

ESTIMATING THE PROBABILITY OF A RARE EVENT OVER A FINITE TIME HORIZON

Pieter-Tjerk de Boer

Faculty for Electrical Engineering,
Mathematics and Computer Science
University of Twente, P.O. Box 217
7500 AE Enschede, THE NETHERLANDS

Pierre L'Ecuyer

IRISA, Rennes
and DIRO, Université de Montreal
C.P. 6128, Succ. Centre-Ville
Montréal (Québec), H3C 3J7, CANADA

Gerardo Rubino, and Bruno Tuffin

IRISA/INRIA, Campus de Beaulieu
35042 Rennes Cedex, FRANCE

ABSTRACT

We study an approximation for the zero-variance change of measure to estimate the probability of a rare event in a continuous-time Markov chain. The rare event occurs when the chain reaches a given set of states before some fixed time limit. The jump rates of the chain are expressed as functions of a rarity parameter in a way that the probability of the rare event goes to zero when the rarity parameter goes to zero, and the behavior of our estimators is studied in this asymptotic regime. After giving a general expression for the zero-variance change of measure in this situation, we develop an approximation of it via a power series and show that this approximation provides a bounded relative error when the rarity parameter goes to zero. We illustrate the performance of our approximation on small numerical examples of highly reliable Markovian systems. We compare it to a previously proposed heuristic that combines forcing with balanced failure biasing. We also exhibit the exact zero-variance change of measure for these examples and compare it with these two approximations.

1 INTRODUCTION

Rare-event simulation is concerned with estimating a performance measure usually expressed as a mathematical expectation, and whose value is strongly affected by certain events that occur rarely. A simple but commonly encountered special case is when the quantity to be estimated is the probability of occurrence of some rare event. When this probability is very small, we may have to simulate the model an excessively large number of times to be able to

estimate it with reasonable relative accuracy, because the rare event occurs only very rarely, by definition.

The best-known way of handling this problem is *importance sampling* (IS): change the probability laws that drive the system, to make the rare event occur more frequently, and multiply the estimator by an appropriate likelihood ratio to recover an unbiased estimator of the quantity of interest. It is well-known that in the case where the estimator is a nonnegative real-valued random variable, there is a change of measure (change of probability law) that gives a zero-variance estimator; that is, the IS estimator becomes a constant (Hammersley and Handscomb 1964, Glynn and Iglehart 1989, Juneja and Shahabuddin 2006). More generally, zero-variance estimators have also been defined for Markov chain models where we want to estimate the probability that the chain reaches a given set of (rare) states before reaching another set of states (Juneja and Shahabuddin 2006), and for finite-state discrete-time chains with a state-dependent cost, where we want to estimate the total expected cost until the chain hits a given set of states (Booth 1987, Kollman et al. 1999).

However, implementing this zero-variance change of measure requires the exact knowledge of the total expected cost-to-go (future costs) from any state that can be visited during the simulation. If we know this, there is no need to perform a simulation in the first place! Nevertheless, attempts to approximate the zero-variance change of measure by various heuristics have been successful in some contexts (Booth 1987, Booth 2001, Kuruganti and Strickland 1997, Kollman et al. 1999, Bolia, Juneja, and Glasserman 2004, Ahamed, Borkar, and Juneja 2006, Juneja and Shahabuddin 2006).

In this article, we express the zero variance change of measure, and develop approximations of the zero-variance importance sampling in the setting of a continuous-time Markov chain (CTMC), in which we want to estimate the probability of reaching a given set of states before a given time horizon (or time limit) t_* . Here we assume that t_* is fixed, but our results will be generalized to the case where it is a random stopping time with fixed distribution (that does not depend on the rarity), under the additional assumption that having reached the stopping time or not can be determined by looking only at the current time and the current state of the CTMC. One application of this is in a reliability setting, where we want to estimate the probability that the system fails before accomplishing its mission (Nakayama and Shahabuddin 2004).

Following Shahabuddin (1994b), our model has a rarity parameter ε and the jump rates are assumed to be polynomial functions of ε . We are interested in the asymptotic behavior when $\varepsilon \rightarrow 0$. For some of the rates, the polynomial may have degree 0, i.e., some rates can be constant as functions of ε . The other rates converge to 0 when $\varepsilon \rightarrow 0$. We assume that over every sample path that leads to the rare event of interest, at least one jump rate is not constant; then the rare-event probability converges to 0 and the relative error of its naive estimator increases to infinity, when $\varepsilon \rightarrow 0$.

In our setting, the zero-variance change of measure can be written in terms of the original jump rates of the chain and the probability $\mu(x, t)$ of reaching the rare event before the time limit if the CTMC is in state x and there remains t units of time. We propose an approximation of the function $\mu(x, t)$ by the first terms of its expansion in powers of ε . We sketch a proof that an IS scheme that uses this approximation gives a bounded relative error in general, and a relative error that converges to 0 as $\mathcal{O}(\sqrt{\varepsilon})$ if all the original jump rates are $\mathcal{O}(\varepsilon)$. We report some numerical experiments with simple models of highly reliable Markovian systems (HRMS). In these experiments, the proposed scheme gives much smaller variance, for small ε , than a combination of forcing and balance failure biasing, recommended in Nakayama and Shahabuddin (2004) for this type of situation.

In the next section, we define our CTMC model and derive the zero-variance sampling scheme. In Section 3, we introduce a rarity parameter for our model, develop our approximation of the zero-variance sampling scheme, discuss how it can be implemented in practice, and show that the corresponding estimator has bounded relative error. Numerical illustrations are given in Section 4. A conclusion follows.

2 CTMC MODEL AND ZERO-VARIANCE SAMPLING

2.1 CTMC Model Over A Finite Time Horizon

Consider a continuous-time Markov chain (CTMC) $\{X_j, j \geq 0\}$ with denumerable state space \mathcal{X} . When the chain is in state x , the jump rate to state x' is $\lambda_{x,x'}$, the total jump rate is $\lambda_x = \sum_{x' \in \mathcal{X}} \lambda_{x,x'} < \infty$ (so the time until the next jump is exponential with mean $1/\lambda_x$), and the next state is x' with probability $p_{x,x'} = \lambda_{x,x'}/\lambda_x$. Jumps to the same state (from x to x) are allowed. We may also have absorbing states, for which $\lambda_x = 0$.

We are interested in estimating the probability that the chain hits a given set of states $\Delta \subset \mathcal{X}$ before some fixed time limit t_* . Let T_j be the time remaining on the clock (t_* minus the current time) at the j th jump of the CTMC, into state X_j . The process starts in state $X_0 \in \mathcal{X}$ with remaining clock time T_0 . Our main interest is for $T_0 = t_*$, but for notational convenience we shall allow T_0 to take any value $t \in \mathbb{R}$. Let

$$\tau = \inf\{j \geq 0 : T_j \leq 0 \text{ or } X_j \in \Delta\}.$$

To avoid fancy complications, we assume that $\mathbb{P}[\tau < \infty] = 1$. Define

$$X = \mathbb{I}[X_\tau \in \Delta \text{ and } T_\tau > 0],$$

where \mathbb{I} is the indicator function. That is, $X = 1$ if the chain hits Δ before the time limit, and $X = 0$ otherwise.

We consider the discrete-time Markov chain (DTMC) $\{(X_j, T_j), j \geq 0\}$, for which the second component of the state indicates the current time. When this DTMC is in state (x, t) , the density of the next state at $(x', t - \delta)$ is given by

$$\pi(x', t - \delta | x, t) = p_{x,x'} \lambda_x \exp[-\lambda_x \delta] \quad (1)$$

for $\delta > 0$, and 0 elsewhere. Note that this density is partitioned into several pieces, one piece for each value of x' , and its total integral over all pieces equals 1.

Let $\mu(x, t)$ be the probability of hitting Δ before the time limit when the DTMC is in state (x, t) . Thus,

$$\mu(x, t) = \begin{cases} 1 & \text{if } x \in \Delta \text{ and } t > 0, \\ 0 & \text{if } t \leq 0, \\ \mathbb{E}[X | X_0 = x, T_0 = t] & \text{otherwise.} \end{cases}$$

For the latter case ($x \notin \Delta$ and $t > 0$), we have the recurrence

$$\begin{aligned} \mu(x, t) &= \mathbb{E}[\mu(X_1, T_1) | X_0 = x, T_0 = t] \\ &= \int_0^t \sum_{x' \in \mathcal{X}} \mu(x', t - \delta) p_{x,x'} \lambda_x \exp[-\lambda_x \delta] d\delta. \quad (2) \end{aligned}$$

Clearly, $\mu(x, t)$ is a nondecreasing function of t for each x . An intuitive argument for this is that enlarging t can only give more opportunity for hitting Δ ; a proof can be made via a path-by-path comparison.

2.2 A Zero-Variance Sampling Scheme

We consider replacing the conditional (transition) density π in (1) by another density g such that $g(x', t - \delta | x, t) > 0$ whenever $\mu(x', t - \delta)\pi(x', t - \delta | x, t) > 0$, until we reach the stopping time τ . The estimator X is then replaced by

$$X_{\text{is}} = X \prod_{i=1}^{\tau} L(X_{i-1}, T_{i-1}, X_i, T_i), \quad (3)$$

where

$$L(X_{i-1}, T_{i-1}, X_i, T_i) = \frac{\pi(X_i, T_i | X_{i-1}, T_{i-1})}{g(X_i, T_i | X_{i-1}, T_{i-1})}$$

and an empty product is assumed to be 1. Thus, the original estimator is weighted by the likelihood ratio that corresponds to the change of densities.

Let $\mathbb{E}_{g,x,t}$ and $\text{Var}_{g,x,t}$ denote the expectation and variance operators under the conditional densities g , from initial state $(X_0, T_0) = (x, t)$. We have

$$\mathbb{E}_{g,x,t}[X_{\text{is}}] = \mu(x, t),$$

i.e., X_{is} is an unbiased estimator of $\mu(x, t)$ under the new densities g . Let $v(x, t)$ be the variance of X_{is} under g , when $X_0 = x$ and $T_0 = t$. We have $v(x, t) = 0$ if $x \in \Delta$ or $t \leq 0$; otherwise,

$$\begin{aligned} v(x, t) &\stackrel{\text{def}}{=} \text{Var}_{g,x,t}[X_{\text{is}}] \\ &= \text{Var}_{g,x,t}[\mathbb{E}_{g,x,t}[X_{\text{is}} | X_1, T_1]] + \mathbb{E}_{g,x,t}[\text{Var}_{g,x,t}[X_{\text{is}} | X_1, T_1]] \\ &= \text{Var}_{g,x,t}[\mu(X_1, T_1)L(x, t, X_1, T_1)] \\ &\quad + \mathbb{E}_{g,x,t}[L^2(x, t, X_1, T_1)v(X_1, T_1)] \\ &= \mathbb{E}_{g,x,t}[\mu^2(X_1, T_1)L^2(x, t, X_1, T_1)] \\ &\quad - \mu^2(x, t) + \mathbb{E}_{g,x,t}[L^2(x, t, X_1, T_1)v(X_1, T_1)] \\ &= \mathbb{E}_{g,x,t}[\mu^2(X_1, T_1) + v(X_1, T_1)]L^2(x, t, X_1, T_1) - \mu^2(x, t). \end{aligned}$$

Suppose now that g is the density g_0 defined by

$$\begin{aligned} g_0(x', t - \delta | x, t) &= \frac{\mu(x', t - \delta)}{\mu(x, t)} \pi(x', t - \delta | x, t) \\ &= \frac{\mu(x', t - \delta)}{\mu(x, t)} p_{x,x'} \lambda_x \exp[-\lambda_x \delta] \end{aligned}$$

for $x' \in \mathcal{X}$ and $0 < \delta < t$, and 0 elsewhere, if $\mu(x, t) > 0$. By integrating with respect to δ and summing over x' , we easily see that these $g_0(\cdot | x, t)$ are probability densities (they integrate to 1). Under this density, the time until the next jump is no longer exponential, so we no longer have a CTMC. This time is nonzero only over the interval $(0, t)$, and it is not a truncated exponential either; in fact, since $\mu(x', t - \delta)$ is decreasing in δ , the right “tail” of the new distribution of the time to the next jump (before truncation) decreases faster than for the exponential distribution. When $\mu(x, t) = 0$, the density is unchanged: $g_0(x', t - \delta | x, t) = \pi(x', t - \delta | x, t)$.

With this choice of g , whenever $g_0(x', t - \delta | x, t) > 0$, we have

$$\begin{aligned} L(x, t, x', t - \delta) &= \begin{cases} \mu(x, t) / \mu(x', t - \delta) & \text{if } \mu(x, t) > 0, \\ 1 & \text{if } \mu(x, t) = 0. \end{cases} \quad (4) \end{aligned}$$

For state pairs $(x, t, x', t - \delta)$ for which $g_0(x', t - \delta | x, t) = 0$, the definition of $L(x, t, x', t - \delta)$ does not matter, because this quantity will then never occur in the estimator. Since (4) holds for g_0 , we have

$$\mathbb{E}_{g_0,x,t}[\mu^2(X_1, T_1)L^2(x, t, X_1, T_1)] = \mu^2(x, t)$$

and we have the simplification:

$$\begin{aligned} v(x, t) &= \mathbb{E}_{g_0,x,t}[(\mu^2(X_1, T_1) \\ &\quad + v(X_1, T_1))L^2(x, t, X_1, T_1)] - \mu^2(x, t) \\ &= \mathbb{E}_{g_0,x,t}[v(X_1, T_1)L^2(x, t, X_1, T_1)]. \end{aligned}$$

Applying induction, we obtain

$$v(x, t) = \mathbb{E}_{g_0,x,t} \left[v(X_\tau, T_\tau) \prod_{i=1}^{\tau} L^2(X_{i-1}, T_{i-1}, X_i, T_i) \right] = 0$$

because $v(X_\tau, T_\tau) = 0$. Thus, a change of measure that satisfies (4) gives a zero-variance estimator of $\mu(x, t)$ for any $(x, t) \in \mathcal{X} \times [0, \infty)$.

3 ASYMPTOTIC ANALYSIS

3.1 Rarity Parameterization

We define an asymptotic rare-event setting by introducing a parameter $\varepsilon \ll 1$ that characterizes rarity. We shall investigate what happens when $\varepsilon \rightarrow 0$. We assume that all jump rates of our CTMC have the form

$$\lambda_{x,x'} = v_{x,x'} \varepsilon^{k_{x,x'}}$$

for some nonnegative constants $v_{x,x'}$ and $k_{x,x'}$ which are bounded uniformly in (x, x') .

Observe that in our setting, dividing the time horizon t_* and all the times T_j by a factor κ , and multiplying all the jump rates by the same factor, gives an equivalent model. Therefore, without loss of generality, we can assume that $t_* = 1$. We do so for the remainder of this article.

A different way of doing the parameterization would be to assume that the rates $\lambda_{x,x'}$ are fixed and that the time horizon has the form $t_* = v\varepsilon^k$ for some positive constants v and k . This parameterization is equivalent to a special case of our setting, with $k_{x,x'} = k$ for all x, x' , so it offers less flexibility. In the next subsection, we start our analysis with this special case, with $k = 1$.

3.2 Simplest Case: All Rates Proportional to ε

We start with the simplified case where $k_{x,x'} = 1$ for all $x, x' \in \mathcal{X}$, so $\lambda_{x,x'} = v_{x,x'}\varepsilon$. Note that the probability that i or more transitions occur before the time limit is $\mathcal{O}(\varepsilon^i)$.

Consider the states x for which $v_{x,m} > 0$ for at least one $m \in \Delta$. For such states, the recurrence (2) implies that

$$\mu(x, t) = \sum_{\{m \in \Delta: v_{x,m} > 0\}} v_{x,m}\varepsilon t + \mathcal{O}(\varepsilon^2).$$

Next, consider the states x for which $v_{x,m} = 0$ for all $m \in \Delta$ but for which there is an $x' \in \mathcal{Y}$ and $m \in \Delta$ such that $v_{x,x'}v_{x',m} > 0$; for those, we find

$$\mu(x, t) = \sum_{\{m \in \Delta, x' \in \mathcal{Y}: v_{x,x'}v_{x',m} > 0\}} v_{x,x'}v_{x',m}\varepsilon^2 t^2 / 2 + \mathcal{O}(\varepsilon^3).$$

This can be easily generalized: for a state x which is exactly ℓ_x transitions away from Δ , we have

$$\mu(x, t) = \frac{A_x \varepsilon^{\ell_x} t^{\ell_x}}{\ell_x!} + \mathcal{O}(\varepsilon^{\ell_x+1}),$$

where A_x is the sum, over all ℓ_x -step paths from state x to Δ , of the products of the coefficients $v_{x,x'}$ of the transitions $x \rightarrow x'$ on that path.

Using only the leading terms of the above expressions for $\mu(x, t)$, we obtain the following approximation to the zero-variance density:

$$\begin{aligned} g_0(x', t - \delta \mid x, t) &= v_{x,x'}\varepsilon \exp[-\varepsilon v_{x,x'}\delta] \frac{A_{x'} \varepsilon^{\ell_{x'}} (t - \delta)^{\ell_{x'}} / \ell_{x'}!}{A_x \varepsilon^{\ell_x} t^{\ell_x} / \ell_x!} (1 + \mathcal{O}(\varepsilon)) \\ &= v_{x,x'}\varepsilon^{1+\ell_{x'}-\ell_x} \frac{A_{x'} \ell_{x'}!}{A_x \ell_x!} \frac{(t - \delta)^{\ell_{x'}}}{t^{\ell_x}} + \mathcal{O}(\varepsilon^{2+\ell_{x'}-\ell_x}). \end{aligned}$$

We shall approximate the zero-variance density by considering only the (normalized) leading term of the above expansion. For transitions that go toward Δ , we have that

$1 + \ell_{x'} - \ell_x = 0$ by construction. For all other transitions, this exponent of ε is positive, so their densities vanish as $\varepsilon \rightarrow 0$.

Computing this approximation requires determining, for each transition (x, x') , the values of A_x and ℓ_x for the states on both sides of the transition. These values can be computed efficiently using a variant of Dijkstra's algorithm to find *all* shortest paths of the chain from state x to the set Δ , where the length of a path is measured by the number of steps. The A_x 's are then computed by adding up the relevant products of coefficients. For very large state spaces, this can become cumbersome and further heuristics could be developed. In certain situations, e.g., for birth-and-death processes, this computation is almost trivial.

Let \hat{g} denote the resulting approximating density. To compute the normalization factor and generate random variates from \hat{g} , we can proceed as follows. We integrate the leading term with respect to δ , over the interval $[0, t]$, for each x' for which this term is nonzero. The sum of these integrals gives the normalization factor. The relative contribution of each integral to the sum gives the probability that the next state X_j is x' under the new density, so we can easily generate the next state first. Then, knowing that $(X_{j-1}, T_{j-1}, X_j) = (x, t, x')$, we can generate T_j by inversion, exploiting the fact that its conditional density at δ is a monomial in δ . This conditional density is uniform over the interval $[0, t]$ when $\ell_{x'} = 0$ (i.e., for direct transitions to Δ), linearly decreasing over the same interval if $\ell_{x'} = 1$ (i.e., for transitions that bring us only one step away from Δ), and so on.

3.3 Generalizing to Mixed Powers of ε

We now generalize the preceding development to the case where $k_{x,x'}$ is allowed to differ from 1. For the moment, we will assume that all $k_{x,x'} > 0$; later on, we will relax this to allow some $k_{x,x'} = 0$. Define ℓ_x as the smallest sum of exponents $k_{z,z'}$ of transition rates, the minimum being taken over all paths from x to Δ ; all paths attaining this minimum are henceforth referred to as *dominant paths*, because they are the most probable paths for ε sufficiently small. Define also $\Gamma_{x,\tau}$ as the set of paths (X_0, \dots, X_τ) going from $X_0 = x$ to $X_\tau \in \Delta$ in exactly τ steps, and for which the sum of exponents is ℓ_x . For $i = 1, 2, \dots$, let

$$A_{x,i} = \sum_{(x_0, \dots, x_i) \in \Gamma_{x,i}} \prod_{j=1}^i v_{x_{j-1}, x_j}.$$

Under appropriate conditions, it can be shown that for all $x \notin \Delta$,

$$\mu(x, t) = \sum_{i \geq 1} \frac{A_{x,i} \varepsilon^{\ell_x} t^i}{i!} + o(\varepsilon^{\ell_x}), \quad (5)$$

that is, the probability of all non-dominant paths is negligible with respect to dominant ones. For this to hold, we need to make sure that the sum over all other paths, which have exponent larger than ℓ_x , induces a probability with exponent larger than ℓ_x too. This is not trivially satisfied because there is generally an infinite number of such paths. Nakayama and Shahabuddin (2004) provide a set of sufficient conditions under which they prove a version of (5) for a specific model of HRMS system. Their results could be generalized.

Our approximation $\hat{\mu}(x, t)$ of $\mu(x, t)$ will be the dominant terms (the sum of terms with the smallest power of ε) in the expansion (5). There may be multiple such terms with different powers of t , if there are distinct values of i for which $\Gamma_{x,i} \neq \emptyset$, i.e., if there are dominant paths from x to Δ with different numbers of transitions. In the following calculations, we will include all such terms for completeness. However, in the numerical illustrations of Section 4 we will only retain the term with the lowest power of t . (This makes sense because we know that t never exceeds 1.) This means that in the following expressions, summations over i are replaced by a single term with the lowest i for which the summand is non-zero. This simplifies the expressions significantly, and makes sampling easier because the resulting densities are monomial. For our examples, this simplification hardly affected the results; the difference in the density plots was almost invisible.

Computing the approximation $\hat{\mu}(x, t)$ could be complicated in general, but there are many nontrivial situations where this is rather easy. For example, for HRMS models where the component failure rates do not depend on the current state (the most typical case), the computations can be quite simple, depending on the general form of Δ .

Keeping only the leading powers of ε leads to the following approximation for the zero-variance density:

$$g_0(x', t - \delta | x, t) \approx v_{x,x'} \varepsilon^{k_{x,x'}} \exp[-\lambda_x \delta] \frac{\sum_{i \geq 1} A_{x',i} \varepsilon^{\ell_{x'}} (t - \delta)^i / i!}{\sum_{i \geq 1} A_{x,i} \varepsilon^{\ell_x} t^i / i!}.$$

As a further approximation, we replace the exponential by 1, to obtain \hat{g} :

$$\hat{g}(x', t - \delta | x, t) = v_{x,x'} \varepsilon^{k_{x,x'} + \ell_{x'} - \ell_x} \frac{\sum_{i \geq 1} A_{x',i} (t - \delta)^i / i!}{\sum_{i \geq 1} A_{x,i} t^i / i!} \quad (6)$$

$$= v_{x,x'} \varepsilon^{k_{x,x'} + \ell_{x'} - \ell_x} \frac{\hat{\mu}(x', t - \delta)}{\hat{\mu}(x, t)}. \quad (7)$$

As in Section 3.2, this function must be multiplied by a normalization constant $\rho(x, t)$, which can be calculated

straightforwardly:

$$\frac{1}{\rho(x, t)} = \sum_{x' \in \mathcal{X}} \int_0^t \hat{g}(x', t - \delta | x, t) d\delta = \sum_{x' \in \mathcal{X}} v_{x,x'} \varepsilon^{k_{x,x'} + \ell_{x'} - \ell_x} \frac{\sum_{i \geq 1} A_{x',i} t^{i+1} / (i+1)!}{\sum_{i \geq 1} A_{x,i} t^i / i!}.$$

Since the terms containing the lowest power of ε are already normalized, $\rho(x, t)$ is of order $1 + \mathcal{O}(\varepsilon)$.

So far, we have assumed that all $k_{x,x'} > 0$. This assumption is needed for (5), and since $\lambda_x = \mathcal{O}(\varepsilon^{\min_{x'} k_{x,x'}})$, it also justifies replacing the exponential term by 1 in approximation (6). In many practical problems however, $k_{x,x'} = 0$ for some transitions which are not on a dominant path; for example, repair transitions in HRMS models. One may still apply the change of measure (6) and (7) to this case, and in fact the resulting simulation turns out to be asymptotically efficient, as will be seen in the sequel.

3.4 Robustness of the Estimators as $\varepsilon \rightarrow 0$

Using a change of measure with the density \hat{g} defined in (6), for a sample path $(X_0, T_0, \dots, X_\tau, T_\tau)$ with $X_\tau \in \Delta$ and $T_\tau > 0$ (so $\mu(X_\tau, T_\tau) = 1$), the IS estimator X_{is} in (3) becomes

$$X_{\text{is}} = \prod_{i=1}^{\tau} \frac{\pi(X_i, T_i | X_{i-1}, T_{i-1})}{\hat{g}(X_i, T_i | X_{i-1}, T_{i-1})} \approx \hat{\mu}(X_0, T_0) \prod_{i=1}^{\tau} e^{-\lambda_{X_{i-1}} (T_i - T_{i-1})} (1 + \mathcal{O}(\varepsilon)),$$

where we have used (1) and (7). The fact that we have a random variable even *in the leading term* of the expansion in ε , instead of just a constant as expected for zero variance, is due to the approximation (6), where the exponential term was removed for ease of implementation. As we are going to see, this results in a relative error that does not decrease to zero when $\varepsilon \rightarrow 0$, but remains bounded.

This estimator can be bounded as follows, for $X_0 = x$ and $T_0 = t$:

$$\exp[-t \max_{x' \in \mathcal{X}} \lambda_{x'}] \hat{\mu}(x, t) \leq X_{\text{is}} \leq \hat{\mu}(x, t).$$

With the proposed change of measure, we have $\mu(X_\tau, T_\tau) = 1$ and X_{is} takes the above form with probability 1. Thus, the relative variance (or squared relative error) of X_{is} when $X_0 = x$ and $T_0 = t$ can be bounded as follows:

$$\frac{\text{Var}_{\hat{g}, x, t}[X_{\text{is}}]}{\mu^2(x, t)} = \frac{\mathbb{E}_{\hat{g}, x, t}[X_{\text{is}}^2]}{\mathbb{E}_{\hat{g}, x, t}[X_{\text{is}}]^2} - 1 \leq \frac{\sup X_{\text{is}}^2}{(\inf X_{\text{is}})^2} - 1 \leq \exp[2t \max_{x' \in \mathcal{X}} \lambda_{x'}] - 1.$$

Thus, the estimator's relative error remains bounded when $\varepsilon \rightarrow 0$. Moreover, in the case where all the rates are $\mathcal{O}(\varepsilon)$ or smaller, i.e., if $k_{x,x'} \geq 1$ for all x, x' , then $\max_{x' \in \mathcal{X}} \lambda_{x'} = \mathcal{O}(\varepsilon)$ and the relative error is then at most $\mathcal{O}(\sqrt{\varepsilon})$. Including the normalization factor $1 + \mathcal{O}(\varepsilon)$ that we have neglected here does not affect these conclusions.

As noted earlier, models with constant rates and a time horizon of order ε can be studied by applying a simple change of variables, which results in all rates being of order ε ; the previous result implies that in this case, the relative error is bounded as $\mathcal{O}(\sqrt{\varepsilon})$. Experimental results in Section 4.1 show a relative error proportional to ε , indicating that the bound may actually be pessimistic.

4 NUMERICAL ILLUSTRATION

We illustrate the use of our approximation of the zero-variance estimator with two toy examples representing HRMS models in 1 and 2 dimensions. The dimension represents the number of component types in the system and the state is the number of components of each type that are down. Each transition represents the failure or the repair of one component. For the one-dimensional example, we consider two variants, one with all transition rates of order ε and one having a transition rate (corresponding to a repair) equal to 1. For both examples, Δ is the set of states where all components of any given type are down, and we take $t_* = 1$.

We compare three simulation approaches: standard simulation (i.e., no importance sampling), forcing+BFB, and our method which uses an approximation to the zero-variance densities. Forcing+BFB (Balanced Failure Biasing) is the classical importance sampling simulation method for the types of HRMS examples we consider; see (Shahabuddin 1994a, Nakayama and Shahabuddin 2004). Briefly, forcing means that the probability distribution for the time until the first failure is replaced by the distribution conditional on that failure happening before the time horizon; and balanced failure biasing means that in any state, all failure transitions are made equally likely and of order 1 rather than order ε . We vary ε from 1 to 10^{-4} . All our simulation results are obtained using 10^5 independent replications.

4.1 Example 1: A Simple Model With a One-Dimensional State Space

We consider a system with 2 components of a single type, with Δ containing the state where both components have failed. Each component has failure rate ε . In variant (a), we assume that the repair rate is also ε ; this allows us to look at the case where all transition rates are small and have the same order, as in Section 3.2, or equivalently the case where the rates are fixed and the time horizon gets small. In variant (b), the repair rate is 1.

Tables 1 and 2 present the simulation results for the three methods, for variants (a) and (b), respectively. For variant (a), the relative error for our method seems to decrease proportionally to ε , which is better than the $\mathcal{O}(\sqrt{\varepsilon})$ bound obtained theoretically. For variant (a), comparing our method with forcing+BFB might be unfair, because the latter method was designed for models where the repair rates were not small, as in variant (b). For this second variant, the results agree very well with the bounded relative error property for our method, as predicted theoretically in Section 3.4, and also for forcing+BFB, as proved in (Nakayama and Shahabuddin 2004). However, the relative error is about 10 times smaller with our method than with forcing+BFB.

Table 1: Simulation Results for Variant (a) of Example 1.

ε	Method	$\mu(0, t_*)$	Relative Error
1	Standard	0.3356	0.0045
1	Forcing+BFB	0.3348	0.0039
1	Our method	0.3356	0.0019
0.1	Standard	0.00851	0.034
0.1	Forcing+BFB	0.00871	0.014
0.1	Our method	0.0087754	0.0001527
0.01	Standard	1.1e-4	0.30
0.01	Forcing+BFB	1.04e-4	0.044
0.01	Our method	0.986779e-4	0.00001495
0.0001	Standard	0	N/A
0.0001	Forcing+BFB	1.2e-8	0.41
0.0001	Our method	0.999867e-8	0.0000001491

Table 2: Simulation Results for Variant (b) of Example 1.

ε	Method	$\mu(0, t_*)$	Relative Error
1	Standard	0.3356	0.0045
1	Forcing+BFB	0.3348	0.0039
1	Our method	0.3355	0.0019
0.1	Standard	0.0066	0.039
0.1	Forcing+BFB	0.00680	0.0062
0.1	Our method	0.006762	0.000707
0.01	Standard	1.0e-4	0.32
0.01	Forcing+BFB	0.734e-4	0.0066
0.01	Our method	0.7292e-4	0.000696
0.0001	Standard	0	N/A
0.0001	Forcing+BFB	7.40e-9	0.0066
0.0001	Our method	7.353e-9	0.000697

4.2 Example 2: An HRMS Model with Two-Dimensional State Space

We consider now a system with 2 types of components and 2 components of each type. The different rates are shown in Figure 1, where (absorbing) states where the system is failed are in grey. Simulation results are given in Table 3. As in the previous example, both forcing+BFB and our method




Figure 1: Markov chain transition graph for Example 2.




Figure 4: Transition probability densities in Example 2.

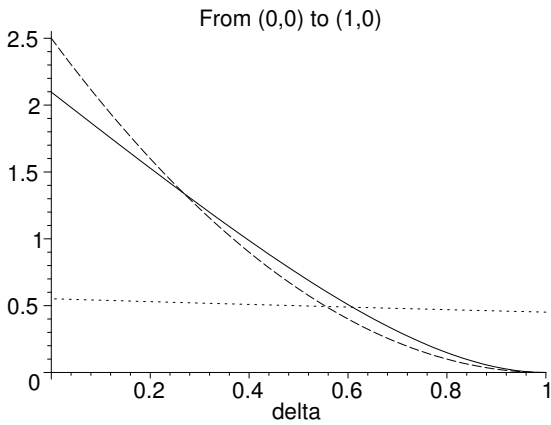


Figure 2: Transition probability densities in Example 2. The solid line is for the zero-variance density, the dashed line for our approximation, and the dotted line for forcing+BFB.




Figure 5: Transition probability densities in Example 2.




Figure 3: Transition probability densities in Example 2.




Figure 6: Transition probability densities in Example 2.

turn out to have bounded relative error, in accordance with theory, and our method wins by a factor of about 30 for the small values of ε .

For this example, we also show graphs of the (partial) probability densities. Figures 2 through 6 show (1) the density g_0 of the zero-variance change of measure (solid line), (2) our approximate density \hat{g} (dashed line), and (3) the density effectively being used by forcing+BFB (dotted line), for some selected transition types (i.e., pairs x, x'); the densities for the other transition types have a similar behavior. Each density is shown as a function of δ for $t = 1$ and $\varepsilon = 0.1$. The zero-variance densities have been calculated by computing the time-dependent state probabilities directly by numerical evaluation of the exponential of the transition rate matrix, see (Kulkarni 1995, pages 266-269) or (de Souza e Silva and Gail 2000): If A is the infinitesimal generator of the CTMC $\{X_j, j \geq 0\}$, the transition probability matrix $P(t) = (P_{x,x'}(t))_{x,x' \in \mathcal{X}}$, with $P_{x,x'}(t)$ the probability of being in x' at time t given that we were in x at time 0, is such that $P(t) = e^{At}$. If states in Δ are absorbing states, $\mu(x, t)$ can then be computed as $\mu(x, t) = \sum_{x' \in \Delta} P_{x,x'}(t)$.

Figures 2, 3, and 4 show the (partial) densities of the transition times on the path that dominates in the limit when $\varepsilon \rightarrow 0$. These figures clearly demonstrate the fact that the densities in our approach (dashed) are monomial, with the direct transition to the failure state having a uniform density, the one before that having a linear density, and the one before that a quadratic density. With the exception of the last (uniform) density, our approximate densities are reasonably similar to the zero-variance densities (solid), whereas the forcing+BFB densities (dotted) are quite different, and typically much lower. This means that forcing+BFB tends to put too little probability mass on the dominant path. Again, the larger difference for the density corresponding to direct transitions to failures stems from the absence of the exponential term in approximation (6).

Figure 5 shows the densities for a transition toward failure but on a non-dominant path, and Figure 6 gives the densities for a repair transition. One sees that forcing+BFB tends to put too much probability on these transitions.

5 CONCLUSION

In this article, we have presented a framework for importance sampling and calculation of the zero variance change of measure in Markov chains, and shown how different HRMS examples fit this framework. For this case, we have shown that a first-order approximation to the zero-variance distribution can easily be obtained and is effective.

This work needs to be pushed forward. First of all, we need to relax as much as possible the conditions used in our results. Second, the HRMS examples in the present paper have very small state spaces for sake of illustrativeness, in order to get the explicit value of the zero variance change of

Table 3: Simulation results for the two-dimensional example.

ε	Method	$\mu((0,0), t_*)$	Relative Error
1	Standard	0.1839	0.0067
1	Forcing+BFB	0.1835	0.0072
1	Our method	0.1845	0.0036
0.1	Standard	0.00012	0.29
0.1	Forcing+BFB	0.000111	0.028
0.1	Our method	0.0001123	0.0028
0.01	Standard	0	N/A
0.01	Forcing+BFB	8.7e-8	0.041
0.01	Our method	8.693e-8	0.00158
0.0001	Standard	0	N/A
0.0001	Forcing+BFB	8.4e-14	0.043
0.0001	Our method	8.402e-14	0.00121

measure. We plan look at practical problems with very large state spaces, such as for instance that in (Nakayama and Shahabuddin 2004). Other applications in insurance risk and queuing networks are also of interest. Last but not least, we plan to investigate the case where the exponential term is not skipped in Equation (6). Sampling the approximate density will then be a little more complicated, but significantly better results can be expected.

ACKNOWLEDGMENTS

This research has been supported by NSERC-Canada grant No. ODGP0110050 and a Canada Research Chair to the second author, as well as EuroFGI Network of Excellence for the third and fourth authors, and INRIA's cooperative research initiative RARE to all authors. The work was done while the first two authors were on leave at IRISA, in Rennes.

REFERENCES

- Ahamed, I., V. S. Borkar, and S. Juneja. 2006. Adaptive importance sampling for Markov chains using stochastic approximation. *Operations Research* 54 (3): 489–504.
- Bolia, N., S. Juneja, and P. Glasserman. 2004. Function-approximation-based importance sampling for pricing American options. In *Proceedings of the 2004 Winter Simulation Conference*, 604–611: IEEE Press.
- Booth, T. E. 1987. Generalized zero-variance solutions and intelligent random numbers. In *Proceedings of the 1987 Winter Simulation Conference*, 445–451: IEEE Press.
- Booth, T. E. 2001. An approximate Monte Carlo adaptive importance sampling method. *Nuclear Science and Engineering* 138 (1): 96–103.
- de Souza e Silva, E., and R. Gail. 2000. Transient solutions for Markov chains. In *Computational Probability*, ed.

- W. K. Grassmann, 43–79. Norwell, Mass.: Kluwer Academic.
- Glynn, P. W., and D. L. Iglehart. 1989. Importance sampling for stochastic simulations. *Management Science* 35:1367–1392.
- Hammersley, J. M., and D. C. Handscomb. 1964. *Monte Carlo methods*. London: Methuen.
- Juneja, S., and P. Shahabuddin. 2006. Rare event simulation techniques: An introduction and recent advances. In *Simulation*, ed. S. G. Henderson and B. L. Nelson, Handbooks in Operations Research and Management Science, 291–350. Amsterdam, The Netherlands: Elsevier. Chapter 11.
- Kollman, C., K. Baggerly, D. Cox, and R. Picard. 1999. Adaptive importance sampling on discrete Markov chains. *Annals of Applied Probability* 9 (2): 391–412.
- Kulkarni, V. G. 1995. *Modeling and analysis of stochastic systems*. London: Chapman & Hall.
- Kuruganti, I., and S. Strickland. 1997. Optimal importance sampling for Markovian systems with applications to tandem queues. *Mathematics and Computers in Simulation* 44:61–79.
- Nakayama, M. K., and P. Shahabuddin. 2004. Quick simulation methods for estimating the unreliability of regenerative models of large highly reliable systems. *Probability in the Engineering and Information Sciences* 18:339–368.
- Shahabuddin, P. 1994a. Fast transient simulation of Markovian models of highly dependable systems. *Performance Evaluation* 20:267–286.
- Shahabuddin, P. 1994b. Importance sampling for the simulation of highly reliable Markovian systems. *Management Science* 40 (3): 333–352.

AUTHOR BIOGRAPHIES

PIETER-TJERK DE BOER received the PhD degree in computer science in 2000 from the University of Twente, The Netherlands. Since 2001, he has been an assistant professor at the department of Electrical Engineering, Mathematics and Computer Science at the University of Twente. In 2007, he was at INRIA in Rennes for 2 months. His research interests include rare event simulation, importance sampling, queueing theory and large deviations theory, with applications to performance analysis of telecommunication networks, as well as several other topics related to telecommunication systems.

PIERRE L'ECUYER is Professor in the Département d'Informatique et de Recherche Opérationnelle, at the Université de Montréal, Canada. He holds the Canada Research Chair in Stochastic Simulation and Optimization. His main research interests are random number generation, quasi-Monte Carlo methods, efficiency improvement via variance

reduction, sensitivity analysis and optimization of discrete-event stochastic systems, and discrete-event simulation in general. He is currently Associate/Area Editor for *ACM Transactions on Modeling and Computer Simulation*, *ACM Transactions on Mathematical Software*, *Statistical Computing*, *International Transactions in Operational Research*, *The Open Applied Mathematics Journal*, and *Cryptography and Communications*. He obtained the *E. W. R. Steacie* fellowship in 1995-97, a *Killam* fellowship in 2001-03, and became an INFORMS Fellow in 2006. His recent research articles are available on-line from his web page: (<http://www.iro.umontreal.ca/~lecuyer>).

GERARDO RUBINO is a senior researcher at INRIA (the French National Institute for Research in Computer Science and Control) where he is the leader of the DIONYSOS (Dependability, Interoperability and performance analysis of networkS) team at IRISA, Rennes. His research interests are in the quantitative analysis of computer and communication systems, mainly using probabilistic models. He also works on the quantitative evaluation of perceptual quality of multimedia communications over the Internet. He is a member of the IFIP WG 7.3. His e-mail address is (rubino@irisa.fr).

BRUNO TUFFIN received his PhD degree in applied mathematics from the University of Rennes 1 (France) in 1997. Since then, he has been with INRIA in Rennes. He spent 8 months as a postdoc at Duke University in 1999. His research interests include developing Monte Carlo and quasi-Monte Carlo simulation techniques for the performance evaluation of telecommunication systems, and developing new Internet-pricing schemes. He is currently Associate Editor for *INFORMS Journal on Computing*.