equal zero infinitely often. In this sense, regeneration at zero has special significance for derivatives.

To formalize these ideas, we return to (3). To write this recursion more compactly, let $D_w \phi$ and $D_u \phi$ be, respectively, $d \times d$ and $d \times l$ matrices of partial derivatives of ϕ with respect to the corresponding arguments. Then (3) simplifies to

$$W'_{n+1}(\theta) = [D_w \phi(W_n(\theta), U_n(\theta))]W'_n(\theta) + [D_u \phi(W_n(\theta), U_n(\theta))]U'_n(\theta). \tag{7}$$

This, in turn, can be re-written as

$$W'_{n+1}(\theta) = X_n(\theta)W'_n(\theta) + Y_n(\theta), \tag{8}$$

where $\{X_n(\theta), n \geq 0\}$ are matrices and $\{Y_n(\theta), n \geq 0\}$ are vectors.

This representation of the IPA estimates is useful in establishing regenerative properties. We first give a result in the setting of Harris ergodic Markov chains, then specialize to classical regeneration. For background on Harris chains, see Asmussen (1987). Since we consider just one value of θ at a time, we suppress the argument.

Theorem 3.2. Suppose $\{(W_n, U_n, U'_n), n \geq 0\}$ is a Harris ergodic Markov chain and let $(\tilde{W}_0, \tilde{U}_0, \tilde{U}'_0)$ have the invariant distribution of this chain. If

$$P(D_w \phi(\tilde{W}_0, \tilde{U}_0) = 0) > 0, \tag{9}$$

then $\{(W_n, W'_n, U_n, U'_n), n \geq 0\}$ is a Harris ergodic Markov chain.

Proof. That $\{(W_n, W'_n, U_n, U'_n), n \geq 0\}$ is Markov follows from (7) and the hypothesis that $\{(W_n, U_n, U'_n), n \geq 0\}$ is Markov. Let $\tilde{X}_n = D_w \phi(\tilde{W}_n, \tilde{U}_n)$ and let $\tilde{Y}_n = D_u \phi(\tilde{W}_n, \tilde{U}_n) U'_n$. By Harris ergodicity, $\{(W_n, U_n, U'_n), n \geq 0\}$ couples with its stationary version at a finite time N_1 , a.s. Subsequently, (X_n, Y_n) coincides with $(\tilde{X}_n, \tilde{Y}_n)$; i.e.,

$$W'_{n+1} = \tilde{X}_n W'_n + \tilde{Y}_n, n \geq N_1.$$

Condition (9) implies that $\tilde{X}_{N_2} = 0$ for some finite $N_2 > N_1$, a.s. Then $W'_{N_2+1} = \tilde{Y}_{N_2}$, regardless of W'_0 ; i.e., the derivatives couple in finite time, for all initial states.

As shown in Glasserman (1992b), condition (9) implies that $\{W'_n, n \geq 0\}$ has a unique stationary distribution, giving a stationary distribution for $\{(W_n, W'_n, U_n, U'_n), n \geq 0\}$. But any Markov chain that admits coupling and has a stationary distribution is Harris ergodic. \square

Similarly, for classical regeneration we have

Theorem 3.3. Suppose, now, that $\{(U_n, U'_n), n \geq 0\}$ are i.i.d., that $\{W_n, n \geq 0\}$ returns to the origin infinitely often, a.s. and that

$$P(D_w \phi(0, U_0) = 0, D_u \phi(0, U_0) U'_0 = 0) > 0. \tag{10}$$

Then $\{(W_n, W'_n), n \geq 0\}$ is regenerative with $W' = 0$ at the regeneration times.

Proof. If $W_n = 0$, then with the strictly positive probability in (10), $W'_{n+1} = 0$ and (W_{n+1}, W'_{n+1}) becomes independent of $\{(W_k, W'_k), k \leq n\}$. \square

With either the Harris ergodicity of Theorem 3.2 or the classical regeneration in Theorem 3.3, we have convergence in distribution of $\{W'_n, n \geq 0\}$ to the stationary distribution \tilde{W}'_0 . If \tilde{W}'_0 is integrable, then

$$n^{-1} \sum_{i=0}^{n-1} W'_i \tag{11}$$

converges almost surely, and this is half of what we need for (6).

3.3 Consistency

Once we have a.s. convergence of time-averages of $\{W'_n(\theta), n \geq 0\}$, the question of consistency reduces to one of interchanging a limit and a derivative. We will give two sets of sufficient conditions for this interchange.

For (6) to hold, we need conditions on the dependence of the inputs across different values of θ . The simplest assumption is

(A4) $\{U_n, n \geq 0\}$ are i.i.d. functions on $[a, b]$.

Naturally, (A4) implies that $\{U'_n, n \geq 0\}$ are also i.i.d. functions. More generally, we could require that $\{(W_n, U_n, U'_n), n \geq 0\}$ be a Harris ergodic function-valued Markov chain.

Our first strong consistency result is based on the method of Hu (1992). Hu shows that limit and derivative can be interchanged under convexity conditions.

Theorem 3.4. Suppose the conditions of Theorem 3.3 hold with (A4) replacing the i.i.d. condition given there. Suppose that (A1) holds with ‘‘Lipschitz’’ replaced by ‘‘convex.’’ Suppose further that ϕ is increasing and convex. If $E[|\tilde{W}_0(\theta)|] < \infty$ for all θ and $E[|\tilde{W}'_0(b)|] < \infty$, then (6) holds at almost every $\theta \in [a, b]$.

Proof. The composition of an increasing convex function with a convex function is convex; thus, every $W_n(\cdot)$ is a.s. convex. A convex function on a closed interval is Lipschitz, hence absolutely continuous, so we have

$$W_n(\theta_2) = W_n(\theta_1) + \int_{\theta_1}^{\theta_2} W'_n(\theta) d\theta, \text{ a.s.,}$$

for all n and θ_1, θ_2 . Now take time averages of both sides and let the time horizon increase to infinity. The result is

$$E[\tilde{W}(\theta_2)] = E[\tilde{W}(\theta_1)] + \lim_{n \rightarrow \infty} \int_{\theta_1}^{\theta_2} n^{-1} \sum_{i=0}^n W'_i(\theta) d\theta.$$

By convexity, $W'_n(\theta) \leq W'_n(b)$, a.s. By dominated convergence, we may therefore interchange limit and integral to get

$$E[\tilde{W}_0(\theta_2)] = E[\tilde{W}_0(\theta_1)] + \int_{\theta_1}^{\theta_2} E[\tilde{W}'_0(\theta)] d\theta,$$

for all θ_1, θ_2 . With $w(\theta) = E[\tilde{W}_0(\theta)]$, this proves that, a.e. on $[a, b]$, $w'(\theta)$ exists and equals $E[\tilde{W}'_0(\theta)]$. \square

The function \max is convex, so Theorem 3.4 can be applied to the Lindley recursion and its generalizations.

Our next result drops the convexity requirement by putting a stronger condition on the Lipschitz property. As in (4), let $k_{\phi,2}$ be a modulus for ϕ as a function of its second argument for any value of its first argument. In the setting of Theorem 3.3, let $\{\tau_k, k \geq 0\}$ be the renewal process of returns to the origin.

Theorem 3.5. Assume the conditions of Theorem 3.3 and (A4). Suppose (A1)–(A3) hold with integrable moduli in (A1). Suppose \tilde{W}_0 and \tilde{W}'_0 are integrable. If $k_{\phi,2} \leq 1$ and $\sup_{\theta} E[\tau_1(\theta) - \tau_0(\theta)] < \infty$, then (6) holds at almost every $\theta \in [a, b]$.

Proof. Consider, for simplicity, the non-delayed case $\tau_0 \equiv 0$. If $\tau_j \leq n < \tau_{j+1}$, then it follows from (4) and the fact that $W'_{\tau_j} = 0$ that

$$\|W'_n(\theta)\| \leq \sum_{i=\tau_j+1}^{\tau_{j+1}} k_{\phi,2}^{|n-i|} K_{U_i}.$$

Hence, if $k_{\phi,2} \leq 1$, taking expectations we get $E[\|W'_n(\theta)\|] \leq E[\tau_1(\theta) - \tau_0(\theta)]E[K_{U_1}]$. By the last hypothesis in the theorem, $E[\|W'_n(\theta)\|]$ is therefore bounded uniformly in n and θ . Arguing much as in Theorem 3.4 (but taking expectations first), this allows us to interchange limit and integral to get

$$\begin{aligned} E[\tilde{W}_0(\theta_2)] - E[\tilde{W}_0(\theta_1)] &= \lim_{n \rightarrow \infty} \int_{\theta_1}^{\theta_2} n^{-1} \sum_{i=0}^n E[W'_i(\theta)] d\theta \\ &= \int_{\theta_1}^{\theta_2} E[\tilde{W}'_0(\theta)] d\theta, \end{aligned}$$

since the limit of (11) is also the limit of its expectation. This implies that $w(\theta) \equiv E[\tilde{W}_0(\theta)]$ is differentiable at almost every θ , with $w'(\theta) = E[\tilde{W}'_0(\theta)]$. \square

3.4 Remarks on Continuous Time

Though we have only considered consistency of IPA for discrete-time processes, similar techniques are useful in continuous time. We briefly outline how.

Let $X = \{X_t, t \geq 0\}$ have a countable state space and suppose X changes state at times $\{\tau_n, n \geq 0\}$. Suppose X depends on θ . Assuming, say,

$$t^{-1} \int_0^t f(X_s(\theta)) ds \rightarrow m(\theta), \tag{12}$$

the question is whether

$$\frac{d}{d\theta} [t_n^{-1} \int_0^{t_n} f(X_s(\theta)) ds] \rightarrow m'(\theta), \tag{13}$$

for some $\{t_n, n \geq 0\}$ increasing to infinity.

A first step in showing (13) is arguing that

$$\frac{d}{d\theta} \int_0^{\tau_n} f(X_s(\theta)) ds = \sum_{i=0}^{n-1} f(Y_n(\theta)) [\tau'_{i+1} - \tau'_i],$$

where Y_n is the state just after the n -th transition. Techniques from Section 3.1 are useful here because this finite-horizon derivative estimator will typically be unbiased only if the state-transition times are Lipschitz functions of θ .

For many discrete-event systems, it is possible to supplement the system state with the time remaining for scheduled events to obtain an augmented process $\{(Y_n, C_n), n \geq 0\}$ that is Markov. Regeneration of this process is useful in establishing (12). To analyze IPA estimators, it is convenient to consider a further augmented process $\{(Y_n, C_n, \Delta_n), n \geq 0\}$, where Δ records information about derivatives of scheduled event times. One way for $\{(Y_n, C_n), n \geq 0\}$ to regenerate is for Y to visit a state in which an entirely new set of events is scheduled; often, this corresponds to a system returning to an empty state. When new events are scheduled, new derivatives are generated, so $\{(Y_n, C_n, \Delta_n), n \geq 0\}$ also regenerates, and this is an important step in verifying (13). As in Section 3.2, we see here a connection between a strong form of regeneration for the original process and regeneration at zero for the derivatives. A detailed treatment of the continuous-time setting is given in Glasserman et al. (1991).

ACKNOWLEDGEMENTS:

The work of the first author was supported by the National Science Foundation under Contract No.

ECS-9216490. The second author was supported by Army Research Office Contract No. DAAL03-91-G-0101 and by the National Science Foundation under Contract No. DDM-9101580.

REFERENCES

- Asmussen, S. 1987. *Applied Probability and Queues*, New York: Wiley.
- Glasserman, P. 1992a. Regenerative Derivatives of Regenerative Sequences. *Adv. Appl. Prob.*, to appear.
- Glasserman, P. 1992b. Stationary Waiting Time Derivatives. *QUESTA*, to appear.
- Glasserman, P., J.Q. Hu, and S.G. Strickland. 1991. Strongly Consistent Steady-State Derivative Estimates. *Prob. Inf. Eng. Sci.* 5:391-413.
- Glynn, P.W. (1989). A GSMP Formalism for Discrete Event Systems. *Proc. IEEE* 77:1423.
- Hu, J.Q. 1992. Convexity of Sample Path Performances and Strong Consistency of Infinitesimal Perturbation Analysis. *IEEE Trans. Automatic Control* 37:258-262.
- Meyn, S.P., and R.L. Tweedie. 1992. *Markov Chains and Stochastic Stability*, Springer-Verlag, New York.

AUTHOR BIOGRAPHIES

PAUL GLASSERMAN is an Associate Professor in the Management Science division of the Columbia Business School. In 1988-90, he was a member of Technical Staff in the Operations Research Department of AT&T Bell Laboratories. His simulation-related interests include variance reduction and gradient estimation.

PETER W. GLYNN is currently an Associate Professor in the Department of Operations Research at Stanford University. His research interests include stochastic systems, computational probability, and simulation.