

IMPORTANCE SAMPLING USING THE SEMI-REGENERATIVE METHOD

James M. Calvin

Department of Computer Science
New Jersey Institute of Technology
Newark, New Jersey 07102 U.S.A.

Peter W. Glynn

Department of Management
Science and Engineering
Stanford University
Stanford, CA 94305 U.S.A.

Marvin K. Nakayama

Department of Computer Science
New Jersey Inst. of Technology
Newark, New Jersey 07102 U.S.A.

ABSTRACT

We discuss using the semi-regenerative method, importance sampling, and stratification to estimate the expected cumulative reward until hitting a fixed set of states for a discrete-time Markov chain on a countable state space. We develop a general theory for this problem and present several central limit theorems for our estimators. We also present some empirical results from applying these techniques to simulate a reliability model.

1 INTRODUCTION

Importance sampling is a variance-reduction technique that, if applied properly, can lead to significant decreases in the variance when estimating performance measures related to rare events. The basic idea is to change the dynamics of the system so as to cause the rare event of interest to occur more frequently. Unbiased estimators are recovered by multiplying the resulting samples by a correction factor known as the likelihood ratio. In the rare-event context, importance sampling is often used in the estimation of the mean of a non-negative random variable Z for which $Z = 0$ with high probability, and $Z > 0$ with small probability. For more details on importance sampling, see Glynn and Iglehart (1989) and Heidelberger (1995).

In this paper, we consider an irreducible, positive-recurrent discrete-time Markov chain (DTMC) X on a discrete state space S , and we examine the estimation of $\eta(w)$, which is the expected cumulative reward until hitting some set of states $S_0 \subset S$ given that the system starts in state

w . One approach to estimating $\eta(w)$ is to run replications, where each replication begins with the system in state w and ends the first time S_0 is hit. When S_0 is a rare set, standard simulation is inefficient since the replications will typically be extremely long and highly variable. But even importance sampling may be ineffective for estimating $\eta(w)$ when using replications because each replication may still be long and the variance of estimators under static importance sampling grows exponentially in the length of a replication (Glynn 1995).

A way to circumvent this problem is to use the regenerative method (Crane and Iglehart 1975; Shedler 1993). A regenerative process has the property that there is an infinite sequence of stopping times, known as regeneration points, at which times the system probabilistically restarts, and for our DTMC X , the successive hitting times to any fixed state form a regeneration sequence. We can break up a sample path of X into i.i.d. cycles based on the regeneration points, and importance sampling need only be used within regenerative cycles.

We can further shorten the time in which importance sampling is applied by using the semi-regenerative method (Calvin, Glynn, and Nakayama 2001). The semi-regenerative method gets its name because of its close relationship to semi-regenerative processes (Çinlar 1975, Section 10.6). To apply this approach to our DTMC X , we fix a set of states A , and we break up a sample path of X into *trajectories* determined by successive entrances into the set A . Importance sampling thus only needs to be applied to individual trajectories, which are shorter than regenerative cycles, and thus may lead to smaller variance

than when using the regenerative method with importance sampling.

The rest of this paper is organized as follows. Section 2 develops the mathematical framework, and in Section 3 we review the estimation of $\eta(w)$ using the regenerative method, both without and with importance sampling. We describe the semi-regenerative method in Section 4, both without and with importance sampling. We also consider combining importance sampling with stratification. We present some empirical results in Section 5.

2 MATHEMATICAL FRAMEWORK

Let $X = \{X_j : j = 0, 1, 2, \dots\}$ be a discrete-time Markov chain on a finite or countably infinite state space S . Let $Q = (Q(x, y) : x, y \in S)$ be the transition probability matrix of X , and let P_x (resp., E_x , Var_x , and Cov_x) denote the probability measure (resp., expectation, variance, and covariance) given that $X_0 = x$, $x \in S$. Let \mathcal{P} be the family of probability measures $\{P_x : x \in S\}$.

Assumption 1. X with transition probability matrix Q is irreducible and positive recurrent.

Under Assumption 1, X has a unique stationary distribution $\pi = (\pi(x) : x \in S)$, which is the row-vector solution to $\pi = \pi Q$ with $\sum_{x \in S} \pi(x) = 1$ and $\pi(x) > 0$ for all $x \in S$.

Let $f : S \rightarrow \mathfrak{R}$ be a “reward” function such that $f(x) \geq 0$, $x \in S$. Let $S_0 \subset S$, and define $\Gamma = \inf\{n \geq 0 : X_n \in S_0\}$. Fix an initial state $w \in S$ with $w \notin S_0$, and our goal is to estimate

$$\eta(w) = E_w \left[\sum_{j=0}^{\Gamma-1} f(X_j) \right],$$

which is the expected cumulative reward until S_0 is hit, given that the chain starts in state w .

3 THE REGENERATIVE METHOD

For $x \in S$, define $\tau_x = \inf\{j \geq 1 : X_j = x\}$. Using $w \in S$ as a “return state,” one can show (e.g., Goyal et al. 1992) that

$$\eta(w) = \frac{E_w[U]}{E_w[V]}, \quad (1)$$

where

$$\begin{aligned} U &= \sum_{j=0}^{(\tau_w \wedge \Gamma)-1} f(X_j), \\ V &= I(\Gamma < \tau_w), \end{aligned}$$

$a_1 \wedge a_2 = \min(a_1, a_2)$ for $a_1, a_2 \in \mathfrak{R}$, and $I(C)$ is the indicator function of an event C , i.e., $I(C) = 1$ if C occurs, and is 0 otherwise.

3.1 Standard Simulation

Using (1) we can apply the regenerative method to estimate $\eta(w)$ by generating independent copies of (U, V) under measure P_w and forming sample means. Specifically, let $T_{w,0} = \inf\{j \geq 0 : X_j = w\}$ and $T_{w,k} = \inf\{j > T_{w,k-1} : X_j = w\}$ for $k \geq 1$. Define $\tau_{w,k} = T_{w,k} - T_{w,k-1}$, for $k \geq 1$. Define $\Gamma_{w,k} = \inf\{j \geq T_{w,k-1} : X_j \in S_0\}$, for $k \geq 1$. Also, define $U_{w,k} = \sum_{j=T_{w,k-1}}^{(T_{w,k} \wedge \Gamma_{w,k})-1} f(X_j)$ and $V_{w,k} = I(\Gamma_{w,k} < T_{w,k})$ for $k \geq 1$. Now fix an integer n large and run a simulation of X up to time $T_{w,n}$, giving a sample path $\{X_j : j = 0, 1, \dots, T_{w,n}\}$. The $(U_{w,k}, V_{w,k})$, $k = 1, 2, \dots, n$, are i.i.d. copies of (U, V) under measure P_w . Then the regenerative estimator of $\eta(w)$ is

$$\tilde{\eta}_n(w) = \frac{(1/n) \sum_{k=1}^n U_{w,k}}{(1/n) \sum_{k=1}^n V_{w,k}}.$$

Let $\mathcal{N}(\kappa, \Phi)$ denote a multivariate normal distribution with mean vector κ and covariance matrix Φ . Also, let $\xrightarrow{\mathcal{D}}$ denote convergence in distribution. We can form an asymptotically valid confidence interval for $\eta(w)$ based on the following central limit theorem, which is a slight variation of one appearing in Shedler (1993), p. 100.

Proposition 1. If Assumption 1 holds and if $E_w[U^2] < \infty$, then

$$n^{1/2}(\tilde{\eta}_n(w) - \eta(w)) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \tilde{\sigma}^2)$$

as $n \rightarrow \infty$, where

$$\begin{aligned} \tilde{\sigma}^2 &= \frac{1}{E_w[V]} (\text{Var}_w[U] - 2\eta(w)\text{Cov}_w(U, V) \\ &\quad + \eta(w)^2 \text{Var}_w[V]). \end{aligned}$$

3.2 Importance Sampling

It turns out that in certain situations, the denominator in (1) may be difficult to estimate (but not the numerator). For example, this is true when $\eta(w)$ corresponds to the expected time to buffer overflow in a stable single-server queue with a large buffer and $w = 0$, and also when $\eta(w)$ is the mean time to failure of a highly reliable system and w is the state with all components operational; see Heidelberger (1995). To understand the difficulty in estimating the denominator in (1) in these settings, note that (under measure P_w) most w -cycles end before hitting S_0 , so $V = I(\Gamma < \tau_w) = 0$ with high probability, say $1 - \epsilon$, where $\epsilon > 0$ is small. Hence,

$E_w[V] = \epsilon$ and since $V^2 = V$, $\text{Var}_w[V] = E_w[V] - (E_w[V])^2 = \epsilon - \epsilon^2 \approx \epsilon$. Thus, the coefficient of variation (CV) of V is approximately $\epsilon^{1/2}/\epsilon = \epsilon^{-1/2} \rightarrow \infty$ as $\epsilon \rightarrow 0$. To see what this means in practice, note that the expected relative half-width of a 95% confidence interval for $E_w[V]$ based on a sample of size n is roughly $1.96 \text{CV}/\sqrt{n} \approx 1.96/\sqrt{\epsilon n}$. Hence, as $\epsilon \rightarrow 0$, the number of samples required to obtain a confidence interval of a specified relative width grows to infinity, so it becomes more and more difficult to estimate $E_w[V]$ as the event $\{\Gamma < \tau_w\}$ becomes rarer. So we need another approach to estimate the denominator.

We now describe the use of importance sampling and the regenerative method to estimate $\eta(w)$. Let $\tilde{\mathcal{F}}_w$ denote the filtration of the process X up to time τ_w with $X_0 = w$. Define \tilde{P}_w to be the probability measure on $\tilde{\mathcal{F}}_w$ for the process X under the transition probability matrix Q given $X_0 = w$. Now suppose that we define another probability measure \tilde{P}_w^* (not necessarily Markovian) on $\tilde{\mathcal{F}}_w$ for X conditional on $X_0 = w$, and let \tilde{E}_w^* and $\tilde{\text{Var}}_w^*$ be the corresponding expectation and variance operators. We need to assume the following.

Assumption 2. \tilde{P}_w is absolutely continuous with respect to \tilde{P}_w^* .

By the Radon-Nikodym theorem (Theorem 32.2 of Billingsley 1995), Assumption 2 guarantees the existence of a non-negative random variable \tilde{L}_w for which

$$\tilde{P}_w(C) = \tilde{E}_w^*[I(C)\tilde{L}_w], \quad C \in \tilde{\mathcal{F}}_w. \quad (2)$$

Equation (2) is known as a *change of measure*, and the random variable $\tilde{L}_w = d\tilde{P}_w/d\tilde{P}_w^*$ is called the *likelihood ratio* (or Radon-Nikodym derivative) of \tilde{P}_w with respect to \tilde{P}_w^* (given $X_0 = w$). For example, if the measure \tilde{P}_w^* is induced by a transition probability matrix $Q^* = (Q^*(x, y) : x, y \in S)$, then Assumption 2 will hold if $Q^*(x, y) = 0$ implies $Q(x, y) = 0$ for all $x, y \in S$, and the likelihood ratio for the regenerative cycle $X_0, X_1, \dots, X_{\tau_w}$ given $X_0 = w$ is

$$\tilde{L}_w = \prod_{j=0}^{\tau_w-1} \frac{Q(X_j, X_{j+1})}{Q^*(X_j, X_{j+1})}.$$

Now (2) implies that we can rewrite (1) as

$$\eta(w) = \frac{E_w[U]}{\tilde{E}_w^*[V\tilde{L}_w]}, \quad (3)$$

which suggests estimating the denominator by using the measure \tilde{P}_w to generate samples of $V\tilde{L}_w$, and this is the basic idea of importance sampling. Note that in (3), we applied a change of measure to only the denominator and not the numerator. This leads to the follow-

ing method known as *measure-specific importance sampling* (Goyal et al. 1992). Fix a simulation budget n , which is the total number of w -cycles to simulate, and fix $0 < \delta < 1$. Using the original measure P_w , generate $n_1 \equiv \lfloor \delta n \rfloor$ regenerative w -cycles to yield n_1 observations of U , which we denote by U_1, U_2, \dots, U_{n_1} . (For $a \in \mathfrak{R}$, $\lfloor a \rfloor$ denotes the greatest integer less than or equal to a .) Independently, use the importance-sampling measure \tilde{P}_w^* to generate $n_2 \equiv \lfloor (1 - \delta)n \rfloor$ regenerative w -cycles to yield n_2 observations of (V, \tilde{L}_w) , which we denote by $(V_1^*, \tilde{L}_1), (V_2^*, \tilde{L}_2), \dots, (V_{n_2}^*, \tilde{L}_{n_2})$. Then our point estimator of $\eta(w)$ is

$$\tilde{\eta}_{n,\delta}^*(w) = \frac{(1/n_1) \sum_{k=1}^{n_1} U_k}{(1/n_2) \sum_{k=1}^{n_2} V_k^* \tilde{L}_k},$$

which satisfies the following central limit theorem from Goyal et al (1992).

Proposition 2. *If Assumptions 1 and 2 hold and if $E_w[U^2] < \infty$ and $\tilde{E}_w^*[V\tilde{L}_w^2] < \infty$, then*

$$n^{1/2}(\tilde{\eta}_{n,\delta}^*(w) - \eta(w)) \xrightarrow{D} \mathcal{N}(0, \tilde{\sigma}_*^2)$$

as $n \rightarrow \infty$, where

$$\tilde{\sigma}_*^2 = \frac{1}{E_w[V]} \left(\frac{\text{Var}_w[U]}{\delta} + \eta(w)^2 \frac{\tilde{\text{Var}}_w^*[V\tilde{L}_w]}{(1 - \delta)} \right).$$

4 THE SEMI-REGENERATIVE METHOD

Calvin, Glynn, and Nakayama (2001) develop another estimator for $\eta(w)$, which we now describe. Fix a set of states $A \subset S$, $A \neq \emptyset$, and we assume that $w \in A$ and $A \cap S_0 = \emptyset$. Set

$$\begin{aligned} T_0 &= \inf\{j \geq 0 : X_j \in A\}; \\ T_k &= \inf\{j > T_{k-1} : X_j \in A\}, \quad k \geq 1; \\ T &= T_1; \\ W_k &= X_{T_k}, \quad k \geq 1. \end{aligned}$$

For each $k \geq 1$, we call the sample path segment $\{X_j : T_{k-1} \leq j \leq T_k\}$ a *trajectory* of the process X .

Proposition 3. *Under Assumption 1, $W = \{W_k : k \geq 0\}$ is an irreducible, positive-recurrent discrete-time Markov chain with state space A .*

The process W is sometimes called the “chain on A .” Define

$$R(x, y) = P_x(X_T = y) \quad (4)$$

for $x, y \in A$, and let $R = (R(x, y) : x, y \in A)$, which is the transition probability matrix of W . Under Assumption 1, Proposition 3 implies the existence of a unique stationary distribution $\nu = (\nu(x) : x \in A)$ for W ; i.e., ν is the row vector satisfying $\nu R = \nu$ with $\sum_{x \in A} \nu(x) = 1$ and $\nu(x) > 0$ for all $x \in A$. Note that $\nu(x) = \pi(x) / \sum_{y \in A} \pi(y)$.

Assumption 3. $|A| = d < \infty$, with $A = \{x_1, x_2, \dots, x_d\}$.

For $x \in A$, note that

$$\begin{aligned} \eta(x) &\equiv E_x \left[\sum_{j=0}^{\Gamma-1} f(X_j) \right] \\ &= E_x \left[\sum_{j=0}^{(T \wedge \Gamma)-1} f(X_j) \right] \\ &+ \sum_{y \in A} E_x [I(X_T = y, \Gamma > T)] \eta(y). \end{aligned}$$

For $x, y \in A$, define

$$\begin{aligned} B(x) &= \sum_{j=0}^{(T \wedge \Gamma)-1} f(X_j), \text{ given } X_0 = x, \\ \phi(x, y) &= I(X_T = y, \Gamma > T), \text{ given } X_0 = x, \end{aligned}$$

and set

$$\begin{aligned} b(x) &= E[B(x)], \\ K(x, y) &= E[\phi(x, y)]. \end{aligned}$$

Let $\eta = (\eta(x) : x \in A)$, $b = (b(x) : x \in A)$, and $K = (K(x, y) : x, y \in A)$, and note that

$$\eta = b + K\eta.$$

Proposition 4. If $|b| < \infty$ and if Assumptions 1 and 3 hold, then $\sum_{m=0}^{\infty} K^m b = (I - K)^{-1} b$ and

$$\eta = (I - K)^{-1} b. \quad (5)$$

When the set $A = \{w\}$, (5) is equivalent to the ratio formula in (1). In general when $|A| > 1$, note the similarities between b and $E_w[U]$, and between $(I - K)^{-1}$ and $1/E_w[V]$.

4.1 Standard Simulation

We now present a semi-regenerative estimator for η based on (5) when simulating one sample path of X up to time T_n . For $x \in A$, define $H_n(x) = \sum_{k=0}^{n-1} I(W_k = x)$. Also, define $T_1(x) = \inf\{j \geq 0 : X_j = x\}$, and for $k \geq 2$, define $T_k(x) = \inf\{j > T_{k-1}(x) : X_j = x\}$. Also, define

$\tilde{T}_k(x) = \inf\{j > T_k(x) : X_j \in A\}$, which is the first time after $T_k(x)$ that X enters A again. Note that the sample-path segment $\{X_j : T_k(x) \leq j < \tilde{T}_k(x)\}$ is the k th trajectory starting in state $x \in A$. For $k \geq 1$, define $\Gamma_k(x) = \inf\{j > T_k(x) : X_j \in S_0\}$. For $x, y \in A$, let

$$\begin{aligned} B_k(x) &= \sum_{j=T_k(x)}^{(\tilde{T}_k(x) \wedge \Gamma_k(x))-1} f(X_j), \\ \phi_k(x, y) &= I(X_{\tilde{T}_k(x)} = y, \Gamma_k(x) > \tilde{T}_k(x)). \end{aligned}$$

Then define the estimators of b and K to be $b_n = (b_n(x) : x \in A)$ and $K_n = (K_n(x, y) : x, y \in A)$, respectively, with

$$\begin{aligned} b_n(x) &= \frac{\sum_{k=0}^{n-1} \sum_{j=T_k}^{(\Gamma_k \wedge \tilde{T}_{k+1})-1} f(X_j) I(W_k = x)}{\sum_{k=0}^{n-1} I(W_k = x)} \\ &= \frac{1}{H_n(x)} \sum_{k=1}^{H_n(x)} B_k(x) \end{aligned}$$

and

$$\begin{aligned} K_n(x, y) &= \frac{\sum_{k=0}^{n-1} I(W_k = x, W_{k+1} = y, \Gamma_k > T_{k+1})}{\sum_{k=0}^{n-1} I(W_k = x)} \\ &= \frac{1}{H_n(x)} \sum_{k=1}^{H_n(x)} \phi_k(x, y), \end{aligned}$$

where $\Gamma_k = \inf\{j > T_k : X_j \in S_0\}$. Then we define our semi-regenerative estimator of η based on a run-length of T_n to be

$$\eta_n = (I - K_n)^{-1} b_n,$$

where $\eta_n = (\eta_n(x) : x \in A)$.

Calvin, Glynn, and Nakayama (2001) show that

$$n^{1/2}(K_n - K, b_n - b) \xrightarrow{\mathcal{D}} (N_1, N_2) \quad (6)$$

as $n \rightarrow \infty$, where N_1 is a normal random matrix and N_2 is a normal random vector, and the following central limit theorem.

Theorem 5. If $E[B(x)^2] < \infty$ for all $x \in A$ and if Assumptions 1 and 3 hold, then

$$n^{1/2}(\eta_n - \eta) \xrightarrow{\mathcal{D}} (I - K)^{-1} N_1 \eta + (I - K)^{-1} N_2$$

as $n \rightarrow \infty$, where (N_1, N_2) is defined in (6). In particular, for each $k = 1, 2, \dots, d$,

$$n^{1/2}(\eta_n(x_k) - \eta(x_k)) \xrightarrow{D} \mathcal{N}(0, \sigma_k^2)$$

as $n \rightarrow \infty$, where

$$\sigma_k^2 = \sum_{i=1}^d \frac{J(x_k, x_i)^2}{v(x_i)} \left[v_i + 2 \sum_{j=1}^d \eta(x_j) s_{ij} + \sum_{j=1}^d \sum_{l=1}^d \eta(x_j) \eta(x_l) \Delta_i(x_j, x_l) \right], \quad (7)$$

$J = (J(x, y) : x, y \in A)$ with $J = (I - K)^{-1}$, $v_i = \text{Var}(B(x_i))$, $s_{ij} = \text{Cov}(\phi(x_i, x_j), B(x_i))$, and $\Delta_i(x_j, x_l) = \text{Cov}(\phi(x_i, x_j), \phi(x_i, x_l))$.

4.2 Importance Sampling

In certain contexts, it may be difficult to estimate $I - K$ using the standard-simulation approach in Section 4.1, so we now discuss how one can apply importance sampling with the semi-regenerative method.

Let $\mathcal{F}_{x,T}$ denote the filtration of the process X up to time T with $X_0 = x$. For $x \in A$, define $P_{x,T}$ to be the probability measure on $\mathcal{F}_{x,T}$ for the process X under the transition probability matrix Q given $X_0 = x$. Now suppose that for each $x \in A$, we define another probability measure $P_{x,T}^*$ (not necessarily Markovian) on $\mathcal{F}_{x,T}$ for X conditional on $X_0 = x$, and let $E_{x,T}^*$ be the corresponding expectation. Also, let P^* (resp., E^* , Var_* , and Cov_*) be the probability measure (resp., expectation, variance, and covariance) induced by the collection of measures $(P_{x,T}^* : x \in A)$. Finally, let P_x^* be the probability measure under P^* given $X_0 = x$, and let E_x^* be its expectation operator. We assume the following.

Assumption 4. For each $x \in A$, $P_{x,T}$ is absolutely continuous with respect to $P_{x,T}^*$.

Assumption 4 guarantees the existence of a non-negative random variable $L(x)$ for which

$$P_{x,T}(C) = E_{x,T}^*[I(C)L(x)], \quad C \in \mathcal{F}_{x,T}. \quad (8)$$

The random variable $L(x) = dP_{x,T}/dP_{x,T}^*$ is the likelihood ratio of $P_{x,T}$ with respect to $P_{x,T}^*$ (given $X_0 = x$). Observe that we allow for the possibility that the $P_{x,T}^*$ are different for different states $x \in A$.

We use the importance-sampling measure P^* to generate a sample path $\{X_j : j \geq 0\}$ of the process X as follows. Given a starting state $x_0 \in A$, set $X_0 = x_0$, so $T_0 = 0$. Then using measure $P_{x_0}^*$, generate a sequence of

states until set A is hit again, thereby yielding the trajectory X_1, X_2, \dots, X_{T_1} . Now from state X_{T_1} , use measure $P_{X_{T_1}}^*$ to generate a sequence of states until A is hit again, yielding $X_{T_1+1}, X_{T_1+2}, \dots, X_{T_2}$. In general, at the k th hit to set A , the process is in state X_{T_k} , and we use measure $P_{X_{T_k}}^*$ to generate a sequence of states until A is hit again, yielding $X_{T_k+1}, X_{T_k+2}, \dots, X_{T_{k+1}}$. We define the process $W = \{W_k : k \geq 0\}$ by letting $W_k = X_{T_k}$.

The process X defined in this way may no longer be a Markov chain since we did not assume any particular structure (other than Assumption 4) for the measure P^* . On the other hand, no matter how the $P_{x,T}^*$, $x \in A$, are defined, the embedded process W is always a Markov chain.

Proposition 6. If Assumptions 1, 3, and 4 hold, then W under measure P^* is an irreducible, positive-recurrent discrete-time Markov chain on A .

Define matrix $\Phi = (\Phi(x, y) : x, y \in A)$ with elements $\Phi(x, y) = P_{x,T}^*(X_T = y)$, and note that Φ is the transition probability matrix of W under the measure P^* . As shown in Proposition 6, Assumptions 1, 3, and 4 ensure that Φ is irreducible and positive recurrent, so Φ has a stationary distribution $\rho = (\rho(x) : x \in A)$; i.e., $\rho\Phi = \rho$ with $\rho(x) > 0$ for each $x \in A$ and $\sum_{x \in A} \rho(x) = 1$.

For $x \in A$, let

$$\begin{aligned} \psi(x, x) &= I(X_T = x, \Gamma \leq T) + I(X_T \neq x), \quad \text{given } X_0 = x, \end{aligned}$$

and for $x, y \in A$, $x \neq y$, let

$$\psi(x, y) = -I(X_T = y, \Gamma > T), \quad \text{given } X_0 = x.$$

Define $M \equiv I - K$, and note that for $x, y \in A$,

$$M(x, y) = E_x[\psi(x, y)] = E_x^*[\psi(x, y)L(x)] \quad (9)$$

by (8). The expression on the far right of (9) suggests using importance sampling to estimate the entries in the matrix M .

We now describe how to apply measure-specific importance sampling to estimate η with the semi-regenerative method. Fix a simulation budget n , which is the total number of trajectories to simulate, and fix $0 < \delta < 1$. Using the original measure P , generate a sample path $\{X_j : j = 0, 1, \dots, T_{n^*}\}$, where $n^* \equiv \lfloor \delta n \rfloor$. From this sample path, for $x \in A$, compute $B_k(x)$, $T_k(x)$, $\tilde{T}_k(x)$, and $\Gamma_k(x)$ as in Section 4.1. Then define the estimator of b to

be $b'_n = (b'_n(x) : x \in A)$ with

$$\begin{aligned} b'_n(x) &= \frac{\sum_{k=0}^{n_1^*-1} \sum_{j=T_k}^{(\Gamma_k \wedge T_{k+1})-1} f(X_j) I(W_k = x)}{\sum_{k=0}^{n_1^*-1} I(W_k = x)} \\ &= \frac{1}{H'_n(x)} \sum_{k=1}^{H'_n(x)} B_k(x), \end{aligned}$$

where $H'_n(x) = \sum_{k=0}^{n_1^*-1} I(W_k = x)$.

Independently of $\{X_j : j = 0, 1, \dots, T_{n_1^*}\}$, use the importance-sampling measure P^* to generate another path $\{X_j^* : j = 0, 1, 2, \dots, T_{n_2^*}\}$, with $n_2^* = \lfloor (1 - \delta)n \rfloor$, where $T_0^* = 0$ and $T_k^* = \inf\{j > T_{k-1}^* : X_j^* \in A\}$. Also, for $x \in A$, define $T_1^*(x) = \inf\{j \geq 0 : X_j^* = x\}$, and let $T_k^*(x) = \inf\{j > T_{k-1}^*(x) : X_j^* = x\}$ for $k \geq 2$, and let $\bar{T}_k^*(x) = \inf\{j > T_k^*(x) : X_j^* \in A\}$ for $k \geq 1$. For $x \in A$, let

$$L_k(x) = \frac{d\zeta(X_{T_k^*(x)}^*, \dots, X_{\bar{T}_k^*(x)}^*)}{d\zeta^*(X_{T_k^*(x)}^*, \dots, X_{\bar{T}_k^*(x)}^*)},$$

where $\zeta(z_0, \dots, z_m)$ (resp., $\zeta^*(z_0, \dots, z_m)$) is the measure of the trajectory (z_0, \dots, z_m) under the original measure P (resp., importance-sampling measure P^*) given $X_0 = z_0$. Note that $d\zeta(z_0, \dots, z_m) = \prod_{j=0}^{m-1} Q(z_j, z_{j+1})$. Now define $\Gamma_k^*(x) = \inf\{j > T_k^*(x) : X_j \in S_0\}$. For $x \in A$, let

$$\begin{aligned} \psi_k(x, x) &= I\left(X_{\bar{T}_k^*(x)}^* = x, \Gamma_k^*(x) \leq \bar{T}_k^*(x)\right) \\ &\quad + I\left(X_{\bar{T}_k^*(x)}^* \neq x\right), \end{aligned}$$

and for $x \neq y$, let

$$\psi_k(x, y) = -I\left(X_{\bar{T}_k^*(x)}^* = y, \Gamma_k^*(x) > \bar{T}_k^*(x)\right).$$

We now define $W^* = (W_k^* : k = 0, 1, 2, \dots)$ with $W_k^* = X_{T_k^*}^*$, and also define $L_k = d\zeta(X_{T_{k-1}^*}^*, \dots, X_{T_k^*}^*)/d\zeta^*(X_{T_{k-1}^*}^*, \dots, X_{T_k^*}^*)$, and $\Gamma_k^* = \inf\{j > T_k^* : X_j^* \in S_0\}$. For $x \in A$, define $H_n^*(x) =$

$\sum_{k=0}^{n_2^*-1} I(W_k^* = x)$, and let

$$\begin{aligned} M_n^*(x, x) &= \left(\sum_{k=0}^{n_2^*-1} I(W_k^* = x)\right)^{-1} \\ &\quad \times \left(\sum_{k=0}^{n_2^*-1} [I(W_k^* = x, W_{k+1}^* = x, \Gamma_k^* \leq T_{k+1}^*) \right. \\ &\quad \left. + I(W_k^* = x, W_{k+1}^* \neq x)] L_k\right) \\ &= \frac{1}{H_n^*(x)} \sum_{k=1}^{H_n^*(x)} \psi_k(x, x) L_k(x), \end{aligned}$$

and for $x \neq y$, let

$$\begin{aligned} M_n^*(x, y) &= \frac{\sum_{k=0}^{n_2^*-1} -I(W_k^* = x, W_{k+1}^* = y, \Gamma_k^* > T_{k+1}^*) L_k}{\sum_{k=0}^{n_2^*-1} I(W_k^* = x)} \\ &= \frac{1}{H_n^*(x)} \sum_{k=1}^{H_n^*(x)} \psi_k(x, y) L_k(x). \end{aligned}$$

Let $M_n^* = (M_n^*(x, y) : x, y \in A)$. Then we define our semi-regenerative estimator of η based on importance sampling to be

$$\eta_{n,\delta}^* = (M_n^*)^{-1} b'_n,$$

where $\eta_{n,\delta}^* = (\eta_{n,\delta}^*(x) : x \in A)$.

It is straightforward to show that

$$n^{1/2}(M_n^* - M) \xrightarrow{\mathcal{D}} N_1^*, \quad n^{1/2}(b'_n - b) \xrightarrow{\mathcal{D}} N_2', \quad (10)$$

where N_1^* and N_2' are normal random matrices with N_1^* and N_2' independent. One can then show the following central limit theorem.

Theorem 7. *Suppose that $E_x[B(x)^2] < \infty$ and $E_{x,T}^*[\psi(x, y)L(x)^2] < \infty$ for all $x, y \in A$. Then if Assumptions 1 and 3 hold,*

$$n^{1/2}(\eta_{n,\delta}^* - \eta) \xrightarrow{\mathcal{D}} (I - K)^{-1} N_1^* \eta + (I - K)^{-1} N_2'$$

as $n \rightarrow \infty$, where N_1^* and N_2' are defined in (10). In particular, for each $k = 1, 2, \dots, d$,

$$n^{1/2}(\eta_{n,\delta}^*(x_k) - \eta(x_k)) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \sigma_{k,*}^2)$$

as $n \rightarrow \infty$, where

$$\sigma_{k,*}^2 = \sum_{i=1}^d J(x_k, x_i)^2 \times \left[\frac{v_i}{\delta v(x_i)} + \sum_{j=1}^d \sum_{l=1}^d \frac{\eta(x_j)\eta(x_l)\Delta_i^*(x_j, x_l)}{(1-\delta)\rho(x_i)} \right],$$

$$v_i = \text{Var}(B(x_i)) \quad \text{and} \quad \Delta_i^*(x_j, x_l) = \text{Cov}_*(\psi(x_i, x_j)L(x_i), \psi(x_i, x_l)L(x_i)).$$

In the previous discussion, using the representation in (9), we employed importance sampling to estimate the matrix M . Alternatively, we can write $K(x, y) = E_x^*[\phi(x, y)L(x)]$, which suggests applying importance sampling to estimate the entries in the matrix K instead. Let K_n^* be the resulting importance-sampling estimator of K , and we then compute the estimator $(I - K_n^*)^{-1}b'_n$ of η . However, in our experiments, implementing this approach led to significantly worse results than the method based on (9).

Although each $B(x)$, $x \in A$, may have a small CV, it may be the case that for certain systems, some states $x \in A$ may be visited only rarely under the original measure P , thus resulting in few (if any) samples of $B(x)$. Hence, we may need to apply importance sampling to ensure that a sufficient number of samples of $B(x)$ are obtained. The resulting method is straightforward, and we omit its development.

4.3 Importance Sampling and Stratification

Another approach is to use a type of stratification in combination with importance sampling. To do this, we fix a computation budget n and constants p_i and q_i , $i = 1, 2, \dots, d$, with $p_i, q_i > 0$ and $\sum_{i=1}^d (p_i + q_i) = 1$. We will use each state $x_i \in A$ as an initial state for trajectories, where we sample $\lfloor p_i n \rfloor$ trajectories using the original measure P_{x_i} , and we sample $\lfloor q_i n \rfloor$ trajectories using the importance-sampling measure $P_{x_i, T}^*$. Specifically, for each $i = 1, 2, \dots, d$, let $\hat{B}_k(x_i)$, $1 \leq k \leq \lfloor p_i n \rfloor$, be i.i.d. copies of $B(x_i)$ under measure P_{x_i} . Let

$$(\hat{\psi}_k(x_i, y), \hat{L}_k(x_i) : y \in A)$$

for $1 \leq k \leq \lfloor q_i n \rfloor$ be i.i.d. copies of

$$(\psi(x, y), L(x_i) : y \in A)$$

under measure $P_{x_i, T}^*$, independent of the $\hat{B}_k(x_i)$. Set

$$\hat{b}_n(x_i) = \frac{1}{\lfloor p_i n \rfloor} \sum_{k=1}^{\lfloor p_i n \rfloor} \hat{B}_k(x_i),$$

$$\hat{M}_n(x_i, y) = \frac{1}{\lfloor q_i n \rfloor} \sum_{k=1}^{\lfloor q_i n \rfloor} \hat{\psi}_k(x_i, y)\hat{L}_k(x_i),$$

for $1 \leq i \leq d$ and $y \in A$, and set $\hat{b}_n = (\hat{b}_n(x) : x \in A)$ and $\hat{M}_n = (\hat{M}_n(x, y) : x, y \in A)$. Finally, define the estimator of η to be

$$\hat{\eta}_n = (\hat{M}_n)^{-1}\hat{b}_n,$$

where $\hat{\eta}_n = (\hat{\eta}_n(x) : x \in A)$.

Theorem 8. *Suppose that $E_x[B(x)^2] < \infty$ and $E_{x, T}^*[\psi(x, y)L(x)^2] < \infty$ for all $x, y \in A$. If Assumptions 1 and 3 hold, then for each $k = 1, 2, \dots, d$,*

$$n^{1/2}(\hat{\eta}_n(x_k) - \eta(x_k)) \xrightarrow{\mathcal{D}} \mathcal{N}(0, \hat{\sigma}_k^2)$$

as $n \rightarrow \infty$, where

$$\hat{\sigma}_k^2 = \sum_{i=1}^d J(x_k, x_i)^2 \times \left[\frac{v_i}{p_i} + \sum_{j=1}^d \sum_{l=1}^d \frac{\eta(x_j)\eta(x_l)\Delta_i^*(x_j, x_l)}{q_i} \right],$$

$$v_i = \text{Var}(B(x_i)) \quad \text{and} \quad \Delta_i^*(x_j, x_l) = \text{Cov}_*(\psi(x_i, x_j)L(x_i), \psi(x_i, x_l)L(x_i)).$$

The estimator $\hat{\eta}_n$ is a type of stratified estimator, in which there are $2d$ strata. Corresponding to each $x_i \in A$ are two strata, one for standard simulation and one for importance sampling. Simulating a trajectory starting from x_i under standard simulation is effectively a sample from the stratum for x_i under standard simulation, and similarly for the importance-sampling measure. Calvin, Glynn, and Nakayama (2001) derive the optimal choice of the stratification weights p_i, q_i , $i = 1, 2, \dots, d$, when estimating a steady-state mean using the semi-regenerative method, and show how to estimate their values.

5 EMPIRICAL RESULTS

We now present preliminary empirical results from applying the techniques discussed in this paper. The model we consider is a reliability system consisting of C types of components, labelled $1, 2, \dots, C$, with each type having the same redundancy r . The operating components in

the system are subject to random failures, and the failure times are exponentially distributed, with failure rate λ_i for components of type i , $i = 1, 2, \dots, C$. Initially, there is one component of each type operational, and the remaining $r - 1$ of each type are spares in hot standby; i.e., if the operational component of type i fails, then a spare, if one is available, immediately takes its place. Failed components are fixed by a single repairman using random-order service with total exponential repair rate μ . In other words, suppose that k_i components of type i , $i = 1, 2, \dots, C$, are currently failed, and let $C_R = \{i : k_i \geq 1\}$. Then the repairman simultaneously repairs one component of each type $i \in C_R$, with the effort $\mu/|C_R|$ devoted to each.

The state space of the system is $S = \{(j_1, j_2, \dots, j_C) : 0 \leq j_i \leq r, i = 1, 2, \dots, C\}$, where for state (j_1, j_2, \dots, j_C) , there are j_i components of type i failed. We consider the embedded DTMC. Let $u = (u_1, u_2, \dots, u_C) \in S$ and $v = (v_1, v_2, \dots, v_C) \in S$ be two generic states, and let 0 denote the state $(0, 0, \dots, 0)$ in which all components are operational. We say that a transition (u, v) is a *failure* (resp., *repair*) *transition* if $u_i + 1 = v_i$ (resp., $u_i - 1 = v_i$) for some i and $u_l = v_l$ for $l \neq i$, and in this case, the failure (resp., repair) transition corresponds to a failure (resp., repair) of a component of type i .

The transition probability matrix Q has the following non-zero entries: $Q(0, v) = \lambda_i / (\sum_{l=1}^C \lambda_l)$ when $(0, v)$ is a failure transition corresponding to a failure of a component of type i ; $Q(u, v) = \lambda_i / (\mu + \sum_{l=1}^C \lambda_l)$ when $u \neq 0$ and (u, v) is a failure transition corresponding to a failure of a component of type i ; $Q(u, v) = \mu / (|C_R(u)|(\mu + \sum_{l=1}^C \lambda_l))$ when (u, v) is a repair transition, where $C_R(u) = \{i : u_i \geq 1\}$.

We assume that the system is operational if and only if there is at least one component of each type operational. We consider the estimation of the expected number of transitions of the embedded discrete-time Markov chain until system failure given some initial state, so $S_0 = \{(j_1, j_2, \dots, j_C) : j_i = r \text{ for some } i\}$ and the reward function $f(z) = 1$, $z \in S$.

In our experiments, we either used importance sampling (IS) or standard simulation (no IS). We implemented importance sampling using balanced failure biasing (Goyal et al. 1992, Shahabuddin 1994). To describe this method, we first define the *failure biasing parameter* θ , $0 < \theta < 1$. (Typically one chooses $0.5 \leq \theta \leq 0.9$ in practice, and in all of our experiments, we let $\theta = 0.5$.) Define the importance-sampling transition matrix Q^* with the following non-zero entries: $Q^*(u, v) = \theta / |C_F(u)|$ when $u \neq 0$ and (u, v) is a failure transition, where $C_F(u) = \{i : u_i < r\}$; $Q^*(u, v) = (1 - \theta) / |C_R(u)|$ when $u \neq 0$ and (u, v) is a repair transition; and $Q^*(0, v) = 1/C$ for failure transitions $(0, v)$. Under balanced failure biasing, within a simulated trajectory, use matrix Q^* to generate transitions until S_0 is hit, and then use Q to generate the rest of the trajectory.

In our experiments we considered three methods to estimate η : the regenerative method (RM), the semi-regenerative method (SR), and a combined method (CM). Used only to estimate $\eta(0)$, CM is implemented as follows. When there are C components, let s_i , $i = 1, 2, \dots, C$, be the state in which there is exactly one component of type i failed and all other components are operational; i.e., $s_i = (0, \dots, 0, 1, 0, \dots, 0) \in \mathfrak{R}^C$, where the 1 is in the i th place. Then we can express

$$\eta(0) = 1 + \sum_{i=1}^C Q(0, s_i) \eta(s_i). \quad (11)$$

When using CM, we estimate each $\eta(s_i)$, $i = 1, 2, \dots, C$, using SR with $A = \{s_1, s_2, \dots, s_C\}$. Then substitute the estimates of $\eta(s_i)$ in (11) to arrive at the CM estimator for $\eta(0)$. The $Q(0, s_i)$ are not estimated since they are known.

For each method we estimated the coefficient of variation (CV) of the resulting estimator by running 1000 independent replications, with each replication consisting of a total of 2×10^6 transitions of the DTMC. Also, we estimated the “numerator” independently of the “denominator,” even when using no importance sampling with the regenerative method. (Here, to simplify the discussion, we use the term “denominator” (resp., “numerator”) to generically denote either the denominator (resp., numerator) in (1) or M (resp., b) in the semi-regenerative representation (5) of η .) When using IS, we only estimated the “denominator” with IS, and the “numerator” was estimated using no IS. For a fixed computation budget, we allotted 75% of the budget to estimating the “denominator” and the rest to the “numerator.”

In our first set of experiments, we took the number of types of components to be $C = 2$, each with redundancy $r = 3$. The failure rates are $\lambda_1 = 0.001$ and $\lambda_2 = 0.002$, and the repair rate is $\mu = 1$. Table 1 contains the results for this parameter set. The first column in Table 1 gives the method used, where SR- d denotes the semi-regenerative method with d states in the set A . The second column gives the states in the set A , where for the regenerative method, $|A| = 1$ and the one state in A is the return state. The third column gives the estimated CV when estimating the value of $\eta(u)$, where u is the first state listed in A (except when using the combined method, in which case $u = 0$); e.g., in the last row, we are estimating η given that $X_0 = s_1 = (1, 0)$. When stratification is used, denoted Strat, we used stratification weights $p_1 = 0.1$, $p_2 = p_3 = 0.45$ when $A = \{(0, 0), (0, 1), (1, 0)\}$, and $p_1 = p_2 = 0.5$ when $A = \{(0, 1), (1, 0)\}$.

In our second set of experiments, we took the number of types of components to be $C = 3$, each with redundancy $r = 3$. The failure rates are $\lambda_1 = 0.001$, $\lambda_2 = 0.002$, and $\lambda_3 = 0.003$, and the repair rate is $\mu = 1$. Table 2 contains

Table 1: Results for Reliability System with $C = 2$ Components

Method	A	CV
RM, no IS	0	0.688
RM IS	0	0.00889
SR-3 IS	$0, s_1, s_2$	0.0184
SR-3 IS, Strat	$0, s_1, s_2$	0.0194
CM-2 IS	s_1, s_2	0.00810
RM, no IS	s_1	0.670
RM IS	s_1	0.0881
SR-2 IS	s_1, s_2	0.00810
SR-2 IS, Strat	s_1, s_2	0.00797

the results for this parameter set. When using stratification, we specified the weights as $p_1 = 0.1, p_2 = p_3 = p_4 = 0.3$ when $A = \{0, s_1, s_2, s_3\}$, and $p_1 = p_2 = p_3 = 1/3$ when $A = \{s_1, s_2, s_3\}$.

Table 2: Results for Reliability System with $C = 3$ Components

Method	A	CV
RM, no IS	0	0.711
RM IS	0	0.0151
SR-4 IS	$0, s_1, s_2, s_3$	0.0149
SR-4 IS, Strat	$0, s_1, s_2, s_3$	0.0142
CM-3 IS	s_1, s_2, s_3	0.0126
RM, no IS	s_1	0.673
RM IS	s_1	0.195
SR-3 IS	s_1, s_2, s_3	0.0126
SR-3 IS, Strat	s_1, s_2, s_3	0.0122

We now provide some analysis of the results in Tables 1 and 2. To simplify the discussion, we call Case 1 the setting in which we are estimating $\eta(0)$ with $C = 2$ (the top half of Table 1), Case 2 is for $\eta(s_1)$ with $C = 2$ (the bottom half of Table 1), Case 3 is for $\eta(0)$ with $C = 3$ (the top half of Table 2), and Case 4 is for $\eta(s_1)$ with $C = 3$ (the bottom half of Table 2). First, note that for all four cases, RM with no importance sampling leads to estimators of about the same quality (as measured by the CV). In all cases these estimators can be improved by applying IS with either RM or SR. If RM with IS is applied, then the estimator in Case 1 is the best, followed by (in ascending order of CV) Cases 3, 2, and 4. Now comparing the amount of improvement gained over RM (with IS) by applying SR (with IS), we see the same ordering of the cases. Actually, SR does worse than RM in Case 1, but in Case 4, SR does significantly better. Thus, it appears that when applying IS, if RM does well, then SR may not help (and may in fact be worse). But when

RM does not do well, SR can lead to improvements, and the benefit increases as the variability of the RM estimator increases. Also, for the cases in which SR is beneficial, stratification may lead to a slight improvement. (Note that we did not attempt to find optimal stratification weights.)

ACKNOWLEDGMENTS

The work of the first author was supported in part by the National Science Foundation under grants DMI-9500173 and DMI-9900117. The second author's research was partially supported by Army Research Office Grant DAAG55-97-1-0377-P0001 and National Science Foundation Grant DMS-9704732. The work of the third author was supported in part by the National Science Foundation under grants DMI-9624469 and DMI-9900117.

REFERENCES

Calvin, J. M., P. W. Glynn and M. K. Nakayama. 2001. The semi-regenerative method of simulation output analysis. Forthcoming.

Çınlar, E. 1975. *Introduction to Stochastic Processes*. Englewood Cliffs, N.J.: Prentice-Hall.

Crane, M. and D. L. Iglehart. 1975. Simulating stable stochastic systems, III: Regenerative processes and discrete-event simulations. *Operations Research* 23: 33-45.

Glynn, P. W. 1995. Importance sampling for Markov chains: Asymptotics for the variance. *Stochastic Models* 10: 701-717.

Glynn, P. W. and D. L. Iglehart. 1989. Importance sampling for stochastic systems. *Management Science* 35: 1367-1393.

Goyal, A., P. Shahabuddin, P. Heidelberger, V.F. Nicola and P.W. Glynn. 1992. A unified framework for simulating Markovian models of highly reliable systems. *IEEE Transactions on Computers* C-41: 36-51.

Heidelberger, P. 1995. Fast simulation of rare events in queueing and reliability models. *ACM Transactions on Modeling and Computer Simulation* 5: 43-85.

Shahabuddin, P. 1994. Importance sampling for the simulation of highly reliable Markovian systems. *Management Science* 40: 333-352.

Shedler, G. S. 1993. *Regenerative Stochastic Simulation*. San Diego: Academic Press.

AUTHOR BIOGRAPHIES

JAMES M. CALVIN is an associate professor in the Department of Computer Science at the New Jersey Institute of Technology. He received a Ph.D. in Operations Research from Stanford University. Besides simulation output anal-

ysis, his research interests include global optimization and probabilistic analysis of algorithms.

PETER W. GLYNN received his Ph.D. from Stanford University, after which he joined the faculty of the Department of Industrial Engineering at the University of Wisconsin-Madison. In 1987, he returned to Stanford, where he is the Thomas Ford Professor of Engineering in the Department of Management Science and Engineering. His research interests include discrete-event simulation, computational probability, queueing, and general theory for stochastic systems.

MARVIN K. NAKAYAMA is an associate professor in the Department of Computer Science at the New Jersey Institute of Technology. He received a Ph.D. in Operations Research from Stanford University. He is a recipient of a CAREER Award from the National Science Foundation. He is an associate editor for *ACM Transactions on Modeling and Computer Simulation* and for *Infornis Journal on Computing*. His research interests include applied probability, statistics, simulation-output analysis, variance-reduction techniques, and gradient estimation.