MANAGEMENT SCIENCE
Vol. 35, No. 11, November 1989
Printed in U.S.A.

SIMULATING DISCOUNTED COSTS*

BENNETT L. FOX AND PETER W. GLYNN
Department of Mathematics, University of Colorado, Denver, Colorado 80204-5300
Department of Operations Research, Stanford University, Stanford, California 94305

We numerically estimate, via simulation, the expected infinite-horizon discounted cost d of running a stochastic system. A naive strategy estimates a finite-horizon approximation to d. We propose alternatives. All are ranked with respect to *asymptotic* variance as a function of computer-time budget and discount rate, when semi-Markov and/or regenerative structure or neither is assumed. In this setting, the naive truncation estimator loses; it may triumph, however, when the computer-time budget is modest, the discount rate is large, and the process simulated is not regenerative or has long cycle lengths.
(DISCOUNTED COSTS; SIMULATION; SEMI-MARKOV PROCESS; REGENERATIVE PROCESS)

1. Introduction

In many settings, discounted costs arise naturally. This paper describes simulation methodologies for estimation of expected discounted costs associated with systems that exhibit stochastic fluctuations. Such techniques are important for numerical computation of discounted costs for stochastic processes in which conventional numerical methods either fail to apply or are inefficient. Examples of such processes include non-Markov processes or infinite state space Markov chains. The discussion given here of simulation algorithms for the discounted cost problem also merits interest to the extent that it provides an excellent vehicle for illustrating several sophisticated "variance reduction" methods for stochastic simulation. These techniques are more accurately called efficiency improvement techniques, as we shall see.

To formulate the estimation problem mathematically, we let $X = \{X(t) : t \geq 0\}$ be a stochastic process taking values in a state space S. Suppose that f, g are two real-valued functions defined on S, in which $f(x)$ represents the cost of running X in state x and $g(x)$ corresponds to the (positive) discount rate in state x. Then

$$D = \int_0^\infty \exp(-V(s))f(X(s))ds \tag{1}$$

is the infinite-horizon discounted cost, where $V(s) = \int_0^s g(X(u))du$. Our goal in this paper is to construct Monte Carlo simulation algorithms for numerically evaluating $d = ED$.

We now describe the layout of the rest of this paper. §2 develops a naive estimator for d based on truncation of the infinite-horizon integral, and studies its relevant theory. In §3, an estimator based on randomizing the truncation point is developed, and it is shown that for large computational budgets, this estimator beats the naive truncation estimator of §2. §4 shows how to exploit semi-Markov process structure to improve the efficiency of the randomized estimator of §3 by "conditioning out" the holding times. In §5, an estimator which makes use of regenerative structure is explored, whereas §6 studies an estimator which utilizes both semi-Markov and regenerative structure to obtain efficiency.

* Accepted by George S. Fishman, former Departmental Editor; received May 13, 1987. This paper has been with the authors 8 months for 2 revisions.

0025-1909/89/3511/1297$01.25

In §7, we compare the asymptotic variances of the above five estimators when the discount rate is small; a small discount rate is natural in many economic settings. §8 analytically calculates the mean square error of the relevant estimators for the $M/D/\infty/\infty$ queue, for an arbitrary discount rate. Finally, in §9, we offer some concluding remarks. When confidence intervals are based on (routinely-constructed) consistent estimators of the variance constants that we give, possibly large computer-time budgets may be needed for the asymptotic theory to "kick in" and thus assure that the desired coverage is achieved; this caveat is not a criticism of our paper, because it applies universally. Unless otherwise indicated, all proofs are deferred to the Appendix to make the paper easier to read.

2. A Truncation Algorithm

Naive Monte Carlo simulation, based directly on (1), is impractical since generating the r.v. D generally requires an infinite amount of computation time. A straightforward alternative truncates the r.v. D at some finite time horizon β, yielding the quantity

$$D(\beta) = \int_0^\beta \exp(-V(s))f(X(s))ds.$$

Given a computational budget t, it is clear that the truncation point β should increase with t. (Observe that a sample mean estimator based on $D(\beta)$ with β fixed converges to $ED(\beta)$, which is in general not equal to d.)

As a consequence, we need to define a sampling plan $\{\beta(t) : t > 0\}$, in which $\beta(t)$ corresponds to the truncation point associated with computer budget t. Assuming that the time required to generate a replicate of $D(\beta)$ is $c_1\beta$ ($c_1 > 0$), we find that the number $n(t)$ of runs completed with budget t is $\lfloor t/c_1\beta(t)\rfloor$. Given t, our estimator for d will then be

$$\delta_1(t) = \begin{cases} \dfrac{1}{n(t)}(D_1(\beta(t)) + D_2(\beta(t)) + \cdots + D_{n(t)}(\beta(t))), & n(t) \geq 1, \\ 0, & n(t) = 0, \end{cases} \tag{2}$$

where $\{D_n(\cdot) : n \geq 1\}$ is a sequence of i.i.d. replicates of $D(\cdot)$.

We now investigate the choice of sampling plan which optimizes the behavior of the estimator $\delta_1(t)$ for a fixed computer budget t. Given the exponential character of discounting, it seems reasonable to expect that the bias of $\delta_1(t)$ behaves like $a\exp(-c\beta(t))$ for some constants a, c. This expectation can be justified when X is a regenerative stochastic process; see Glynn and Whitt (1988). In fact, when the discount rate is constant (i.e. $g \equiv \alpha$), $c = \alpha$. In any case, it then follows that the mean square error (MSE) of $\delta_1(t)$ is given approximately by

$$E(\delta_1(t) - d)^2 \approx c_1(\mathrm{var}\, D)\frac{\beta(t)}{t} + a^2\exp(-2c\beta(t)).$$

The choice of $\beta(t)$ which minimizes the above MSE expression is

$$\beta^*(t) = \lambda_0^* + \lambda_1^* \log t \tag{3}$$

where $\lambda_0^* = -\log(c_1 \mathrm{var}\, D/2a^2c)/2c$, $\lambda_1^* = 1/2c$.

Theorem 1 below shows that the above approximations can be justified rigorously; its proof relies heavily on the general theory of replication estimators of the form (2), as described in Fox and Glynn (1989). The following (reasonable) assumptions will be needed:

H1. f, g are strictly positive functions on S.

H2. $0 < \mathrm{var}\, D < \infty$

H3. $b(\beta) \equiv d - ED(\beta) \sim ae^{-c\beta}$ as $\beta \to \infty$ for some constants a, c.

We require that f be strictly positive merely to simplify the technical statements of the theorems presented in this paper. It is not necessary and can be replaced by suitable (cumbersome) absolute integrability hypotheses on f.

THEOREM 1. *Assume H1–H3 and suppose* β^* *is defined by* (3). *Then*:
(i) *If* $\beta(t) = \beta^*(t)$, *then*

$$\sqrt{\frac{t}{\log t}}(\delta_1(t) - d) \Rightarrow \left(\frac{c_1 \operatorname{var} D}{2c}\right)^{1/2} N(0, 1).$$

(ii) *If* $\beta(t)/\beta^*(t) \to \kappa > 1$ *as* $t \to \infty$, *then*

$$\sqrt{\frac{t}{\log t}}(\delta_1(t) - d) \Rightarrow \left(\frac{\kappa c_1 \operatorname{var} D}{2c}\right)^{1/2} N(0, 1).$$

(iii) *If* $\beta(t)/\beta^*(t) \to \kappa < 1$ *as* $t \to \infty$, *then*

$$\sqrt{\frac{t}{\log t}}\,|\delta_1(t) - d| \Rightarrow \infty.$$

Part (iii) forces the selection $\beta(t) \geq \beta^*(t)$ for large enough t. On the other hand, if the constant κ appearing in (ii) is strictly greater than one, the variance of the limiting normal r.v. is greater than that obtained when $\beta = \beta^*$. We conclude that Theorem 1 shows that the asymptotically optimal choice of sampling plan is $\beta = \beta^*$.

Implementing this choice requires determining a and c. [A glance at the proof of Theorem 1 shows that in fact c is the crucial parameter, in the sense that if $\beta(t) = \log t/2c + \xi$, then convergence result (i) always ensues, regardless of the choice of ξ.] Theorem 1 indicates that if one is to guess a choice for c, it is better to underestimate c than overestimate it. In particular, suppose that one uses $\beta(t) = \log t/2c' + \xi$ with $c' < c$. Then $\delta_1(t)$ will satisfy relation (ii) with $\kappa = c/c'$; on the other hand, if $c' > c$, $\delta_1(t)$ has the poor convergence structure associated with (iii). An underestimate c' of c is always available when $\Lambda = \inf\{g(x) : x \in S\} > 0$, namely $c' = \Lambda$.

Theorem 1 also shows that even if β is chosen optimally, the best possible rate of convergence is $\sqrt{\log t/t}$. This is unsatisfactory in comparison to the canonical rate of $1/\sqrt{t}$ typical of Monte Carlo simulation. Thus, the straightforward truncation approach of this section appears inefficient for large computational budgets t, and the investigation of alternative algorithms is warranted. Heuristic adjustments to $\beta^*(t)$ may be appropriate when the computer-time budget is only moderate.

EXAMPLE. Let $X = (X(t) : t \geq 0)$ represent the number of customers in a $G/G/\infty/\infty$ queueing system. Assume that $f(x) = x$ and that $g(x) = \alpha$. Then $c = \alpha$ and the above analysis suggests that with computer budget t, we should simulate $\lfloor 2\alpha t/\log t \rfloor$ i.i.d. replicates of the r.v. $D(\log t/2\alpha)$, thereby yielding the estimator

$$\delta_1(t) = \frac{1}{\lfloor 2\alpha t/\log t \rfloor} \sum_{i=1}^{\lfloor 2\alpha t/\log t \rfloor} \int_0^{(\log t/2\alpha)} e^{-\alpha s} X_i(s)\, ds.$$

3. A Randomized Estimator

A principal difficulty in estimating d is that the naive Monte Carlo estimator based on replicates of D is inadmissible: it requires infinite time to simulate a single observation of D. Hence, it is clear that one must carefully consider the computational effort required per observation in order to properly assess the efficiency of an estimator. Hammersley and Handscomb (1964) proposed evaluating the efficiency of a Monte Carlo procedure via the formula

$$\text{Efficiency} = (\text{Time})^{-1} \cdot (\text{Variance})^{-1} \qquad \text{where}$$

Time = expected computation time per observation,
Variance = variance per observation.
Glynn and Whitt (1986) rigorously justify this criterion, even when an observation and the work to get it are correlated. Thus, the efficiency of an estimator may be improved by reducing computation time per observation and/or reducing variance. An important implication of this observation is that the efficiency of an estimator may be improved by increasing the variance per observation provided that the computational time required per observation is appropriately decreased. The estimator proposed in this section has precisely this property. Specifically, the variance per observation is greater than that of D, but the observations can be generated in finite time.

Suppose R is an exponential r.v. with mean one, which is independent of X. Set

$$\tilde{D}(1) = \int_0^{V^{-1}(R)} f(X(t))dt$$

where $V^{-1}(\cdot) = \inf\{t > 0 : V(t) \geq \cdot\}$; we call $\tilde{D}(1)$ a randomized estimator since it involves adding additional randomness to the probability space. Note that $\tilde{D}(1)$ requires simulating X only up to time $V^{-1}(R)$, and can therefore be generated in finite time if:

H4. $V(\infty) \equiv \int_0^\infty g(X(t))dt = \infty$ a.s.

When $g \equiv \alpha$, then $EV^{-1}(R) = 1/\alpha$ and var $V^{-1}(R) = 1/\alpha^2$. Thus, both the expected work per run and the variance per run are (generally significantly) affected by α.

The following proposition shows that efficient estimation of d can be based on $\tilde{D}(1)$, since $E\tilde{D}(1) = d$.

PROPOSITION 1. *Assume H1, H2, H4. Then*

$$D = E\{\tilde{D}(1)|X\}, \qquad \text{so that} \qquad E\tilde{D}(1) = d.$$

Standard properties of conditional expectation guarantee that var $D \leq$ var $\tilde{D}(1)$, so that the variance per observation is increased by using an estimator based on $\tilde{D}(1)$. To analyze the efficiency of $\tilde{D}(1)$, we will obtain a central limit theorem (CLT) for the corresponding estimator. Let $\{(\tilde{D}_n(1), V_n^{-1}(R_n)) : n \geq 1\}$ be a sequence of i.i.d. copies of $(\tilde{D}(1), V^{-1}(R))$. Given t units of computation time, the number of observations generated is

$$N_1(t) = \max\{n \geq 0 : c_1(V_1^{-1}(R_1) + \cdots + V_n^{-1}(R_n)) \leq t\}$$

(disregarding overhead for generating the R_i's) and the estimator $\delta_2(t)$ available after t units of computational effort have been expended is

$$\delta_2(t) = \begin{cases} \frac{1}{N_1(t)}(\tilde{D}_1(1) + \cdots + \tilde{D}_{N_1(t)}(1)), & N_1(t) \geq 1, \\ 0, & N_1(t) = 0. \end{cases} \tag{4}$$

EXAMPLE (continued). Let $R_1\ R_2, \cdots$ be i.i.d. exponential r.v.'s having unit mean and let $N_1(t) = \max\{n \geq 0 : R_1 + \cdots + R_n \leq \alpha t/c\}$ be the corresponding counting process. (In fact, $N_1(t)$ is a Poisson r.v. with mean $\alpha t/c_1$.) Thus, if $X_1, X_2, \cdots$ are i.i.d. replicates of X, the estimator $\delta_2(t)$ takes the form

$$\delta_2(t) = \frac{1}{N_1(t)} \sum_{i=1}^{N_1(t)} \int_0^{R_i/\alpha} X_i(s)ds$$

when $N_1(t) \geq 1$.

Theorem 2 shows that $\delta_2(t)$ converges at rate $t^{-1/2}$; it can also be used, in a straightforward way, to obtain confidence intervals for d.

THEOREM 2. *Assume H*1, *H*4, *and* $E\tilde{D}(1)^2 < \infty$. *Then*

$$t^{1/2}(\delta_2(t) - d) \Rightarrow (c_1 \operatorname{var} \tilde{D}(1) \cdot EV^{-1}(R))^{1/2} N(0, 1).$$

Furthermore, $\operatorname{var} \tilde{D}(1) = 2 \int_0^\infty \int_0^t E\{\exp(-V(t)) f(X(s)) f(X(t))\} ds dt - d^2$ *and* $EV^{-1}(R) = \int_0^\infty E \exp(-V(t)) dt$.

This theorem confirms the efficiency criterion specified by (3), in the sense that the asymptotic variance of the limiting normal r.v. is precisely the reciprocal of the efficiency given by (3). In the next three sections, we will describe estimation algorithms that will increase the efficiency of $\delta_2(t)$ by reducing the variance of $\tilde{D}(1)$ without increasing the average amount of time required to generate an observation; the improved efficiency will be obtained by utilizing special stochastic structure in X.

4. Discrete-Time Conversion for Semi-Markov Processes

In this section, we construct an efficient estimator for d which exploits semi-Markov process (SMP) type structure; the idea is to eliminate some of the variance in $\tilde{D}(1)$ by conditioning on the embedded discrete-time process which describes the sequence of states visited by X. This "discrete-time conversion" is similar in spirit to the estimator discussed in Fox and Glynn (1986) for estimation of steady-state quantities associated with SMP's. It "undoes" some of the variance increase due to randomization.

Specifically, we assume in this section the existence of a discrete-time process $Y = \{Y_n : n \geq 0\}$ taking values in S and a strictly increasing sequence of random times $\{S_n : n \geq 0\}$ such that

H5. (i) $X(t) = \sum_{n=0}^{\infty} Y_n I(S_n \leq t < S_{n+1})$ where $S_0 = 0$.

(ii) $\{\beta_n : n \geq 0\}$ is conditionally independent given Y, where $\beta_n = S_{n+1} - S_n$.

(iii) $P\{\beta_n \in dt \mid Y\} = F(Y_n, Y_{n+1}, dt)$ for some family of distributions F indexed by $S \times S$.

H5 generalizes the notion of SMP, since we do not require here that Y be a Markov chain. (See, e.g., Çinlar (1975) for further discussion of SMP's.)

To apply discrete-time conversion, we let $N(t) = \max\{n \geq 0 : S_n \leq t\}$ be the number of transitions of X by time t, and set $M = N(V^{-1}(R))$. From H5, we get

$$\begin{aligned}\tilde{D}(1) &= \sum_{j=0}^{M-1} \int_{S_j}^{S_{j+1}} f(X(t)) dt + \int_{S_M}^{V^{-1}(R)} f(X(t)) dt \\ &= \sum_{j=0}^{M-1} f(Y_j)\beta_j + f(Y_M)(V^{-1}(R) - S_M).\end{aligned}$$

The "discrete-time" estimator of this section is based on $\tilde{D}(2) = E\{\tilde{D}(1) \mid Y, M\}$, that is

$$\tilde{D}(2) = \sum_{j=0}^{M-1} f(Y_j) E\{\beta_j \mid Y, M\} + f(Y_M) E\{V^{-1}(R) - S_M \mid Y, M\}. \tag{5}$$

Let $\varphi(x, y, \lambda)$ be the Laplace transform of the distribution $F(x, y, dt)$ defined by

$$\varphi(x, y, \lambda) = \int_{[0,\infty)} e^{-\lambda t} F(x, y, dt).$$

It is easy to show, using a dominated convergence argument, that the derivative $\varphi'(x, y, \lambda)$ with respect to λ exists for all positive arguments λ. Let $\{(\varphi_k, \varphi_k') : k \geq 0$ be the sequence defined by

$$\varphi_k = \varphi(Y_k, Y_{k+1}, g(Y_k)), \qquad \varphi_k' = \varphi'(Y_k, Y_{k+1}, g(Y_k)).$$

With this notation in hand, the next proposition calculates the conditional expectations appearing in (5), as well as the conditional distribution of M given Y.

PROPOSITION 2. *Assume H1, H4, and* $E\tilde{D}(1) < \infty$. *Then, for* $k < m$

$$P\{M \geq m \mid Y\} = \prod_{j=0}^{m-1} \varphi_j,$$

$$E\{\beta_k \mid Y, M = m\} = -\varphi_k'/\varphi_k,$$

$$E\{V^{-1}(R) - S_M \mid Y, M\} = \frac{(1 + g(Y_M)\varphi_M' - \varphi_M)}{g(Y_M)(1 - \varphi_M)}.$$

As a consequence of Proposition 2, *we find that*

$$\tilde{D}(2) = -\sum_{j=0}^{M-1} f(Y_j)\varphi_j'/\varphi_j + f(Y_M)\frac{(1 + g(Y_M)\varphi_M' - \varphi_M)}{g(Y_M)(1 - \varphi_M)}. \tag{6}$$

Formula (6) shows that we get $\tilde{D}(2)$ by generating Y up to time M, where M is generated by using the conditional distribution given in Proposition 2. The following algorithm can be used to produce r.v.'s with the distribution of $\tilde{D}(2)$; its validity follows immediately from (6), noting that M is generated by "inversion".

Algorithm A. 1. Generate a random variate U, uniform on $(0, 1)$.
2. Generate Y_0.
3. Set $m \leftarrow 0$, $\Lambda \leftarrow 1$, $\Gamma \leftarrow 0$.
   Comment: now $\Lambda = P\{M \geq 0 \mid Y\}$.
4. Generate Y_{m+1}.
5. Set $\Lambda \leftarrow \Lambda\varphi_m$.
   Comment: now $\Lambda = P\{M \geq m + 1 \mid Y\}$.
6. If $U > \Lambda$, then
   (i) Set $D \leftarrow f(Y_m)\dfrac{(1 + g(Y_m)\varphi_m' - \varphi_m)}{g(Y_m)(1 - \varphi_m)} - \Gamma$.
   Comment: now $M = m$.
   (ii) exit.
7. Else,
   (i) set $\Gamma \leftarrow \Gamma + f(Y_m)\varphi_m'/\varphi_m$.
   (ii) set $m \leftarrow m + 1$.
   (iii) go to step 4.

An estimator $\delta_3(t)$ based on a sequence $\{(\tilde{D}_n(2), M_n) : n \geq 1\}$ of i.i.d. replicates of $\tilde{D}(2)$ can be constructed analogously to $\delta_2(t)$ (see (4)). The estimator $\delta_3(t)$ so defined is a sample mean of $N_2(t)$ observations of $\tilde{D}(2)$, where $N_2(t)$ is the number of observations generated in t units of computer time. To a first approximation, $N_2(t) = \max\{n \geq 0 : c_2(M_1 + \cdots + M_n) \leq t\}$ where c_2 is the computer time required to increment m by one in Algorithm A. (This disregards the set-up time to generate M, and the fact that the effort required to execute steps 4 to 7 of Algorithm A depends on the random states occupied at times m and $m + 1$.)

EXAMPLE (continued). We illustrate this estimator by specializing the $G/G/\infty/\infty$ queue so that the interarrival and service time distributions are exponential with parameters λ and μ, respectively. The resulting $M/M/\infty/\infty$ queue is a semi-Markov process. (In fact, it is a continuous-time Markov chain.) In this case,

$$\varphi(x, y, \tau) = \frac{\lambda + x\mu}{(\lambda + \tau + x\mu)}$$

for $x, y \in R^+ = \{0, 1, 2, \cdots\}$ and $\tau \geq 0$. If $g(x) = \alpha$ for $x \geq 0$ and $f(x) = x$, the ratio φ_j'/φ_j takes the form $\varphi_j'\varphi_j = -(\lambda + \alpha + Y_j\mu)^{-1}$. Also, it turns out that

$$\frac{(1 + g(Y_M)\varphi_M' - \varphi_M)}{g(Y_M)(1 - \varphi_M)} = \frac{1}{\lambda + \alpha + Y_M\mu}.$$

Letting $(Y(1), M_1), (Y(2), M_2), \cdots$ be i.i.d. replicates of (Y, M), our estimator (based on t units of computational effort) can be expressed as

$$\delta_3(t) = \frac{1}{N_2(t)} \sum_{i=1}^{N_2(t)} \sum_{j=0}^{M_i} \frac{Y_j(i)}{(\lambda + \alpha + Y_j(i)\mu)},$$

where the r.v. M_i has conditional distribution

$$P\{M_i \geq m \mid Y(i)\} = \prod_{j=0}^{m-1} \frac{(\lambda + Y_j(i)\mu)}{(\lambda + \alpha + Y_j(i)\mu)}.$$

Note that this estimator requires only simulating the embedded chain Y (the exponential holding times needed to simulate X are unnecessary). The discrete-time Markov chain Y has transition matrix

$$P\{Y_{n+1} = j \mid Y_n = i\} = \begin{cases} \dfrac{\lambda}{\lambda + n\mu}, & j = i + 1, \\ \dfrac{n\mu}{\lambda + n\mu}, & j = i - 1, \quad i \geq 1, \\ 0, & i = j \quad \text{or} \quad |i - j| \geq 2. \end{cases}$$

The following CLT describes the behavior of the estimator $\delta_3(t)$, and can be used to construct confidence intervals for d.

THEOREM 3. *Assume H1, H4, and* $E\tilde{D}(2)^2 < \infty$. *Then*

$$t^{1/2}(\delta_3(t) - d) \Rightarrow (c_2 EM \cdot \text{var}\, \tilde{D}(2))^{1/2} \cdot N(0, 1)$$

as $t \to \infty$.

The proof of this result follows immediately from §5 of Glynn and Whitt (1986). Since $\tilde{D}(2) = E\{\tilde{D}(1) \mid Y, M\}$, it follows by the principle of conditional Monte Carlo (see Bratley, Fox, and Schrage 1987, §2) that var $\tilde{D}(2) \leq$ var $\tilde{D}(1)$. Thus, the estimator $\delta_3(t)$ is obtained from $\delta_2(t)$ by reducing the variance per observation. However, as Theorems 2 and 3 point out, an efficiency increase is obtained only if $(c_2 EM)/(c_1 EV^{-1}(R)) \leq$ var $\tilde{D}(1)/$var $\tilde{D}(2)$.

To fully understand this condition, note that var $\tilde{D}(1)/$var $\tilde{D}(2)$ reflects the degree to which randomness in $\tilde{D}(1)$ is due to the holding times β_j, as opposed to the embedded sequence Y. On the other hand, the ratio $c_2 EM/c_1 EV^{-1}(R)$ describes the complexity of generating a $\tilde{D}(2)$ observation relative to a $\tilde{D}(1)$ variate. Observe that both types of observations require generating Y up to time M; the difference is that $\tilde{D}(1)$ additionally requires generating the holding times β_j, while $\tilde{D}(2)$ involves the Laplace transform quantities φ_j and φ_j'. If the $F(x, y, dt)$'s are distributions having Laplace transforms that are easily numerically evaluated (as is the case with gamma r.v.'s, for example), then the (possible) increase in effort involved in passing from $\tilde{D}(1)$ to $\tilde{D}(2)$ should be modest; in these circumstances, $\delta_3(t)$ is more efficient than $\delta_2(t)$. For a more detailed comparison of "discrete-time" estimators with their "continuous-time" analogs, see §2 of Fox and Glynn (1986).

5. Estimation for Regenerative Processes

We assume now that X is a (possibly) delayed regenerative process with regeneration times $0 \le T_0 < T_1 < \cdots$ (if X is nondelayed, set $T_0 = 0$.); see Çinlar (1975) for a definition and examples. Thus, we do not require in this section that X satisfy the semi-Markov hypothesis H5. Let $T_{-1} = 0$. The independence of regenerative cycles implies that

$$d = EA(0) + EC(0)EK(0) \qquad \text{where} \tag{7}$$

$$A(i) = \int_0^{\tau_i} \exp\left(-\int_0^t g(X(T_{i-1}+s))ds\right) f(X(T_{i-1}+t))dt,$$

$$C(i) = \exp\left(-\int_0^{\tau_i} g(X(T_{i-1}+t))dt\right),$$

$$K(i) = \int_0^{\infty} \exp\left(-\int_0^t g(X(T_{i-1}+s))ds\right) f(X(T_i+t))dt,$$

and $\tau_i = T_i - T_{i-1}$. A similar analysis of $EK(0)$ shows that

$$EK(0) = EA(1) + EC(1)EK(1).$$

But $K(1)$ has the same distribution as $K(0)$ by the regenerative property, so $EK(0) = EK(1)$. We conclude that $EK(0) = EA(1)\cdot(1 - EC(1))^{-1}$. Substituting into (7) yields

$$d = EA(0) + EC(0)EA(1)\cdot(1 - EC(1))^{-1}. \tag{8}$$

Equation (8) suggests that d can be estimated by simulating regenerative cycles. Since each regenerative cycle can be generated in finite time, independently of g, we will avoid the problems inherent in trying to generate D explicitly, or, when the discount rate is small, in randomizing as in §§3 and 4. (See also §7.) In the discounting context, it is important to allow the possibility that X is a delayed regenerative process (as opposed to steady-state simulation). For example, if one is asked to compute the discounted cost for a Markov chain initiated with a distribution concentrated on more than one point, this generalization would be required.

Since (8) involves two different types of cycles (delayed and nondelayed), it offers the possibility to stratify the computation effort so as to maximize the efficiency of the resulting estimators. Given a computational budget t, we allocate a proportion p to generating pairs $(C(0), A(0))$ and a proportion $q = 1 - p$ to simulating the pairs $(C(1), A(1))$ from the nondelayed cycle. An estimator $\delta_4(t)$ is then obtained by substituting the resulting sample means in (8).

To be precise, let $\{(C_n(i), A_n(i), \tau_{ni}) : n \ge 1\}$ $(i = 0, 1)$ be two independent sequences of i.i.d. random vectors where $(C_n(i), A_n(i))$ shares the same distribution as $(C(i), A(i))$, and where τ_{ni} represents the length of the corresponding cycle used to obtain $(C_n(i), A_n(i))$. Thus, if we set $p_0 = p$, $p_1 = q$, then $N^i(t) = \max\{n \ge 0 : c_1(\tau_{1i} + \cdots + \tau_{ni}) \le p_i t\}$ is the number of type i cycles completed by time t. Put

$$(\bar{C}_t(i), \bar{A}_t(i)) = \begin{cases} \dfrac{1}{N^i(t)}((C_1(i), A_1(i)) + \cdots + (C_{N^i(t)}(i), A_{N^i(t)}(i))), & N^i(t) \ge 1, \\ 0, & N^i(t) = 0. \end{cases}$$

Then the estimator $\delta_4(t)$ is given by

$$\delta_4(t) = \bar{A}_t(0) + \bar{C}_t(0)\bar{A}_t(1)\cdot(1 - \bar{C}_t(1))^{-1}.$$

EXAMPLE (continued). For our $M/M/\infty/\infty$ example, suppose that $X(0) = 0$ and that we use the state 0 as our regeneration state. Then X is a nondelayed regenerative process. If we set $T(0) = 0$ and $T(n) = \tau_{11} + \cdots + \tau_{n1}$ for $n \geq 1$, $T(n)$ is the instant at which the process X enters 0 for the nth time. The estimator $\delta_4(t)$ can then be written as

$$\delta_4(t) = \frac{\sum_{i=1}^{N^1(t)} \int_0^{\tau_{1i}} e^{-\alpha s} X(T(i-1)+s)ds/N^1(t)}{1 - (\sum_{i=1}^{N^1(t)} e^{-\alpha\tau_{1i}}/N^1(t))}.$$

To analyze the behavior of this estimator, we derive a CLT for $\delta_4(t)$. (Again, this can also be used to produce confidence intervals for d.) We require that:
H6. $E\tau_i < \infty$ $(i = 0, 1)$.

THEOREM 4. *Assume H1, H2, and H6. Then, for* $0 < p < 1$,

$$t^{1/2}(\delta_4(t) - d) \Rightarrow (\sigma_0^2/p + \sigma_1^2/q)^{1/2} N(0, 1)$$

as $t \to \infty$, *where*

$$\sigma_0^2 = c_1 \operatorname{var}(A(0) + C(0) \cdot EK(1)) \cdot E\tau_0,$$

$$\sigma_1^2 = c_1 \left(\frac{EC(0)}{1 - EC(1)}\right)^2 \operatorname{var}(A(1) + C(1) \cdot EK(1)) \cdot E\tau_1.$$

To optimize the performance of $\delta_4(t)$, we select p to minimize the asymptotic variance term $\delta_0^2/p + \sigma_1^2/q$. It is easily verified that the minimizer is given by $p^* = \sigma_0(\sigma_0 + \sigma_1)^{-1}$ (provided $\sigma_0 + \sigma_1 > 0$) where $\sigma_i = \sqrt{\sigma_i^2}$, in which case the corresponding variance is $(\sigma_0 + \sigma_1)^2$. To compare the efficiency of the estimator with the previous ones, in particular the deterministic truncation estimator, it is useful to relate the coefficients defining δ_0^2 and σ_1^2 to var D appearing in Theorem 1.

PROPOSITION 3. *Assume H1 and H2. Then*

$$\operatorname{var} D = \operatorname{var}(A(0) + C(0) \cdot EK(1)) + \frac{EC(0)^2}{1 - EC(1)^2} \operatorname{var}(A(1) + C(1) \cdot EK(1))).$$

To aid in comparison, note that $EC(0)^2 \geq (EC(0))^2$ and $E(1 - C(1)^2) \geq (1 - EC(1))^2$. (For the second inequality, $0 \leq C(1) \leq 1$ so $EC(1) \geq EC(1)^2$. Hence $1 - EC(1) \leq 1 - EC(1)^2$. But since $0 \leq EC(1) \leq 1$, $(1 - EC(1))^2 \leq 1 - EC(1)$.) In the nondelayed case where $C(0) = 1$, $A(0) = 0$, we choose $q = 1$ (obviously). Theorems 1 and 4 then suggest that

$$\operatorname{var} \delta_1(t)/\operatorname{var}\delta_4(t) \sim \frac{\log t \cdot (1 - EC(1))^2}{2bE\tau_1 \cdot (1 - EC(1)^2)}. \tag{9}$$

We conclude that if $t \gg \exp(2bE\tau_1)$, it is better to use $\delta_4(t)$.

6. Estimation for Regenerative Semi-Markov Processes

In this section, we illustrate how the methods of §4 and 5 can be combined to obtain an estimator $\delta_5(t)$ which exhibits the best features of $\delta_3(t)$ and $\delta_4(t)$. In particular, $\delta_5(t)$ exploits the regenerative structure of X while "filtering out" the variance in $\delta_4(t)$ due to holding time randomness; the latter property is achieved by using discrete-time conversion.

Returning to the set-up of §3, we now assume that the embedded sequence Y is regenerative. Thus, we require that Y possess regeneration times $0 \leq U_0 < U_1 < \cdots$ and set $U_{-1} = 0$, $\eta_i = U_i - U_{i-1}$. By the conditional independence of the β_j's given Y, it follows that the random times $T_i = S_{U_i}$ are regeneration times for X. Hence, (8) is valid for the T_i's; as an immediate consequence, we obtain the identity

$$d = E\tilde{A}(0) + E\tilde{C}(0)E\tilde{A}(1)\cdot(1 - E\tilde{C}(1))^{-1} \tag{10}$$

where $\tilde{A}(i) = E\{A(i)\,|\,Y\}$, $\tilde{C}(i) = E\{C(i)\,|\,Y\}$. To compute the conditional expectations appearing in (10), observe that

$$E\{C(i)\,|\,Y\} = E\left\{\prod_{j=U_{i-1}}^{U_i-1} \exp(-g(Y_j)\beta_j)\,|\,Y\right\} = \prod_{j=U_{i-1}}^{U_i-1} \varphi_j \qquad \text{and}$$

$$\begin{aligned} E\{A(i)\,|\,Y\} &= E\left\{\sum_{j=U_{i-1}}^{U_i-1} \int_{S_j}^{S_{j+1}} \exp\left(-\int_0^t g(X(T_{i-1}+s))ds\right) f(Y_j)\,|\,Y\right\} \\ &= E\left\{\sum_{j=U_{i-1}}^{U_i-1} \prod_{k=0}^{j-1} \exp(-g(Y_k)\beta_k) \int_0^{\beta_j} \exp(-g(Y_j)t)dt f(Y_j)\,|\,Y\right\} \\ &= \sum_{j=U_{i-1}}^{U_i-1} \prod_{k=0}^{j-1} \varphi_k(1-\varphi_j)\frac{f(Y_j)}{g(Y_j)}\,. \end{aligned}$$

Given the above formulas, it is straightforward to generate the pairs $(\tilde{C}(i), \tilde{A}(i))$ by simulating the sequence Y. As in §4, the computational effort may be assigned so that a fraction p_i of the total time t is delegated to generation of pairs $(\tilde{C}(i), \tilde{A}(i))$ $(i = 0, 1)$. An estimator $\delta_5(t)$ can then be constructed analogously to $\delta_4(t)$.

We can derive a CLT for $\delta_5(t)$ which describes its convergence and can be used for confidence interval estimation; the proof is analogous to that of Theorem 4 and is therefore omitted. The result (Theorem 5 below) assumes that the computational effort required to generate $(\tilde{C}(i), \tilde{A}(i))$ is $c_3\eta_i$. The constant c_3 reflects the difficulty of simulating the chain Y and numerically evaluating the φ_j's. We do not assume that $c_2 = c_3$ since the discrete-time algorithm of §4 also involves numerical evaluation of the derivatives of the φ_j's, which may be harder. For continuous-time Markov chains, however, both φ_j and φ_j' have simple closed forms.

H7. $E\eta_i < \infty$ $(i = 0, 1)$.

THEOREM 5. *Assume H1, H2, and H7. Then, for* $0 < p < 1$,

$$t^{1/2}(\delta_5(t) - d) \Rightarrow (\tilde{\sigma}_0^2/p + \tilde{\sigma}_1^2/q)^{1/2}N(0, 1)$$

as $t \to \infty$, *where*

$$\tilde{\sigma}_0^2 = c_3\cdot \operatorname{var}(E\{A(0) + C(0)\cdot EK(1)\,|\,Y\})\cdot E\eta_0,$$

$$\tilde{\sigma}_1^2 = c_3\left(\frac{EC(0)}{1 - EC(1)}\right)^2 \operatorname{var}(E\{A(1) + C(1)\cdot EK(1)\,|\,Y\})\cdot E\eta_1.$$

EXAMPLE (continued). We illustrate this estimator with our $M/M/\infty/\infty$ example. If $X(0) = 0$, then the embedded Markov chain Y (see §4) has initial state 0. If we choose 0 as our regeneration state, $U_0 = 0$ and U_i is the transition on which Y visits 0 for the ith time. Let $N_5(t) = \max\{n \ge 0 : c_3U_n \le t\}$ be the process which counts the number of regenerative cycles completed with t units of computational effort. Noting that Y is a nondelayed regenerative sequence when $Y_0 = 0$, the estimator $\delta_5(t)$ takes the form

$$\delta_5(t) = \frac{\sum_{i=1}^{N_5(t)} \sum_{j=U_{i-1}}^{U_i-1} (1-\varphi_j)Y_j \prod_{k=0}^{j-1} \varphi_k / N_5(t)}{1 - (\sum_{i=1}^{N_5(t)} \prod_{j=U_{i-1}}^{U_i-1} \varphi_j / N_5(t))}$$

where $\varphi_j = (\lambda + Y_j\mu)/(\lambda + \alpha + Y_j\mu)$.

The principle of conditional Monte Carlo again guarantees that $\tilde{\sigma}_0^2 \le \sigma_0^2$, $\tilde{\sigma}_1^2 \le \sigma_1^2$; the amount of variance reduction depends on the extent to which the randomness of D is due to the holding times. As an immediate consequence, we find that if $c_1 E\tau_1 \approx c_3 E\eta_1$, $\delta_5(t)$ is more efficient than $\delta_4(t)$. The following proposition relates $\tilde{\sigma}_0^2$ and $\tilde{\sigma}_1^2$ to var $(E\{D|Y\})$; its proof is similar to Proposition 3 and is omitted.

PROPOSITION 4. *Assume H1 and H2. Then*

$$\operatorname{var}(E\{D|Y\}) = \operatorname{var}(\tilde{A}(0) + \tilde{C}(0)\cdot EK(1)) + \frac{E\tilde{C}(0)^2}{1 - E\tilde{C}(1)^2} \operatorname{var}(\tilde{A}(1) + \tilde{C}(1)\cdot EK(1)).$$

This result can be used to compare the efficiency of $\delta_1(t)$ to $\delta_5(t)$ when Y is nondelayed. By arguing as in (9), we find that

$$\frac{\operatorname{var}\delta_1(t)}{\operatorname{var}\delta_5(t)} \sim \frac{\log t\cdot c_1 \operatorname{var} D\cdot(1 - EC(1))^2}{2bE\eta_1\cdot c_s\cdot \operatorname{var}(E\{D|Y\})\cdot(1 - E\tilde{C}(1)^2)}.$$

Here, we find that if $t \gg \exp(2bc_3 E\eta_1 \operatorname{var}(E\{D|Y\})/(\operatorname{var} D\cdot c_1))$, $\delta_5(t)$ is more efficient than $\delta_1(t)$.

7. Analysis of Efficiency for Small Discount Rates

In this section, we study the relative efficiencies for small discount rates of the methods considered above. Smallish discount rates arise naturally in many economic contexts (e.g. low inflation rate settings) and a considerable literature has developed on this topic. (See, for example, Veinott (1969) or Whitt (1972).)

To make our analysis precise, let

$$V_\alpha(t) = \alpha \int_0^t g(X(s))ds$$

where g does not depend on α and set $d(\alpha) = ED_\alpha$, where

$$D_\alpha = \int_0^\infty \exp(-V_\alpha(t))f(X(t))dt.$$

We are interested in the efficiency of our five estimators for $D(\alpha)$ when α is small. Given Theorems 1 through 5, we examine the asymptotic behavior of the scaling constants appearing in front of the limiting normal r.v. These scaling constants determine the width of the confidence interval associated with a given method, and consequently one wishes to choose estimators for which the scaling constants are as small as possible.

Our subsequent mathematical analysis requires:

H8. X is a (possibly) delayed regenerative process with regeneration times $0 \le T_0 < T_1 < \cdots$.

H9. $E(Y_i(f)^4 + Y_i(g)^4) < \infty$ $(i = 0, 1)$, where

$$Y_i(f) = \int_{T_{i-1}}^{T_i} f(X(s))dx, \qquad Y_i(g) = \int_{T_{i-1}}^{T_{i-1}} g(X(s))dx.$$

Although the results stated here require the regenerative structure for the proofs, it seems likely that the same asymptotic behavior holds for more general classes of processes. This belief is supported by some of the more general limit theorems appearing in Glynn and Whitt (1988).

To state the following theorem, we add an α-dependence to all the r.v.'s and constants appearing in Theorems 1 to 5. For example, $\tilde{D}_\alpha(1)$ is defined as $\int_0^{V_\alpha^{-1}(R)} f(X(t))dt$.

THEOREM 6. *Assume H1–H8. Then*:
(a) $d(\alpha) \sim (1/\alpha)(r(f)/r(g))$.
(b) var $D_\alpha \sim (1/2\alpha)(\sigma^2/r(g))$.
(c) $c(\alpha) \sim \alpha r(g)$.
(d) var $\tilde{D}_\alpha(1) \sim (1/\alpha^2)(r^2(f)/r^2(g))$.
(e) $EV_\alpha^{-1}(R) \sim (\alpha r(g))^{-1}$.
(f) var $\tilde{D}_\alpha(2) \sim (1/\alpha^2)(r^2(f)/r^2(g))$.
(g) $EM(\alpha) \sim (\alpha r(g))^{-1}(E\eta_1/E\tau_1)$.
(h) $(\sigma_0(\alpha) + \sigma_1(\alpha))^2 \sim (c_2/\alpha^2)(\sigma^2/r(g))$.
(i) $(\tilde{\sigma}_0(\alpha) + \tilde{\sigma}_1(\alpha))^2 \sim (c_3/\alpha^2)(\tilde{\sigma}^2/r(g))(E\eta_1/E\tau_1)$.
as $\alpha \downarrow 0$, *where* $r(f) = EY_1(f)/E\tau_1$, $\tau(g) = EY_1(g)/E\tau_1$,

$$\sigma^2 = \frac{\text{var}\,(r(g)Y_1(f) - r(f)Y_1(g))}{E\tau_1}\cdot\frac{1}{r(g)^3},$$

$$\tilde{\sigma}^2 = \frac{\text{var}\,(E\{r(g)Y_1(f) - r(f)Y_1(g)\,|\,Y\})}{E\tau_1}\cdot\frac{1}{r(g)^3}.$$

Given a computation budget of (at least) moderate size t, the above theorem tells us that if the discount rate is small, then we can expect that

$$\text{var}\,\delta_1(t) \approx \frac{\log t}{t}\cdot\frac{c_1}{\alpha^2}\cdot\frac{\sigma^2}{4r(g)}, \tag{11a}$$

$$\text{var}\,\delta_2(t) \approx \frac{1}{t}\cdot\frac{c_1}{\alpha^3}\cdot\frac{r^2(f)}{r^3(g)}, \tag{11b}$$

$$\text{var}\,\delta_3(t) \approx \frac{1}{t}\cdot\frac{c_2}{\alpha^3}\cdot\frac{r^2(f)}{r^3(g)}\cdot\frac{E\eta_1}{E\tau_1}, \tag{11c}$$

$$\text{var}\,\delta_5(t) \approx \frac{1}{t}\cdot\frac{c_1}{\alpha^2}\cdot\frac{\sigma^2}{r(g)}, \tag{11d}$$

$$\text{var}\,\delta_5(t) \approx \frac{1}{t}\cdot\frac{c_3}{\alpha^2}\cdot\frac{\tilde{\sigma}^2}{r(g)}\cdot\frac{E\eta_1}{E\tau_1}. \tag{11e}$$

Assuming that $c_jE\eta_1 \le c_1E\tau_1$ for $j = 2, 3$ (i.e. the cost of simulating a regenerative cycle in discrete time is less than or equal to the cost of simulating a cycle in continuous time), the above analysis suggests that we can order (for small discount rates) the estimators in order of decreasing preference as follows: $\delta_5(t)$, $\delta_4(t)$, $\delta_1(t)$, $\delta_3(t)$, $\delta_2(t)$.

The above results also show that the discounting problem does *not* get harder as $\alpha \downarrow 0$, provided that we take advantage of regenerative structure. Suppose that we wish to construct a $100(1-\gamma)\%$ confidence interval for $d(\alpha)$ with half-width equal to $\epsilon\%$ of $d(\alpha)$. If the estimator $\delta_5(t)$ is used, the computational effort $t(\alpha)$ required for this relative width confidence interval is given approximately by

$$t(\alpha) \approx \frac{z^2(\gamma)\tilde{\sigma}^2 c_3\cdot E\eta_1 r(g)}{\epsilon^2 E\tau_1 r(f)^2}$$

where $z(\gamma)$ solves $P\{N(0,1) \le z(\gamma)\} = 1 - \gamma/2$. Since the right-hand side does not depend on α, this shows that the discounting problem does not get harder, in a relative error sense, as the discount rate is driven to zero.

8. An Example: The $M/D/\infty/\infty$ Queue

In this section, we analytically compute the mean square error of the estimators $\delta_1(t)$, $\delta_2(t)$, and $\delta_4(t)$ for a special case of the $G/G/\infty/\infty$ queue described earlier in the paper.

Specifically, we shall analyze the $M/D/\infty/\infty$ queue, in which the arrival process is Poisson with rate λ and service times take on the constant value Δ. (This queue is more analytically tractable than $M/M/\infty/\infty$.) For our example, we let $f(x) = x$, $g(x) = \alpha$, and assume $c_1 = c_2 = c_3$.

If we let $\eta = (\eta(s) : s \geq 0)$ be the Poisson arrival process to the queue and $X(s)$ is the number of customers in the system at time s, it is clear that $X(s) = \eta(s) - \eta(s - \Delta)$, provided that we define $\eta(s) = 0$ for $s < 0$. Hence,

$$\begin{aligned} D_\alpha &= \int_0^\infty e^{-\alpha s} X(s) ds \\ &= \int_0^\Delta e^{-\alpha s} \eta(s) ds + \int_\Delta^\infty e^{-\alpha s} (\eta(s) - \eta(s - \Delta)) ds \\ &= (1 - e^{-\alpha\Delta}) \int_0^\infty e^{-\alpha s} \eta(s) ds. \end{aligned}$$

Then $d(\alpha) = ED_\alpha = (1 - e^{-\alpha\Delta}) \int_0^\infty e^{-\alpha s} \lambda s ds = (1 - e^{-\alpha\Delta})\lambda/\alpha^2$. Also, using the independent increments property of η, find that

$$\begin{aligned} \operatorname{var} D_\alpha &= (1 - e^{-\alpha\Delta})^2 2 \int_0^\infty \int_s^\infty e^{-\alpha s - \alpha u} \operatorname{cov}(\eta(s), \eta(u)) du ds \\ &= (1 - e^{-\alpha\Delta})^2 2 \int_0^\infty \int_s^\infty e^{-\alpha s - \alpha u} \lambda s du ds \\ &= (1 - e^{-\alpha\Delta})^2 \lambda / 2\alpha^3. \end{aligned}$$

Similar, but more involved, calculations show that if we choose $\beta > 2\Delta$,

$$\operatorname{var}\left(\int_0^\beta e^{-\alpha s} X(s) ds\right) = (1 - e^{-\alpha\Delta})^2 \lambda / 2\alpha^3 - \frac{\lambda e^{-2\alpha\beta}}{\alpha^2} \left(\frac{(e^{\alpha\Delta} - 1)}{\alpha} - \Delta\right). \quad (12)$$

We find that if $\beta > \Delta$,

$$\operatorname{bias}\left(\int_0^\beta e^{-\alpha s} X(s) ds\right) = \frac{\lambda \Delta e^{-\alpha\beta}}{\alpha}. \quad (13)$$

The theory of §2 suggests that we choose the truncation point $\beta = (\log t)/2\alpha$ when the computer budget equals t. The computer budget just equals the total simulated time over all replications. The corresponding number of replications is then $\lfloor 2\alpha t/\log t \rfloor$. Using (12) and (13), we find that if $t > \exp(4\alpha\Delta)$,

$$\operatorname{MSE}(\delta_1(t)) \approx \frac{\log t}{2\alpha t} (1 - e^{-\alpha\Delta})^2 \frac{\lambda}{2\alpha^3} + \frac{1}{t} \frac{\lambda^2 \Delta^2}{\alpha^2} - \frac{(\log t)}{2\alpha^3 t^2} \lambda \left(\frac{(e^{\alpha\Delta} - 1)}{\alpha} - \Delta\right). \quad (14)$$

(We write $\approx$ only because we have replaced the integer $\lfloor 2\alpha t/\log t \rfloor$ by $2\alpha t/\log t$.) To analyze the mean square error of $\delta_2(t)$, we observe that if $s \leq u$, then

$$EX(s)X(u) = \lambda^2 (s \wedge \Delta)(u \wedge \Delta) + \lambda((s - u + \Delta) \vee 0).$$

Straightforward integrations then prove that (see Theorem 2)

$$\begin{aligned} t\operatorname{MSE}(\delta_2(t)) \approx{}& (1 - e^{-\alpha\Delta}) \frac{\lambda^2}{\alpha^5} (5 + e^{-\alpha\Delta}) - \frac{6\lambda^2}{\alpha^4} \Delta e^{-\alpha\Delta} \\ &- \frac{\lambda^2 \Delta^2}{\alpha^3} e^{-\alpha\Delta} + \frac{2\lambda\Delta}{\alpha^3} (1 - e^{-\alpha\Delta}) - \frac{\lambda \Delta^2}{\alpha^2} e^{-\alpha\Delta}. \quad (15) \end{aligned}$$

To compute the mean square error of the regenerative estimator, we let 0 be our regeneration state. Then, $T(1) = \inf\{t \geq 0 : X(t-) \neq 0, X(t) = 0\}$ is the first regeneration time for X. Let $E_0(\cdot)$ ($\mathrm{var}_0(\cdot)$) represent the expectation (variance) operators, conditional on $X(0) = 0$.

THEOREM 7. *For* $\tau \geq 0$, $E_0 \exp(-\tau T(1)) = \lambda(\lambda + \tau \exp((\lambda + \tau)\Delta))^{-1}$.
Proposition 3 then proves that

$$\mathrm{var}_0(A(1) + C(1) \cdot EK(1)) = \mathrm{var}_0(D_\alpha) \cdot (1 - E_0 e^{-2\alpha T(1)})$$
$$= (1 - e^{-\alpha\Delta})^2 \cdot \frac{\lambda}{\alpha^2} \cdot \frac{e^{(\lambda+2\alpha)\Delta}}{\lambda + 2\alpha e^{(\lambda+2\alpha)\Delta}}.$$

Recalling that the stationary distribution π of X is Poisson with parameter $\lambda\Delta$, it follows that $\pi_0 = e^{-\lambda\Delta}$. Hence, $E_0 T(1) = (\lambda\pi_0)^{-1} = e^{\lambda\Delta}/\lambda$. Thus (see Theorem 4),

$$\mathrm{MSE}(\delta_4(t)) \approx (1 - e^{-\alpha\Delta})^2 \cdot \frac{1}{\alpha^4 t} \cdot \frac{(\lambda + \alpha e^{(\lambda+\alpha)\Delta})^2}{(\lambda + 2\alpha e^{(\lambda+2\alpha)\Delta})}. \tag{16}$$

Since X is not a semi-Markov process, estimators δ_3 and δ_5 are inapplicable.

Table 1 uses formulas (14), (15), and (16) to calculate the mean square error of estimators δ_1, δ_2, and δ_4 for various choices of the parameters t, Δ, and α. The table shows that if Δ is not too large, the regenerative estimator δ_4 is best. Recalling that the length of a regenerative cycle is $e^{\lambda\Delta}/\lambda$, we see that the regenerative cycle length $E_0 T(1)$ grows very rapidly with Δ. The lesson here is that the regenerative estimator is the method of choice when the regenerative cycles are not too long. We also see that if the discount rate α is large, the randomized estimator δ_2 wins, as suggested by (11b). Perhaps surprisingly, Table 1 indicates that δ_1 is competitive with δ_2 and δ_4 across the entire range of parameter values. We found that our choice of $\beta^*(t) = (\log t)/2\alpha$ gave reasonable performance in this example. (Recall that Theorem 1 shows that setting $\beta^*(t) = (\log t)/2\alpha + \xi$ is asymptotically optimal, for any choice of ξ.) The table also shows that each of our (applicable) estimators can have the lowest mean square error and that their mean square errors can differ widely in particular cases.

9. Conclusions

Table 2 reviews the basic properties of the estimators considered in this paper.

Our $M/D/\infty/\infty$ example has the fixed-time truncation estimator δ_1 performing well at essentially all combinations of t, α, and Δ, except that when Δ is small (and hence the regeneration cycles short) the regenerative estimator wins by a large margin. The

TABLE 1

			Mean Square Errors		
Δ	α	t	Truncated Estimator δ_1	Randomized Estimator δ_2	Regenerative Estimator δ_4
0.20	0.003	10^5	0.17	10.3	0.044
1.0	0.30	10^5	3.5×10^{-4}	4.71×10^{-4}	9.22×10^{-5}
10.0	0.03	10^7	0.045	0.37	2.75
20.0	0.30	10^{11}	5.22×10^{-8}	3.45×10^{-8}	8.94×10^{-2}
20.0	0.03	10^8	1.60×10^{-2}	0.14	18289.9

TABLE 2

Estimators	Truncated Estimator δ_1	Randomized Estimator δ_2	Randomized Semi-Markov Estimator δ_3	Regenerative Estimator δ_4	Regenerative Semi-Markov Estimator δ_5
truncates deterministically	Yes	No	No	No	No
truncates randomly	No	Yes	Yes	No	No
"conditions out" holding times when process is semi-Markov	No	No	Yes	No	Yes
uses regenerative structure	No	No	No	Yes	Yes
becomes less efficient as discount rate decreases	No	Yes	Yes	No	No
variance per run	see below	δ_3 beats δ_2		δ_5 beats δ_4	
overall efficiency	always loses for large enough computer budgets	δ_2, δ_3 always lose to δ_4, δ_5 for small enough discount rate		depends on expected work per run as well as variance per run; see text	

randomized estimator δ_2 occasionally wins (by a slight margin) when the discount rate is (artificially) high; its performance is otherwise terrible. When Δ is large (recall that the average length of a regeneration cycle here is $e^{\lambda\Delta}/\lambda$), the regenerative method breaks down. Too much should not be inferred from one example. However, the example indicates that fixed-time truncation is the least volatile estimator, in accordance with intuition.

The real contribution of this paper is for problems with computer budgets ample enough to afford pilot runs to get "ballpark" estimates of all auxiliary parameters required for our estimators and their respective variance constants. The work to compute the Laplace transforms can be estimated easily. If they require little work (and the process is semi-Markov and regenerative), we would narrow the set $\mathcal{S}$ of estimators to be compared in trial runs to δ_1 and δ_5. If the process is regenerative but not semi-Markov, we would narrow $\mathcal{S}$ to δ_1 and δ_4. If the process is semi-Markov but not regenerative, we would consider δ_1 and δ_3. If the process is neither regenerative nor semi-Markov, then the only possibilities are δ_1 and δ_2.

Estimating the variance constants gives the mean square error parametrically in the computer-time budget (assuming that budget is large). We can thus rank our estimators for large budgets, much larger than that spent on pilot runs. A direct comparison of sample mean square errors based on pilot runs could be misleading, because δ_2, δ_3, δ_4, and δ_5 converge at the canonical rate $t^{-1/2}$ whereas δ_1 converges slower (at rate $t^{-1/2} \log t$). We can also get the mean square errors parametrically in the discount rate via (11a)–(11e); this can be used as an additional theoretical guide for choosing the "production run" estimator.[1]

[1] We thank the referees for their comments and suggestions, which helped improve this paper.
Part of this research was carried out while Dr. Fox was on leave at the School of Operations Research and Industrial Engineering at Cornell University from the Department of Computer Science and Operations Research, University of Montreal.
Dr. Glynn's research was performed at University of Wisconsin–Madison and was supported by the National Science Foundation under Grant ECS-8404809 and the U.S. Army under Contract No. DAA29-84-K-0030.

Appendix

PROOF OF THEOREM 1. First, observe that the positivity of f shows that the bias $b(\beta)$ is positive for all β. Furthermore, by H2, it is evident that $b(\beta)$ converges to zero as $\beta \to \infty$, and thus a and b must be positive finite constants.

We now apply the results of Fox and Glynn (1989) to obtain the theorem; it is easily checked that their hypotheses are in force. Their Proposition 1 states that $\beta(t) \to \infty$ as $t \to \infty$ is necessary for consistency of $\delta_1(t)$, while their Theorem 2 proves that

$$q(t)(\delta_1(t) - d) \Rightarrow N(0, 1) + \gamma \tag{A1}$$

as $t \to \infty$, where $q(t) = (t/c_1\beta(t) \operatorname{var} D(\beta(t)))^{1/2}$, and $\gamma = \lim_{t\to\infty} q(t)b(\beta(t))$. Since $\beta(t) \to \infty$, it is evident that

$$\operatorname{var} D(\beta(t))/\operatorname{var} D \to 1 \tag{A2}$$

as $t \to \infty$. Furthermore,

$$t^{1/2}b(\beta(t)) \sim \frac{b(\beta(t))}{a \exp(-c\beta(t))} \cdot a \exp\left[-c\left(\beta(t) - \frac{1}{2c}\log t\right)\right]. \tag{A3}$$

Hence, if $\beta(t) = \beta^*(t)$, it follows that $t^{1/2}b(\beta(t))$ converges to a finite constant, so that $\gamma = 0$; part (i) is then obtained by using (A1) and (A2). Similarly, for part (ii), $t^{1/2}b(\beta(t)) \to 0$ so that $\gamma = 0$ and the result again follows immediately from (A1) and (A2). Finally for (iii), (A3) shows that for t sufficiently large,

$$t^{1/2}b(\beta(t)) \geq \frac{a}{2}\exp\left[\frac{1}{4}(1-\kappa)\log t\right] = \frac{a}{2}t^{(1-\kappa)/4}$$

so that $\gamma = \infty$, yielding the result.

PROOF OF PROPOSITION 1. We can write $\tilde{D}(1)$ as

$$\tilde{D}(1) = \int_0^\infty I(V^{-1}(R) > t)f(X(t))dt$$
$$= \int_0^\infty I(R > V(t))f(X(t))dt.$$

The result then follows from Lemma 1 below, by noting that the independence of X and K proves that $E\{I(R > V(t))f(X(t))|X\} = f(X(t))P\{R > V(t)|X\} = f(X(t))\exp(-V(t))$.

LEMMA 1. *Let Z be a nonnegative process on a probability space $(\Omega, \mathcal{F}, P)$. If $\mathcal{G}$ is a sub-σ-field of $\mathcal{F}$, then*

$$E\left\{\int_0^\infty Z(t)dt \mid \mathcal{G}\right\} = \int_0^\infty E\{Z(t)|\mathcal{G}\}dt \qquad a.s.$$

PROOF OF LEMMA 1. We use the defining relation for conditional expectation, as given on p. 298 of Chung (1974). Note that $\int_0^\infty E\{Z(t)|\mathcal{G}\}dt$ is a $\mathcal{G}$-measurable r.v. such that if $A \in \mathcal{G}$,

$$E\left(\int_0^\infty E\{Z(t)|\mathcal{G}\}dt \cdot I(A)\right) = \int_0^\infty E(I(A)E\{Z(t)|\mathcal{G}\})dt$$
$$= \int_0^\infty E(I(A)Z(t))dt$$
$$= E\left(\int_0^\infty Z(t)dt \cdot I(A)\right);$$

the first and third equalities use Fubini's theorem, whereas the second follows from the defining relation for $E\{Z(t)|\mathcal{G}\}$. We have therefore demonstrated that $\int_0^\infty E\{Z(t)|\mathcal{G}\}dt$ satisfies the defining relation for $E\{\int_0^\infty Z(t)dt|\mathcal{G}\}$, proving the result.

PROOF OF THEOREM 2. The CLT for $\delta_2(t)$ follows immediately from §5 of Glynn and Whitt (1986). For the expression for var $\tilde{D}(1)$, note that

$$E\tilde{D}(1)^2 = 2E\left\{\int_0^{V^{-1}(R)} \int_0^t f(X(s))f(X(t))dsdt\right\}$$

$$= 2E\left\{\int_0^\infty \int_0^t I(R > V(t))f(X(s))f(X(t))dsdt\right\}.$$

$$= \int_0^\infty \int_0^t E\{I(R > V(t))f(X(s))f(X(t))dsdt\}.$$

But $E\{I(R > V(t))f(X(s))f(X(t))|X\} = f(X(s))f(X(t))\cdot P\{R > V(t)\} = f(X(s))f(X(t))\exp(-V(t))$, yielding the formula. A similar proof gives the expression for $EV^{-1}(R)$.

PROOF OF PROPOSITION 2. For the first formula, note that

$$\begin{aligned} P\{M \geq m|Y\} &= P\{V^{-1}(R) > S_m|Y\} \\ &= P\{R > V(S_m)|Y\} \\ &= E\{P\{R > V(S_m)|X\}|Y\} \\ &= E\{\exp(-V(S_m))|Y\} \\ &= \prod_{j=0}^{m-1} E\{\exp(-g(Y_j)\beta_j)|Y\} \\ &= \prod_{j=0}^{m-1} \varphi(Y_j, Y_{j+1}, g(Y_j)), \end{aligned}$$

from which the result follows. For the second expression, observe that $E\{\beta_k|Y, M = m\} = E\{\beta_k I(M = m)|Y\}/P\{M = m|Y\}$. To analyze the numerator, note that for $k < m$,

$$\begin{aligned} E\{\beta_k I(M \geq m)|Y\} &= E\{E\{\beta_k I(R > V(S_m))|X\}|Y\} \\ &= E\{\beta_k \exp(-V(S_m))|Y\} \\ &= -\varphi_k' \cdot \prod_{\substack{j=0 \\ j\neq k}}^{m-1} \varphi_j \end{aligned}$$

so that $E\{\beta_k I(M = m)|Y\} = -\varphi_k' \cdot (1 - \varphi_m)\cdot \prod_{j=0; j\neq k}^{m-1} \varphi_j$. This, when combined with the first formula, yields the second identity.

The proof of the third formula follows a similar pattern. We write $E\{V^{-1}(R) - S_M|Y, M = m\} = E\{(V^{-1}(R) - S_M)I(M = m)|Y\}/P\{M = m|Y\}$; again, we handle the denominator using the first formula. For the numerator, we note that $E\{V^{-1}(R)I(M = m)|Y\}$ can be expressed as

$$\int_0^\infty P\{V^{-1}(R) > t; M = m|Y\}dt = \int_0^\infty P\{R > V(t); R > V(S_m)|Y\}dt - \int_0^\infty P\{R > V(t); R > V(S_{m+1})/|Y\}dt. \quad \text{(A4)}$$

By conditioning on X, we find that $P\{R > \max(V(t), V(S_k))|Y\} = E\{\exp(-\max(V(t), V(S_k)))|Y\}$. So by Lemma 1

$$\begin{aligned} \int_0^\infty P\{R > \max(V(t), V(S_k))|Y\}dt &= E\left\{\int_0^\infty \exp(-\max(V(t), V(S_k)))dt\right\} \\ &= E\{S_k \exp(-V(S_k))|Y\} + E\left\{\int_{S_k}^\infty \exp(-V(t)dt|Y\right\}. \quad \text{(A5)} \end{aligned}$$

On the other hand, $E\{S_m I(M = m)|Y\} = E\{S_m I(V(S_m) \leq R < V(S_{m+1}))|Y\} = E\{S_m(\exp(-V(S_m)) - \exp(-V(S_{m+1})))|Y\}$. Combining this with (A4) and (A5) shows that the numerator equals

$$E\left\{\exp(-V(S_m))\left(\int_0^{\beta_m}\exp(-g(Y_m)t)dt-\beta_m\exp(-g(Y_m)\beta_m)\right)\mid Y\right\}=\prod_{i=0}^{m-1}\varphi_i\left(\frac{1}{g(Y_m)}(1-\varphi_m)+\varphi_m'\right).$$

Dividing by the denominator gives the third formula.

PROOF OF THEOREM 4. Standard weak convergence arguments prove that

$$\delta_4(t)=[\bar{A}_t(0)+\bar{C}_t(0)EK(1)]+\left[\frac{EC(0)}{1-EC(1)}(\bar{A}_t(1)-EK(1)(1-\bar{C}_t(1)))\right]+o_p(t^{-1/2})$$

where $o_p(t^{-1/2})$ represents a process $\chi(t)$ such that $t^{1/2}\chi(t)\Rightarrow 0$. The random time-change results of §5 of Glynn and Whitt (1986) can now be applied to the bracketed terms above to obtain the result. (To show that the respective variances are finite, see the proof of Proposition 3.)

PROOF OF PROPOSITION 3. The regenerative structure of X proves that ($\overset{\mathcal{D}}{=}$denotes equality in distribution)

$$D\overset{\mathcal{D}}{=}A(0)+C(0)K(1)$$

where $(A(0), C(0))$ is independent of $K(1)$. Squaring both sides and taking expectations, we get

$$ED^2=EA(0)^2+2EA(0)C(0)EK(1)+EC(0)^2EK(1)^2. \tag{A6}$$

Since all terms on the right-hand side are positive, we see that H2 implies the finiteness of all the qualities appearing there. We apply the same analysis to $K(1)$:

$$K(1)\overset{\mathcal{D}}{=}A(1)+C(1)K(2).$$

Using the fact that $K(2)\overset{\mathcal{D}}{=}K(1)$, we get $EK(1)^2=(EA(1)^2+2EA(1)C(1)\cdot EK(1))(1-EC(1)^2)^{-1}\cdot(EC^2(1)<1$ by H1). Substituting this into (A6) yields the result, after algebraic simplification.

PROOF OF THEOREM 6. For (a) we use (8) and let $\alpha\downarrow 0$. A Taylor expansion gives

$$\begin{aligned}A_\alpha(0)&=\int_0^{T_0}\exp(-\alpha V(t))f(X(t))dt\\&=\int_0^{T_0}\left[1-\alpha V(t)+\frac{\alpha^2}{2}V^2(t)e^{-\gamma V(t)}\right]f(X(t))dt\end{aligned}$$

where $\gamma=\gamma(\alpha)\in(0,\alpha)$. Since $V^2(t)\exp(-\gamma(\alpha)V(t))f(X(t))\le V^2(\tau)f(X(t))$ uniformly in α and

$$\int_0^{T_0}V^2(\tau)f(X(t))dt=Y_0(f)Y_0(g)^2$$

is integrable by H9, it follows from the dominated convergence theorem that

$$EA_\alpha(0)=EY_0(f)-\alpha E\left\{\int_0^{T_0}V(t)f(X(t))dt\right\}+\frac{\alpha^2}{2}E\left\{\int_0^{T_0}V^2(t)f(X(t))dt\right\}+o(\alpha^2). \tag{A7}$$

Similarly, one can show that

$$EC_\alpha(0)=1-\alpha EY_0(y)+\frac{\alpha^2}{2}EY_0(y)^2+o(\alpha^2).$$

Corresponding expressions for $EA_\alpha(1)$ and $EC_\alpha(1)$ lead immediately to (a). For (b), (h), (i), we use Propositions 3 and 4 and arguments similar to the above. Relation (c) can be found in Glynn and Whitt (1988).

For (e), observe that

$$V_\alpha^{-1}(R)\overset{\mathcal{D}}{=}L_\alpha(0)+I_\alpha(0)Q_\alpha(1) \tag{A8}$$

where $L_\alpha(0)=V_\alpha^{-1}(R)\wedge T_0$, $I_\alpha(0)=I(V_\alpha^{-1}(R)>T_0)$, $Q_\alpha(1)=V_\alpha^{-1}(r)-T_0$, and $(L_\alpha(0), I_\alpha(0))$ is independent of $Q_\alpha(1)$. Furthermore, on $\{V_\alpha^{-1}(R)>T_0\}$,

$$Q_\alpha(1)\overset{\mathcal{D}}{=}L_\alpha(1)+I_\alpha(1)Q_\alpha(2) \tag{A9}$$

where $L_\alpha(1)=(V_\alpha^{-1}(R)\wedge T_1)-T_0$, $I_\alpha(1)=I(V_\alpha^{-1}(R)>T_1)$, $Q_\alpha(2)=V_\alpha^{-1}(R)-T_1$, $(L_\alpha(1), I_\alpha(1))$ is independent of $Q_\alpha(2)$, and the distribution of $Q_\alpha(i)$ conditional on $\{V_\alpha^{-1}(R)>T_i\}$ is independent of i ($i=0$, 1). Taking expectations in (A8) and (A9) and using the independence leads to an expression similar to (8). One then expands the expectations in a manner similar to (A7) to obtain (e). Results (d), (f), and (g) are proved using decompositions analogous to (A8) and (A9), followed by Taylor expansions for small α.

PROOF OF THEOREM 7. We let $E_1(\cdot)$ denote the expectation operator, conditional on $X(0)=1$ and service starting on the single customer at time $t = 0$. Then, since interarrivals are exponential with parameters λ, we find that

$$E_0 \exp(-\tau T(1)) = \frac{\lambda}{\lambda + \tau} E_1 \exp(-\tau T(1)).$$

If V is the arrival time of the next customer, we see that

$$\begin{aligned} E_1 e^{-\tau T(1)} &= E_1\{e^{-\tau T(1)}; V > \Delta\} + E_1\{e^{-\tau T(1)}; V \leq \Delta\} \\ &= E_1\{e^{-\tau\Delta}; V > \Delta\} + \int_0^\Delta E_1\{e^{-\tau T(1)}; V \in dv\} \\ &= e^{-\tau\Delta} P_1\{V > \Delta\} + \int_0^\Delta E_1 e^{-\tau(T(1)+v)} \lambda e^{-\lambda v} dv \\ &= e^{-(\tau+\lambda)\Delta} + E_1 e^{-\tau T(1)} \frac{\lambda}{\lambda + \tau}(1 - e^{-(\lambda+\tau)\Delta}). \end{aligned}$$

Solving for $E_1 e^{-\tau T(1)}$, we get $E_1 e^{-\tau T(1)} = (\lambda + \tau)/(\lambda + \tau e^{(\lambda+\tau)\Delta})$. Thus, $E_0 e^{-\tau T(1)} = \lambda/(\lambda + \tau e^{(\lambda+\tau)\Delta})$.

References

BRATLEY, P., B. L. FOX AND L. E. SCHRAGE, *A Guide to Simulation* (2nd ed), Springer-Verlag, New York, 1987.

CHUNG, K. L., *A Course in Probability Theory*, Academic Press, New York, 1974.

ÇINLAR, E., *Introduction to Stochastic Processes*, Prentice-Hall, Englewood Cliffs, NJ, 1975.

FOX, B. L. AND P. W. GLYNN, "Discrete-time Conversion for Simulating Semi-Markov Processes," *Oper. Res. Lett.*, (1986), 191–196.

——— AND ———, "Replication Schemes for Limiting Expectations," *Probability in the Engineering and Informational Sciences,* 3 (1989), 299–318.

GLYNN, P. W. AND W. WHITT, "The Efficiency of Simulation Estimators," technical report, AT&T Bell Laboratories, Holmdel, NJ, 1986.

——— AND ———, "Limit Theorems for State-Dependent Discounted Costs," forthcoming technical report.

HAMMERSLEY, J. M. AND D. C. HANDSCOMB, *Monte Carlo Methods*, Chapman and Hall, New York, 1964.

VEINOTT, JR., A. F., "Discrete Dynamic Programming with Sensitive Discount Optimality Criteria," *Ann. Math. Statist.*, 40 (1969), 1635–1660.

WHITT, W., "Stochastic Abelian and Tauberian Theorems," *Z. Wahrsch. verw. Gerbiete*, 22 (1972), 251–267.