

IMPORTANCE SAMPLING FOR MARKOV CHAINS : COMPUTING VARIANCE AND DETERMINING OPTIMAL MEASURES

Indira Kuruganti
Stephen Strickland

Department of Systems Engineering
University of Virginia
Charlottesville, VA 22903

ABSTRACT

In this paper we describe several computational algorithms useful in studying importance sampling (IS) for Markov chains. Our algorithms compute optimal IS measures and evaluate the estimate variance for a given measure. As knowledge of the optimal IS measure implies knowledge of the quantity to be estimated, our algorithms produce this quantity as a by-product. Since effective IS measures must often closely approximate the optimal measure, the use of these algorithms for small problems may produce insights that lead to effective measures for larger problems of actual interest. We consider two classes of problems: hitting times and fixed-horizon costs.

1 INTRODUCTION

Most investigators who experiment with importance sampling quickly discover that the “holy grail” of variance reduction is highly elusive. One seems as apt to make things worse as to improve them. While theoretical results provide guidance in some problem domains (e.g. single station queues, highly reliable Markovian systems - see Heidelberger (1995) for a review), little support exists in other situations.

The few general results in importance sampling indicate that any IS measure with “nice” variance properties (e.g. bounded relative error) must closely approximate the optimal IS measure in many aspects (e.g. support, relative distribution over the sample space). Thus it seems reasonable to compute optimal IS measures for “small” problems in the hope that the insights may be applied to the larger problems of actual interest. In particular, computationally reasonable approximations of the optimal measure may have good variance behavior.

In this paper we describe several computational algorithms useful in studying importance sampling (IS) for Markov chains. Our algorithms compute optimal IS measures and evaluate the estimate variance

for a given measure. Since knowledge of the optimal IS measure implies knowledge of the quantity to be estimated, our algorithms produce this quantity as a by-product. Our algorithms employ dynamic-programming-like ideas and compute IS measures in a recursive form amenable to event-by-event simulation.

We consider two classes of problems: hitting times and fixed-horizon costs. Typical examples of hitting time problems include estimating the mean-time-to-overflow in finite capacity Markovian queueing systems, or the mean-time-to-failure in a system with exponential component lifetimes and repair times. Fixed-horizon costs occur when estimating inventory costs or queueing times over a fixed time period.

The paper is organized as follows. In Section 2 we review basic concepts of importance sampling and summarize some earlier results on the properties of the optimal (zero-variance) importance sampling scheme. In Section 3 we discuss the problem of estimating the first passage time to state F starting from state O . We present different approaches of finding the optimal dynamic IS scheme (Section 3.1) and develop a useful tool for assessing the variance of a constructed scheme without using simulation (Section 3.2). Next, in Section 4 we consider the problem of estimating the mean performance measure over a fixed (finite) time horizon. We discuss two different formulations of this problem - a multiplicative formulation in Section 4.1 and an additive formulation in Section 4.2. In each case, we present the optimal IS scheme and develop a tool for variance computation. We conclude in Section 5 with a mention of some heuristics for constructing “good” IS schemes suggested by the results of this paper.

2 BACKGROUND

Let X denote a stochastic process defined on the probability space $(\Omega, \mathcal{F}, \mathbf{P})$, and assume that we wish

to estimate the mean $\alpha \equiv \mathbf{E}[g(x)]$ via simulation. The **naive simulation approach** for estimating α involves generating M independent sample paths x_1, \dots, x_M and forming the estimator

$$\hat{\alpha} = \frac{1}{M} \sum_{i=1}^M g(x_i)$$

Since $\mathbf{E}_{\mathbf{P}}[g(x_i)] = \alpha$ the naive estimator is always unbiased.

To motivate the need for importance sampling, let us consider the problem of estimating a rare event probability α . The performance measure in this case is $g(x) = 1_A(x)$ (where $1_A(\cdot)$ is the indicator function of A). Here, the naive estimator has standard deviation $\sigma_{\hat{\alpha}} = \sqrt{\alpha(1-\alpha)/M}$. So, the relative error $\sigma_{\hat{\alpha}}/\alpha$ grows without bound as $\alpha \rightarrow 0$. Equivalently, as the event becomes rarer, the number of samples needed to estimate α to a fixed precision grows without bound. Thus, for many practical situations where simulation is warranted, direct simulation of the system is infeasible.

Importance sampling (IS) is based on the observation that

$$\begin{aligned} \mathbf{E}_{\mathbf{P}}[g(x)] &= \int g(x) d\mathbf{P}(x) \\ &= \int g(x) \frac{d\mathbf{P}(x)}{d\tilde{\mathbf{P}}(x)} d\tilde{\mathbf{P}}(x) \\ &= \mathbf{E}_{\tilde{\mathbf{P}}}[g(x)\mathcal{L}(x)] \end{aligned} \tag{1}$$

where $\tilde{\mathbf{P}}$ is any measure absolutely continuous w.r.t. $g(\cdot)\mathbf{P}$ and $\mathcal{L}(x) = d\mathbf{P}(x)/d\tilde{\mathbf{P}}(x)$ is the Radon-Nikodym derivative of \mathbf{P} w.r.t. $\tilde{\mathbf{P}}$ (often called the likelihood ratio) Glynn and Iglehart (1989). Absolute continuity requires that $\tilde{\mathbf{P}}(x) > 0$ whenever $g(x)\mathbf{P}(x) > 0$.

From (1) it is clear that we can form a new (unbiased) estimator for α , namely,

$$\tilde{\alpha} = \frac{1}{M} \sum_{i=1}^M g(x_i)\mathcal{L}(x_i)$$

where the x_i are sampled from the measure $\tilde{\mathbf{P}}$. The variance of this estimator is given by

$$\sigma_{\tilde{\alpha}}^2 = \frac{\mathbf{E}_{\tilde{\mathbf{P}}}[(g(x)\mathcal{L}(x))^2] - \alpha^2}{M} = \frac{m_{\tilde{\mathbf{P}}}^2 - \alpha^2}{M}. \tag{2}$$

Here $m_{\tilde{\mathbf{P}}}^2$ represents the second moment of the estimator $g(x_i)\mathcal{L}(x_i)$ corresponding to a single sample. If $\tilde{\mathbf{P}}$ is chosen properly, then $\tilde{\alpha}$ will have significantly lower variance than $\hat{\alpha}$ - often orders of magnitude lower. However, a poor choice of $\tilde{\mathbf{P}}$ can lead to increased (even infinite) variance - a fact most recently

illustrated in a paper by Andradóttir et. al. (1995). Thus, choosing $\tilde{\mathbf{P}}$ properly is the key issue in employing importance sampling.

Since the goal is to minimize the variance of $\tilde{\alpha}$, we also make the following

DEFINITION 1 An **optimal IS measure** is one which produces a zero-variance estimate of α .

Note that if

$$d\tilde{\mathbf{P}}(x) = g(x) \cdot d\mathbf{P}(x)/\mathbf{E}_{\mathbf{P}}[g(x)] =: d\mathbf{P}^*, \tag{3}$$

then

$$g(x) \cdot \mathcal{L}(x) = \mathbf{E}_{\mathbf{P}}[g(x)]$$

and the variance is zero. Thus, an optimal IS measure always exists. Note that \mathbf{P}^* assigns probability proportional to a path's contribution to α , thereby giving rise to the name importance sampling.

3 HITTING TIMES

Consider a continuous-time Markov chain defined on a finite state space S . Given an initial state O , we wish to estimate the hitting time τ_F for a set of states $F \subset S$, $O \notin F$ under the assumption that probability of reaching F from O is very small. Since we are only interested in the *first passage* to F , we will model F as a single absorbing state. Such a model may be used for example to determine the mean time to system failure in a highly reliable system or the mean time to buffer overflow in a queuing system.

If τ_O is the first passage time to O , and $\alpha = \mathbf{P}[\tau_F < \tau_O]$, then for the regenerative process considered here, $\tau_F = \mathbf{E}_{\mathbf{P}}[\min(\tau_O, \tau_F)]/\alpha$ (see Bratley, Fox, and Schrage (1987) and Heidelberger (1995) for details). This result can be heuristically explained as follows. Since visits to F from O are rare, each sample path which starts in O and eventually ends in F may visit O multiple times in the intermediate epochs. Suppose in an arbitrary sample path, there are M "cycles" (i.e., sections of the sample path beginning in state O and ending with first passage into either state O or F). Let the time duration of the i -th cycle be denoted by T_i . Then, since the system probabilistically regenerates itself at each re-visit to O , $\{T_1, T_2, \dots, T_M\}$ are i.i.d. variables. Thus, the time until failure is the sum of a random number (M) of i.i.d. random variables and by Wald's lemma we have,

$$\tau_F = \mathbf{E}_{\mathbf{P}}\left[\sum_{i=1}^M T_i\right] = \mathbf{E}_{\mathbf{P}}[T_i]\mathbf{E}_{\mathbf{P}}[M].$$

The ratio formula for τ_F follows from the fact that $\mathbf{E}_{\mathbf{P}}[M] = 1/\alpha$ and since F is reached very rarely, $\min(\tau_O, \tau_F) \approx \tau_O$ most of the time. Hence, we can write,

$$\mathbf{E}_{\mathbf{P}}[T_i] = \mathbf{E}_{\mathbf{P}}[\min(\tau_O, \tau_F)] \approx \mathbf{E}_{\mathbf{P}}[\tau_O].$$

Now, estimating τ_O involves only a few state transitions and moreover, most cycles end in state O in a typical simulation. So we can easily estimate $\mathbf{E}_{\mathbf{P}}[\min(\tau_O, \tau_F)]$ using the naive simulation approach.

Thus the problem of estimating τ_F essentially reduces to that of estimating the rare event probability $\alpha \equiv \mathbf{P}[A] = \mathbf{E}_{\mathbf{P}}[1_A(x)]$. Here, A is the set of sample paths that begin in O and end in F without visiting O along the way. Since S is finite, A is either finite or at most countably infinite.

It is clear that α does not depend on the time spent in each state; rather it depends on the particular sequence of states visited in a given simulation run. Hence, we can replace the continuous time Markov process by the corresponding embedded discrete time Markov chain (DTMC) $\{X_t, t = 0, 1, 2, \dots\}$ with state space S and one-step transition probability matrix $P = [p_{s,t}]_{m \times m}$. In the remaining sections on this problem we will work with this DTMC.

The optimal IS scheme is given by:

$$\mathbf{P}^*(x) = \begin{cases} \frac{\mathbf{P}(x)}{\alpha} & \text{for } x \in A \\ 0 & \text{otherwise.} \end{cases}$$

A Markov chain construction of \mathbf{P}^* is given in the following theorem:

THEOREM 1 *Suppose the system behaves as a DTMC with state space S and transition matrix P . Then the probability assignment of (3) corresponds to a DTMC whose transition probability matrix P^* has elements:*

$$p^*(s, t) = \frac{p(s, t)\gamma(t)}{\gamma(s)} \quad (4)$$

where $\gamma(s)$ denotes the probability that the system reaches F starting from state s without ever returning to the initial state O .

For proof of this theorem see Kuruganti and Strickland (1995).

Thus if $\gamma(s)$ for each state s in S be known then we can compute the optimal scheme using (4). Since $\gamma(O) = \alpha$, this result is not useful for importance sampling; however, it shows that \mathbf{P}^* can be realized by a Markov chain.

3.1 Recursive Calculation of the Optimal IS Measure

Observe that

$$\gamma(s) = \lim_{N \rightarrow \infty} \gamma_{\leq N}(s), s \in S - F$$

where $\gamma_{\leq N}(s)$ denotes the probability that the system first reaches F starting from state s in N steps or less without ever returning to the initial state O .

Let $\gamma_N(s)$ denote the corresponding probability of hitting F in exactly N steps without returning to O . Clearly, $\gamma_1(s) \equiv p(s, F)$ and $\gamma_N(F) = 1$ for $N = 1, 2, \dots$. Set

$$q(s, t) = \begin{cases} 0 & t = O \text{ or } s = F, t \neq F \\ 1 & s = t = F \\ p(s, t) & \text{otherwise} \end{cases}$$

and let Q be the matrix with $q(s, t)$ as its elements.

Now, if the system is in state $s \in S - F$, then:

$$\begin{aligned} \gamma_{\leq N}(s) &= q(s, F) + \sum_{t \notin F} q(s, t)\gamma_{\leq N-1}(t) \\ &= \sum_t q(s, t)\gamma_{\leq N-1}(t). \end{aligned} \quad (5)$$

As in dynamic programming methods (e.g., see Bertsekas (1987)), we can regard $\gamma_{\leq N}(s)$ as the ‘‘cost-to-go’’ at the N -th stage. (Note that we are simply accumulating costs without the optimization at each stage that occurs in dynamic programming.)

Suppose $\Gamma_{\leq N} = [\gamma_{\leq N}]_{m \times 1}$ and $\Gamma_0 = [1_F(s)]_{m \times 1}$. Then, equation (5) becomes

$$\Gamma_{\leq N} = Q\Gamma_{\leq N-1} = Q^N\Gamma_0$$

Assuming that the expected hitting time is bounded, we have $\gamma_{\leq N} \rightarrow \gamma$ as $N \rightarrow \infty$ (see Kemeny and Snell (1976) for a detailed explanation). So if $\Gamma = [\gamma(s)]_{m \times 1}$

$$\Gamma = Q\Gamma$$

which can be solved by a variety of methods. Γ can also be approximated accurately by $Q^N\Gamma_0$ for large N at a computational cost of $O(\log N)$.

3.2 Variance Computation

The recursive approach above can be adapted to compute the estimate variance for any IS scheme characterized by a transition matrix \tilde{P} . From (2), this amounts to computing

$$\mathbf{E}_{\tilde{\mathbf{P}}} [(1_A(x)\mathcal{L}(x))^2] = \mathbf{E}_{\tilde{\mathbf{P}}} [(1_A(x)\mathbf{P}(x)/\tilde{\mathbf{P}}(x))].$$

Let $\psi_N(s)$ be the contribution to $m_{\mathbb{P}}^2$ of all paths starting from state s in $S - F$ which hit F in exactly N steps. And let $\psi_{\leq N}(s) = \sum_1^N \psi_j(s)$. As before, let \tilde{Q} be a matrix whose elements are given by

$$\tilde{q}(s, t) = \begin{cases} 0 & t = O \text{ or } s = F, t \neq F \\ 1 & s = t = F \\ \tilde{p}(s, t) & \text{otherwise} \end{cases}$$

Then

$$\begin{aligned} \psi_{\leq N}(s) &= \frac{q(s, F)^2}{\tilde{q}(s, F)} + \sum_{t \notin F} \frac{q(s, t)^2}{\tilde{q}(s, t)} \psi_{\leq N-1}(t) \\ &= \sum_t \frac{q(s, t)^2}{\tilde{q}(s, t)} \psi_{\leq N-1}(t). \end{aligned} \quad (6)$$

The desired second moment is given by :

$$m_{\mathbb{P}}^2 = \psi(O) = \lim_{N \rightarrow \infty} \psi_{\leq N}(O).$$

Letting Λ be the matrix with elements

$$l(s, t) = \begin{cases} \frac{q(s, t)^2}{q(s, t)} & \text{if } q(s, t) > 0 \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

and $\Psi = [\psi]_{m \times 1}$ we can write (6) as

$$\psi_{\leq N} = \Lambda \psi_{\leq N-1} = \Lambda^N \psi_0.$$

where $\psi_0(s) = 1_F(s)$. If the IS scheme has finite variance, so $\psi_{\leq N} \rightarrow \psi$ as $N \rightarrow \infty$, then $\Psi = [\psi(s)]_{m \times 1}$ is given by

$$\Psi = \Lambda \Psi.$$

Note though that this tool does not address the effect of the IS measure on the computational cost of each sample. We ignore this issue here because in most practical situations of interest, with an appropriate choice for $\tilde{\mathbb{P}}$ the reduction (relative to the naive approach) in number of samples needed to estimate α far exceeds the increased cost of generating each sample under importance sampling (see Sadowsky (1993)).

4 FIXED HORIZON COST FUNCTIONS

Let τ be an arbitrary (but fixed) number of time epochs and $Y = g(X_0, X_1, \dots, X_\tau)$ some performance measure of interest. Suppose we wish to estimate the mean $\alpha_\tau = E_P[Y]$. Let $\tilde{P} = [\tilde{p}(s, t)]_{m \times m}$ be the one-step transition probability matrix used during importance sampling and x_i the state visited at the i -th epoch. For simplicity, we assume that x_0 is fixed under both P and \tilde{P} i.e., the system always

starts in the same initial state. (For a more general treatment we can include the initial probability distribution vectors used under the original and importance sampling measures respectively).

From (3) the optimal (zero-variance) importance sampling scheme for this problem is given by :

$$\mathbf{P}^*(x_0, x_1, \dots, x_\tau) = \frac{\mathbf{P}(x_0, x_1, \dots, x_\tau) g(x_0, x_1, \dots, x_\tau)}{\alpha_\tau} \quad (8)$$

If the function $g(X_0, X_1, \dots, X_n)$ can be expressed in terms of a recursive equation, then we can exploit some of the techniques developed in the first segment of the paper to find the optimal IS scheme and compute the IS variance. In particular, we show how this may be accomplished for two formulations which cover many applications of practical interest :

Multiplicative:

$$g(X_0, X_1, \dots, X_\tau) = \prod_{i=0}^{\tau} g(X_i)$$

Additive:

$$g(X_0, X_1, \dots, X_\tau) = \sum_{i=0}^{\tau} g(X_i).$$

Let $G = [g(s)]_{m \times 1}$. In both cases we assume that the performance measure of interest is a positive quantity i.e., $g(s) > 0$ for every state s in the state space S . Since the multiplicative form is easier to work with we first examine this in the following section.

4.1 Multiplicative Cost

In this formulation, the mean performance measure α_τ is given by :

$$\alpha_\tau = \sum_{x_1, x_2, \dots, x_\tau} \prod_{i=0}^{\tau} g(x_i) \prod_{i=0}^{\tau-1} p(x_i, x_{i+1})$$

which can be calculated recursively as explained below. Define

$$\begin{aligned} h_j(x_j) &= \sum_{x_1, \dots, x_{j-1}} \prod_{i=0}^j g(x_i) \prod_{i=0}^{j-1} p(x_i, x_{i+1}) \\ &= \sum_{x_{j-1}} h_{j-1}(x_{j-1}) p(x_{j-1}, x_j) g(x_j) \end{aligned} \quad (9)$$

with $h_0(x_0) = g(x_0) \equiv \alpha_0$ and

$$\alpha_j = \sum_{x_j} h_j(x_j), \quad j = 1, 2, \dots, \tau \quad (10)$$

(recall that we assume a fixed starting state x_0). In matrix form, the recursive equation (9) can be written as :

$$H_j = DP^T H_{j-1}, \quad j = 1, 2, \dots, \tau \quad (11)$$

where D is a $m \times m$ diagonal matrix with elements $g(s)$ along the principal diagonal and P^T is the transpose of P . Here $H_j = [h(s)]_{m \times 1}$ is a vector that is recursively accumulated.

4.1.1 Optimal Scheme

Comparing equations (9) and (10) with (5) we note that in this dynamic programming formulation, α_j behaves as the “cost-to-reach” state x_j , the state visited at the j -th epoch. So, by analogy with (4) we hypothesize the following form for the elements of the Markovian optimal IS transition matrix P^* :

$$p^*(x_j, x_{j+1}) = \frac{\alpha_j p(x_j, x_{j+1}) g(x_{j+1})}{\alpha_{j+1}} \quad (12)$$

where $j = 0, 1, \dots, \tau - 1$. Note that non-negativity of $p^*(x_j, x_{j+1})$ is assured by the positive form assumed for the function $g(\cdot)$. Moreover, equations (9) and (10) ensure that every row of P^* sums to one. Since (12) expresses the optimal IS transition probabilities in terms of the known simulation inputs P and G , the optimal IS scheme can be computed and actually implemented in a simulation.

That this is indeed the optimal IS scheme can be verified as follows. Let (x_0, \dots, x_τ) be a particular sample path generated under P^* . Now:

$$\begin{aligned} & \mathbb{P}^*(x_0, x_1, \dots, x_\tau) \\ &= \prod_{i=0}^{\tau-1} p^*(x_i, x_{i+1}) \\ &= \frac{\alpha_0 p(x_0, x_1) g(x_1)}{\alpha_1} \frac{\alpha_1 p(x_1, x_2) g(x_2)}{\alpha_2} \\ & \quad \dots \frac{\alpha_{\tau-1} p(x_{\tau-1}, x_\tau) g(x_\tau)}{\alpha_\tau} \\ &= \frac{\prod_{i=0}^{\tau-1} g(x_i) \prod_{i=0}^{\tau-1} p(x_i, x_{i+1})}{\alpha_\tau}, \text{ and} \\ & \mathbb{P}(x_0, x_1, \dots, x_\tau) = \prod_{i=0}^{\tau-1} p(x_i, x_{i+1}). \end{aligned}$$

So, under \mathbb{P}^* each sample leads to the following I.S. estimate for α_τ :

$$\begin{aligned} & g(X_0, X_1, \dots, X_\tau) \mathcal{L}(X_0, X_1, \dots, X_\tau) \\ &= \left(\prod_{i=0}^{\tau-1} g(x_i) \right) \frac{\mathbb{P}^*(x_0, x_1, \dots, x_\tau)}{\mathbb{P}^*(x_0, x_1, \dots, x_\tau)} = \alpha_\tau. \end{aligned}$$

which is the (constant) true value. Hence the variance of the IS estimator is zero i.e., \mathbb{P}^* is the optimal IS scheme.

4.1.2 Variance Computation

As in Section 3.2 we can compute the variance of an arbitrary IS scheme for this problem using dynamic programming and consequently develop a matrix-based formulation.

Equation (11) suggests a matrix-based method for finding $\alpha_\tau \equiv h_0(x_0)$. So, using the same arguments as in Section 4.1 we can derive the recursive matrix equation:

$$V_j = D^2 \Lambda^T V_{j-1}, \quad j = 1, 2, \dots, \tau$$

where $V_0 \equiv G_2 \equiv [g(s)^2]_{m \times 1}$, $m_{\mathbb{P}}^2 \equiv \sum_{x_\tau} v_\tau(x_\tau)$ and Λ^T is the transpose of the matrix Λ defined in (7). The vector $V_j = [v_j(x_j)]_{m \times 1}$ can be recursively computed.

Using induction, it can be shown that:

$$\begin{aligned} V_j &= [D^2 \Lambda^T]^j V_0, \quad j = 1, 2, \dots, \tau \\ \Rightarrow V_\tau &= [D^2 \Lambda^T]^\tau G_2. \end{aligned}$$

So, the variance of an arbitrary (Markovian) IS scheme for this problem can be found using :

$$\bar{\sigma}^2 = \sum_{x_\tau} v_\tau(x_\tau) - \alpha_\tau^2. \quad (13)$$

4.2 Additive Cost

Now assume that:

$$\begin{aligned} \alpha_\tau \equiv E_P[Y] &= \\ & \sum_{x_1, \dots, x_\tau} \left(\sum_{i=0}^{\tau-1} g(x_i) \right) \prod_{i=0}^{\tau-1} p(x_i, x_{i+1}). \end{aligned}$$

Typical examples include the net payoff from a fixed number of successive gambling trials and the total waiting time experienced by a single customer at different stations in a queueing network. This representation can also be used when the performance measure of interest is the number of times a particular state (or set of states) F is visited in a simulation run of fixed length. In this case $g(x_i) = 1_F(x_i)$. (Note that this last example does not satisfy our earlier assumption of strictly positive values for $g(x_i)$.)

4.2.1 Optimal Scheme

Equation (8) and the fact that

$$\alpha_\tau \equiv E_{\mathbb{P}} \left[\sum_{i=0}^{\tau} g(x_i) \right] = \sum_{i=0}^{\tau} E_{\mathbb{P}} [g(x_i)]$$

imply that

$$\mathbb{P}^*(x_0, x_1, \dots, x_{j+1}) = \frac{\left\{ \sum_{i=0}^{j+1} g(x_i) \right\} \prod_{i=0}^{j+1} p(x_i, x_{i+1})}{\sum_{i=0}^{j+1} E_{\mathbb{P}} [g(x_i)]}.$$

But by the Markov property, this probability is given by:

$$\mathbb{P}^*(x_0, x_1, \dots, x_{j+1}) = \mathbb{P}^*(x_0, x_1, \dots, x_j) P^*(x_j, x_{j+1}).$$

We therefore conclude that

$$\begin{aligned} p^*(x_j, x_{j+1}) &= \frac{\mathbb{P}^*(x_0, x_1, \dots, x_{j+1})}{\mathbb{P}^*(x_0, x_1, \dots, x_j)} \\ &= \left(\frac{\sum_{i=0}^{j+1} g(x_i)}{\sum_{i=0}^j g(x_i)} \right) \left(\frac{\prod_{i=0}^j p(x_i, x_{i+1})}{\prod_{i=0}^{j-1} p(x_i, x_{i+1})} \right) \\ &= \left(\frac{\sum_{i=0}^j E_{\mathbb{P}} [g(x_i)]}{\sum_{i=0}^{j+1} E_{\mathbb{P}} [g(x_i)]} \right) \\ &= p(x_j, x_{j+1}) \\ &= \left(\frac{\sum_{i=0}^j E_{\mathbb{P}} [g(x_i)]}{\sum_{i=0}^{j+1} E_{\mathbb{P}} [g(x_i)]} \right) \left(\frac{\sum_{i=0}^{j+1} g(x_i)}{\sum_{i=0}^j g(x_i)} \right) \\ &= p(x_j, x_{j+1}) \left(\frac{\alpha_j}{\alpha_{j+1}} \right) \left(1 + \frac{g(x_{j+1})}{\sum_{i=0}^j g(x_i)} \right). \end{aligned}$$

Since all the components of the above expression can be computed off-line or along the sample path, this $P^* = [p^*(s, t)]_{m \times m}$ can be used for simulation. But note that the transition probabilities at epoch i depend on the sample path history up to i . Thus it appears that the optimal IS measure cannot be realized by a Markov chain. Also, observe that as $j \rightarrow \infty$,

$$\frac{\alpha_j}{\alpha_{j+1}} \rightarrow 1 \text{ and } \frac{g(x_{j+1})}{\sum_{i=0}^j g(x_i)} \rightarrow 0,$$

so

$$p^*(x_j, x_{j+1}) \rightarrow p(x_j, x_{j+1}).$$

Thus, as the sample path length increases, the optimal IS transition matrix P^* approaches the original matrix P . This is consistent with the result derived

in Andradóttir et. al (1995) that for fixed time-horizon problems, for best results, the IS transition matrix \tilde{P} must approach the original transition matrix P as the sample path length (i.e., time-horizon) increases.

4.2.2 Variance Computation

We begin by establishing that α_τ can be recursively computed. So, we define

$$\begin{aligned} h_j(x_j) &= \sum_{x_1, \dots, x_{j-1}} \sum_{i=0}^j g(x_i) \prod_{i=0}^{j-1} p(x_i, x_{i+1}) \\ &= \sum_{x_{j-1}} h_{j-1}(x_{j-1}) p(x_{j-1}, x_j) \\ &\quad + P^j(x_0, x_j) g(x_j) \end{aligned}$$

with $h_0(s) = g(s) 1_{x_0}(s)$ and $\alpha_j = \sum_{x_j} h_j(x_j)$. Note that $P^j(x_0, x_j)$, the total probability of going from x_0 to x_j in exactly j steps, is an element of the row for x_0 in the matrix P^j . So, in matrix form the above recursive equation is:

$$H_j = P^T H_{j-1} + D P^j I_0, \quad j = 1, 2, \dots, \tau \quad (14)$$

where D is a diagonal $m \times m$ matrix with elements $g(s)$ along the principal diagonal and $I_0 = [1_{x_0}(s)]_{m \times 1}$. Thus the vector $H_j = [h_j(x_j)]_{m \times 1}$ (and hence α_j) can be recursively computed using (14).

Similarly, we can show that the I.S. variance can be recursively computed by defining

$$\begin{aligned} v_j(x_j) &= \sum_{x_1, \dots, x_{j-1}} \sum_{i=0}^j g(x_i)^2 \prod_{i=0}^{j-1} \frac{p(x_i, x_{i+1})^2}{\tilde{p}(x_i, x_{i+1})} \\ &= \sum_{x_{j-1}} v_{j-1}(x_{j-1}) l(x_{j-1}, x_j) \\ &\quad + \Lambda^j(x_0, x_j) g^2(x_j) \\ &\quad + 2g(x_j) \tilde{h}_j(x_j) \end{aligned}$$

where the last term itself has a recursive structure as indicated below:

$$\begin{aligned} \tilde{h}_j(x_j) &= \sum_{x_1, \dots, x_{j-1}} \left(\sum_{i=0}^{j-1} g(x_i) \right) \prod_{i=0}^{j-1} \frac{p(x_i, x_{i+1})^2}{\tilde{p}(x_i, x_{i+1})} \\ &= \sum_{x_{j-1}} \Lambda^{j-1}(x_0, x_{j-1}) g(x_{j-1}) l(x_{j-1}, x_j) \\ &\quad + \sum_{x_{j-1}} \tilde{h}_{j-1}(x_{j-1}). \end{aligned}$$

Writing these results in matrix form we get:

$$V_j = \Lambda^T V_{j-1} + \Lambda^j G_2 + 2D\tilde{H}_j \\ \text{for } j = 1, 2, \dots, \tau$$

where

$$\tilde{H}_j = \Lambda^T \left(\tilde{H}_{j-1} + D(\Lambda^{j-1})^T I_0 \right) \\ \text{So, } \tilde{h}_j = \Lambda^{Tj} \tilde{h}_0 + \sum_{i=1}^{j-1} \Lambda^{Ti} D(\Lambda^{j-i})^T I_0.$$

Here $\tilde{h}_0 = [g(s)^2 1_{x_0}(s)]_{m \times 1}$, $V_0 \equiv [g(s)^2 1_{x_0}(s)]_{m \times 1}$, $G_2 \equiv [g(s)^2]_{m \times 1}$, and $m_{\mathbf{p}}^2 \equiv \sum_{x_\tau} v_\tau(x_\tau)$. So, the variance of an arbitrary (Markovian) IS scheme for this problem can be once again computed using Equation (13).

5 CONCLUSIONS

In this paper we have presented several recursive algorithms for constructing optimal IS measures and computing the variance of arbitrary IS measures. (A different non-recursive approach for finding the optimal IS measure for the hitting time problem was presented in Kuruganti and Strickland (1995).) While these techniques are limited to relatively "small" problems, they may prove useful for developing effective schemes for large problems, either through approximation of the optimal scheme or through exploratory evaluation of trial schemes for small prototypes.

For example, when estimating the first passage time to F , using approximations of the methods developed in Section 3.1 we can generate guesstimates for $\gamma(s)$ and substitute these in (4) to construct a sub-optimal yet computationally attractive IS measure that is dynamic (state-dependent). Another possible approach is to replace the original model by a simpler model (with fewer states and/or simpler transition structure) for which $\gamma(s)$ can be computed easily. Then, we could use these results to approximate $\gamma(s)$ for the original more complex model.

Approximations of the recursive method of computing $\gamma(s)$ (and thence the optimal change of measure) may also lead to heuristic strategies for constructing sub-optimal but computationally attractive IS simulation schemes with desirable asymptotic properties. One strategy in this context involves considering at each stage of the infinite-horizon dynamic programming algorithm the contributions from only the dominant most likely paths to failure. In an earlier paper, it was established by Strickland (1993)

that the resultant IS scheme behaves like the optimal IS scheme in the limit as the rare event of interest has vanishingly small probability. An alternative technique is to use a myopic approach i.e., restrict attention to only sample paths of length N or less (where N is a small positive integer) in calculating $\gamma(s)$ via equation (5). Similar approximation heuristics may also be investigated for the fixed-horizon problem considered in this paper.

These techniques potentially present a simple and intuitive alternative to the approaches inspired by large deviations theory that have thus far been discussed in the literature - see for example Cottrell et. al (1983), Parekh and Walrand (1989) and Sadowsky (1993).

ACKNOWLEDGMENTS

The work of the authors was supported in part by the National Science Foundation under contract No. ECS-9209902.

REFERENCES

- Andradóttir, S., Heyman, D.P., and Teunis, J.Ott, 1995, "On the Choice of Alternative Measures in Importance Sampling with Markov Chains," *Operations Research*, Vol. 43, No. 3, May-June 1995.
- Bertsekas, D.P., 1987, *Dynamic Programming : Deterministic and Stochastic Models*, Prentice Hall, New Jersey.
- Bratley, P., Fox, B., and Schrage, L., 1987, *A Guide to Simulation*, 2nd Ed., Springer-Verlag, New York.
- Cottrell, M., Fort, J. and Malgouyres, G., 1983, "Large Deviations and Rare Events in the Study of Stochastic Algorithms," *IEEE Transactions on Automatic Control*, Vol. AC-28, No. 9, pp. 907-920.
- Glynn, P., and Iglehart, D., 1989, "Importance Sampling for Stochastic Simulation," *Management Science*, Vol. 35, No. 11, pp. 1367-1392.
- Heidelberger, P., 1995, "Fast Simulation of Rare Events in Queuing and Reliability Models," *ACM Transactions on Modeling and Computer Simulation*, Vol. 5, No. 1, pp 43-85.
- Kemeny, J.G., and Snell, J.L., 1976, *Finite Markov Chains*, Springer-Verlag, New York.
- Kuruganti, I., and Strickland, S.G., 1995, "Optimal Importance Sampling For Markovian Systems," *Proceedings of the 1995 IEEE Conference on Systems, Man and Cybernetics*.

- Parekh, S., and Walrand, J., 1989, "A Quick Simulation Method for Excessive Backlogs in Networks of Queues," *IEEE Transactions on Automatic Control*, Vol. 34, No.1, pp. 54-66.
- Sadowsky, J., S., 1993, "On the Optimality and Stability of Exponential Twisting in Monte Carlo Estimation," *IEEE Transactions on Information Theory*, Vol. 39, pp. 119-128.
- Strickland, S., 1993, "Optimal Importance Sampling for Quick Simulation of Highly Reliable Markovian Systems," *Proceedings of the 1993 Winter Simulation Conference*, pp 437-444.

AUTHOR BIOGRAPHIES

INDIRA KURUGANTI is a doctoral candidate in the Department of Systems Engineering at the University of Virginia. Her research interests include Markov processes and rare event simulation.

STEPHEN G. STRICKLAND has been an Assistant Professor in the Department of Systems Engineering at the University of Virginia since 1990. His research interests include gradient/sensitivity estimation and rare event simulation.