

## EFFECTIVE APPROXIMATION OF ZERO-VARIANCE SIMULATION IN A RELIABILITY SETTING

Pierre L'Ecuyer

Bruno Tuffin

Département d'Informatique et de Recherche Opérationnelle  
Université de Montréal, C.P. 6128, Succ. Centre-Ville  
Montréal (Québec), H3C 3J7, CANADA

IRISA-INRIA, Campus de Beaulieu  
35042 RENNES Cedex, FRANCE

### ABSTRACT

For every stochastic simulation model, there is in theory a way of changing the probability laws that drive the system so that the resulting IS estimator has zero variance. This optimal estimation scheme is generally impractical to implement, but it can be possible to approximate it in an effective way. When the model is described by a discrete-time Markov chain that evolves up to some random stopping time, the zero-variance change of measure can be written exactly in terms of a value function that gives the expected cost-to-go from any state of the chain, so it can be approximated by approximating this value function. We detail this approach and show how it can be effectively used to estimate the reliability of a highly-reliable multicomponent system with Markovian behavior. In our implementation, we start with a very simple crude approximation, use it in a first-order IS scheme to obtain a better approximation at a few selected states, interpolate in between, and use this interpolation in our final (second-order) IS scheme. In numerical illustrations, our approach outperforms the popular IS heuristics previously proposed for this class of problems.

### INTRODUCTION

Any discrete-event simulation model can be represented as a discrete-time Markov chain (DTMC), say  $\{Y_j, j \geq 0\}$ , with very large (usually infinite and high-dimensional) state space  $\mathcal{Y}$ . By putting enough information in the state, this chain can also be defined so that the simulation terminates when (and only when) the chain reaches a given (fixed) subset of the state space, say  $\Delta \subset \mathcal{Y}$ . Suppose there is a cost  $c(y, y')$  whenever the chain goes from state  $y$  to state  $y'$  at any given step, and we want to estimate  $\mu(y_0) = \mathbb{E}[X \mid Y_0 = y_0]$ , where  $X$  is the total cost until termination, and the conditioning means that the chain starts from a given initial state  $Y_0 = y_0$ .

The standard Monte Carlo method estimates  $\mu(y_0)$  by simulating  $n$  independent copies of the DTMC under its original probability law, and averages the  $n$  copies of  $X$  (Fishman

1996, Law 2007). A confidence interval on  $\mu(y_0)$  is usually computed by assuming that this average is approximately normally distributed (relying on the central-limit theorem) and estimating its variance by the sample variance of the  $n$  copies of  $X$ , divided by  $n$ . There are many situations, however, where the required accuracy calls for a very large value of  $n$ , sometimes so large that the simulation time becomes excessive. This is particularly true in the context of rare-event simulation, where certain events having a large influence on the expected cost occur very rarely (Heidelberger 1995, Juneja and Shahabuddin 2006). If an important rare event happens only once every billion simulation on average, for example, then  $n$  must be at least several billion before we start to obtain an estimator with some meaning.

Variance reduction methods have been designed to solve or alleviate this type of problem (Fishman 1996). In this paper, we focus on one of these methods, called *importance sampling* (IS), which consists in changing the probability laws that drive the system, and multiplying the final estimator by an appropriate correction factor to recover an unbiased estimator of the original expectation (Glynn and Iglehart 1989, Juneja and Shahabuddin 2006). This correction factor turns out to be the likelihood ratio (or Radom-Nikodym derivative) between the old and new probabilities. A major difficulty in applying this method is to figure out how to change the probabilities so that the IS estimator has much smaller variance than the original one.

In theory, there is actually a way of changing the transition probabilities of the chain so that the new estimator has zero variance! In a nutshell, it suffices to multiply the transition probability (or density) from the current state  $y$  to any other state  $y'$  by the ratio of the expected total cost from now on if we go next to state  $y'$ , divided by the (unconditional) expected total cost from now on. Details will be given in the next section. Special cases are examined in Booth (1987), Kollman et al. (1999), Juneja and Shahabuddin (2006), for example. So if we would know the expected total cost until absorption as a function of the initial state  $y$ ,  $\mu(y)$ , we would be able to define a zero-variance sampling scheme. This

function  $\mu(y)$  is called the *cost-to-go* function, or Bellman function, in dynamic programming. The catch is that in practice, we do not know this function. If we knew it, there would be no need for simulation!

But not all hope is lost. By obtaining a rough-cut approximation of the cost-to-go function  $\mu$ , and using it in place of  $\mu$  to approximate the zero-variance change of measure, we may get an IS scheme that achieves a significant variance reduction compared with standard Monte Carlo.

In this paper, we first explain in detail (in the next section) how to define the zero-variance sampling scheme for a DTMC. Then, in the following section, we select a specific class of problems in which the goal is to estimate the reliability of a highly-reliable multicomponent system with Markovian behavior, and show how to approximate the zero-variance change of measure for that class of problems. The models we consider are known in the literature under the name of highly-reliable Markovian systems (HRMS) (Shahabuddin 1994b, Cancela et al. 2002, Nakayama and Shahabuddin 2004). According to our numerical examples in the last section, the new method is much more effective than the IS heuristics previously proposed for this class of problems and are applicable to large problem instances.

## MARKOV CHAIN MODEL AND ZERO-VARIANCE SAMPLING

We consider a discrete-time Markov chain (DTMC)  $\{Y_j, j \geq 0\}$  with general state space  $\mathcal{Y}$ , transition kernel  $P$ , and nonnegative one-step cost function  $c: \mathcal{Y}^2 \rightarrow [0, \infty)$ . When the chain is in state  $Y_{j-1} = y \in \mathcal{Y}$ , the next state  $Y_j$  obeys a probability law defined by  $\mathbb{P}[Y_j \in B \mid Y_{j-1} = y] = P(B \mid y)$  for all (measurable)  $B \subseteq \mathcal{Y}$ , and a transition cost  $c(y, Y_j)$  is incurred. The state space contains a set of absorbing states  $\Delta \subset \mathcal{Y}$  from which the transition cost is zero:  $P(\{y\} \mid y) = 1$  and  $c(y, y) = 0$  for all  $y \in \Delta$ . Let  $\tau = \inf\{j : Y_j \in \Delta\}$ , the number of steps until absorption,

$$X = \sum_{j=1}^{\tau} c(Y_{j-1}, Y_j),$$

the total cost until absorption, and

$$\mu(y) = \mathbb{E}[X \mid Y_0 = y],$$

the expected total cost when starting in state  $y$ . We assume that  $\mathbb{E}[\tau \mid Y_0 = y] < \infty$  and  $\mu(y) < \infty$  for all  $y \in \mathcal{Y}$ . We also assume implicitly that the required measurability conditions hold so that all the expressions introduced here and later are well defined.

The function  $\mu: \mathcal{Y} \rightarrow [0, \infty)$  satisfies the recurrence

$$\mu(y) = \mathbb{E}[c(y, Y_1) + \mu(Y_1) \mid Y_0 = y] \quad (1)$$

$$= \int_{\mathcal{Y}} [c(y, y_1) + \mu(y_1)] dP(y_1 \mid y) \quad (2)$$

for all  $y \in \mathcal{Y}$ .

We consider changing the transition kernel  $P$  for another kernel  $Q$  such that  $Q(B \mid y) > 0$  whenever  $\int_B [c(y, y_1) + \mu(y_1)] dP(y_1 \mid y) > 0$ . The estimator  $X$  is then replaced by

$$X_{\text{is}} = \sum_{j=1}^{\tau} c(Y_{j-1}, Y_j) \prod_{i=1}^j L(Y_{i-1}, Y_i), \quad (3)$$

where  $L(Y_{i-1}, Y_i) = (dP/dQ)(Y_i \mid Y_{i-1})$  and  $dP/dQ$  is the Radon-Nikodym derivative of  $P$  with respect to  $Q$  (Glynn and Whitt 1989, Heidelberger 1995, Bucklew 2004, Juneja and Shahabuddin 2006). Let  $\mathbb{E}_{Q,y}$  and  $\text{Var}_{Q,y}$  denote the expectation and variance operators under the measure (transition kernel)  $Q$ , from initial state  $Y_0 = y$ . We have

$$\mathbb{E}_{Q,y}[X_{\text{is}}] = \mu(y),$$

i.e.,  $X_{\text{is}}$  is an unbiased estimator of  $\mu(y)$  under the new kernel  $Q$ . The variance of  $X_{\text{is}}$  under  $Q$ , for  $Y_0 = y$ , is

$$\begin{aligned} V(y) &\stackrel{\text{def}}{=} \text{Var}_{Q,y}[X_{\text{is}}] \\ &= \text{Var}_{Q,y}[\mathbb{E}_{Q,y}[X_{\text{is}} \mid Y_1]] + \mathbb{E}_{Q,y}[\text{Var}_{Q,y}[X_{\text{is}} \mid Y_1]] \\ &= \text{Var}_{Q,y}[(c(y, Y_1) + \mu(Y_1))L(y, Y_1)] \\ &\quad + \mathbb{E}_{Q,y}[L^2(y, Y_1)V(Y_1)] \\ &= \mathbb{E}_{Q,y}[(c(y, Y_1) + \mu(Y_1))^2 L^2(y, Y_1)] - \mu^2(y) \\ &\quad + \mathbb{E}_{Q,y}[L^2(y, Y_1)V(Y_1)] \\ &= \mathbb{E}_{Q,y}[(c(y, Y_1) + \mu(Y_1))^2 + V(Y_1)]L^2(y, Y_1) \\ &\quad - \mu^2(y). \end{aligned}$$

Suppose now that we choose the measure  $Q$  in a way that for each  $y$  such that  $\mu(y) > 0$ ,

$$dQ(y_1 \mid y) = dP(y_1 \mid y)(c(y, y_1) + \mu(y_1))/\mu(y). \quad (4)$$

By integrating with respect to  $y_1$ , we easily see that  $Q(\cdot \mid y)$  thus defined is a probability measure (it integrates to 1). If  $\mu(y) = 0$ , we simply take  $Q(\cdot \mid y) = P(\cdot \mid y)$ . With this choice of  $Q$ , whenever  $dQ(y_1 \mid y) > 0$ , we have

$$L(y, y_1) = \begin{cases} \mu(y)/(c(y, y_1) + \mu(y_1)) & \text{if } \mu(y) > 0, \\ 1 & \text{if } \mu(y) = 0. \end{cases} \quad (5)$$

For state pairs  $(y, y_1)$  for which  $dQ(y_1 | y) = 0$ ,  $L(y, y_1)$  may be undefined, but this does not matter because  $L(y, y_1)$  will then never occur in the estimator. If (5) holds, then

$$\mathbb{E}_{Q,y}[(c(y, Y_1) + \mu(Y_1))^2 L^2(y, Y_1)] = \mu^2(y)$$

and we have the simplification:

$$\begin{aligned} & V(y) \\ &= \mathbb{E}_{Q,y}[(c(y, Y_1) + \mu(Y_1))^2 + V(Y_1)] L^2(y, Y_1) - \mu^2(y) \\ &= \mathbb{E}_{Q,y}[V(Y_1) L^2(y, Y_1)]. \end{aligned}$$

Applying induction, we obtain

$$V(y) = \mathbb{E}_{Q,y} \left[ V(Y_\tau) \prod_{j=1}^{\tau} L(Y_{j-1}, Y_j) \right] = 0$$

because  $V(Y_\tau) = 0$ . We conclude that a change of measure that satisfies (5) gives a zero-variance estimator of  $\mu(y)$  for any  $y \in \mathcal{Y}$ .

As a special case, suppose that the state space  $\mathcal{Y}$  is finite, so that the DTMC has discrete transition probabilities  $p(y_1 | y) = \mathbb{P}[Y_1 = y_1 | Y_0 = y]$ , which we replace by new probabilities  $q(y_1 | y) = \mathbb{Q}[Y_1 = y_1 | Y_0 = y]$ , such that  $q(y_1 | y) > 0$  whenever  $[c(y, y_1) + \mu(y_1)]p(y_1 | y) > 0$ . In this case, we have  $L(y, y_1) = p(y_1 | y)/q(y_1 | y)$ . To satisfy (5), it suffices to choose the probabilities

$$q(y_1 | y) = p(y_1 | y)(c(y, y_1) + \mu(y_1))/\mu(y) \quad (6)$$

if  $\mu(y) > 0$  and  $q(y_1 | y) = p(y_1 | y)$  otherwise.

Since  $\mu$  is generally unknown, the idea is to approximate it by some function  $v$  that can easily be computed during the simulation. This function can also be modified or “learned” along the way, from the results of the simulation runs. For each run, we use the function  $v$  that we have so far (which is our best estimate of  $\mu$ ), and plug it in (6) or (4) in place of  $\mu$  to define the change of probabilities for the IS estimator. That is, in the case of a finite-state Markov chain, we use

$$q(y_1 | y) = p(y_1 | y)(c(y, y_1) + v(y_1))/v(y) \quad (7)$$

if  $v(y) > 0$  and  $q(y_1 | y) = p(y_1 | y)$  otherwise.

One class of methods to approximate  $\mu$  collect information from the simulation to “learn” about  $\mu(y)$  for each state  $y \in \mathcal{Y}$ . Ahamed et al. (2006) do this using a stochastic approximation approach to update the individual entries  $\mu(x)$ . Booth (1982; 1987) has also been using this type of approach to learn  $\mu$  for both splitting and IS; this is implemented in the Los Alamos simulation software. However,

this type of approach has limitations: for real-life problems for which Monte Carlo methods are used, the state space is usually much too large to think of storing an approximation of  $\mu(y)$  for each state  $y$ . For this reason, we will not consider it any further in this paper. What we propose instead is to construct a leaner approximation of  $\mu$ , that can be computed quickly at any state  $y$ , and which requires much less storage than directly estimating  $\mu(y)$  for each state  $y$ .

## A MODEL OF HIGHLY RELIABLE MARKOVIAN SYSTEM

We consider an HRMS model as in Shahabuddin (1994b;a), Nakayama and Shahabuddin (2004), among others. The system has  $c$  types of components, with  $n_i$  identical components of type  $i$ , for  $i = 1, \dots, c$ . Each component is either in a failed state or an operational state. The *state of the system* is represented by a vector  $y = (y^{(1)}, \dots, y^{(c)})$ , where  $y^{(i)}$  is the number of *failed* components of type  $i$ . Thus, we have a finite state space  $\mathcal{Y}$  of cardinality  $(n_1 + 1) \cdots (n_c + 1)$ . This system is assumed to evolve as a continuous-time Markov chain (CTMC), with jump rate  $\lambda(y, y')$  from state  $y$  to state  $y'$ . We suppose that  $\mathcal{Y}$  is partitioned in two subsets  $\mathcal{U}$  and  $\mathcal{F}$ , where  $\mathcal{U}$  is a decreasing set (i.e., if  $y \in \mathcal{U}$  and  $y \geq y' \in \mathcal{F}$ , then  $y' \in \mathcal{U}$ ) that contains the state  $\mathbf{0} = (0, \dots, 0)$  in which all the components are operational.

The CTMC has an embedded DTMC with transition probabilities  $p(y, y')$ . We assume that state  $\mathbf{0}$  can be reached from any state  $y \in \mathcal{U}$ . Define  $\mu(y) = \mathbb{P}[\tau_{\mathcal{F}} < \tau_{\mathbf{0}} | Y_0 = y]$ , where  $\tau_{\mathcal{F}} = \inf\{j \geq 0 : Y_j \in \mathcal{F}\}$  and  $\tau_{\mathbf{0}} = \inf\{j > 0 : Y_j = \mathbf{0}\}$ . Our goal is to estimate  $\mu(\mathbf{0})$ . Note that  $\mu(y) = 1$  when  $y \in \mathcal{F}$ . Estimating  $\mu(\mathbf{0})$  is required in the estimation of several performance measures, such as the mean time to failure, via a regenerative approach (Shahabuddin 1994a).

This estimation can be especially difficult when  $\mu(\mathbf{0})$  is extremely small, which occurs frequently in the case of highly reliable systems. Then, the transitions that correspond to component failures have very small probabilities, while those that correspond to repairs are not so small. In theoretical analysis of such systems, the failure probabilities are often parameterized by some parameter  $\varepsilon$  so that these probabilities converge to 0 when  $\varepsilon \rightarrow 0$  while the repair probabilities remain bounded, and one is interested in the asymptotic behavior of the estimator when  $\varepsilon \rightarrow 0$ . One usually seeks estimators with the highly desirable property of *bounded relative error* (BRE) as a function of  $\varepsilon$ , which means that the variance divided by the square of  $\mu(\mathbf{0})$  remains bounded when  $\varepsilon \rightarrow 0$ .

As a rough-cut approximation of the function  $\mu$  in this model, we start with the following. For any state  $y \in \mathcal{U}$ , let  $\Gamma(y)$  be the set of all paths  $\pi = (y = y_0 \rightarrow y_1 \rightarrow \dots \rightarrow y_k)$  going from state  $y$  to the set  $\mathcal{F}$ , where  $y_j \in \mathcal{U} \setminus \{\mathbf{0}\}$  for  $j = 1, \dots, k-1$ ,  $p(y_{j-1}, y_j) > 0$  for  $j = 1, \dots, k$ , and  $y_k \in \mathcal{F}$ . Each path  $\pi \in \Gamma(y)$  has original probability

$$p(\pi) = \prod_{j=1}^k p(y_{j-1}, y_j)$$

and we have that  $\mu(y) = \sum_{\pi \in \Gamma(y)} p(\pi)$ . However, the latter sum is usually too complicated to compute in practice, because it involves too many paths (there is often an infinite number of paths, because the paths may contain an unlimited number of cycles within  $\mathcal{U}$ ). A very crude estimate is to replace the sum by the maximum; i.e., approximate  $\mu(y)$  by its lower bound

$$v_0(y) = \max_{\pi \in \Gamma(y)} p(\pi).$$

Computing this  $v_0(y)$  amounts to computing the shortest path from  $y$  to  $\mathcal{F}$ , where the length of the directed link from  $y'$  to  $y''$  is  $-\log p(y', y'')$  for any pair of states  $(y', y'')$ , or equivalently the longest path where the length of a path is its probability (the product of probabilities of its one-step transitions). It can be computed cheaply even for complicated structures. When such a shortest path from some state  $y$  is computed, one can store it in a hash table together with all other shortest paths that are computed simultaneously or have been computed previously. Whenever we need  $v_0(y)$  for some state  $y$ , we first check the hash table to see if it has already been computed.

This  $v_0$  would do well in the cases where a single path dominates the sum, which sometimes occurs when all transitions toward  $\mathcal{F}$  have a very small probability. But this lower bound on  $\mu(y)$  could still underestimate the true value by a significant factor. An easy improvement is to take the sum over the most dominant paths (only a few of them) instead of just considering the single most dominant one. These dominant paths can often be selected by exploiting our knowledge of the structure of the system.

In L'Ecuyer and Tuffin (2006), we had good luck with the following simple type of correction (this was used there in a splitting algorithm): estimate  $\mu(\mathbf{0})$  in preliminary runs with some initial IS strategy, and compute the exponent  $\alpha$  such that  $(v_0(\mathbf{0}))^\alpha$  equals this estimate. Then, replace the estimate  $v_0(y)$  by  $v_1(y) = (v_0(y))^\alpha$  for all  $y \in \mathcal{U}$ . This function  $v_1$  matches  $\mu$  for  $y \in \mathcal{F}$  and matches its estimate at  $y = \mathbf{0}$ . In between, it uses an exponential interpolation, motivated by the idea that  $\mu(y)$  typically decreases exponentially (roughly) when  $y$  gets away from  $\mathcal{F}$ .

This idea could be refined. One possibility is to replace  $\alpha$  by a state-dependent correction exponent  $\alpha(y)$ . Here we shall consider the following form for  $\alpha(y)$ :

$$\alpha(y) = 1 + [\alpha(\mathbf{0}) - 1] \frac{\log v_0(y)}{\log v_0(\mathbf{0})},$$

where  $\alpha(\mathbf{0})$  is the value of  $\alpha$  as in the previous paragraph. The rationale for this form is that the correction exponent is needed usually because  $v_0(y)$  accounts for only a few paths and disregards many other ways of reaching  $\mathcal{F}$  from  $y$ . Since the set of paths leading to  $\mathcal{F}$  is generally richer when we are farther from  $\mathcal{F}$ , it appears sensible to have a correction exponent that changes progressively from 1 when we are very close to  $\mathcal{F}$ , to  $\alpha(\mathbf{0})$  when we are in state  $\mathbf{0}$ , and reflects the “distance” to  $\mathcal{F}$  for the states in between. We denote the resulting approximation by  $v_2(y) = (v_0(y))^{\alpha(y)}$ .

Among other possibilities, instead of estimating  $\mu(y)$  only at  $y = \mathbf{0}$ , we can estimate it directly, via IS, over a finite subset of states  $\mathcal{E} \subset \mathcal{U}$ , in preliminary runs. For example, in an HRMS model,  $\mathcal{E}$  could be the set of states where no more than one component is failed, or the set of states where no more than two components are failed (depending on the structure and size of the model). For each state  $y \in \mathcal{E}$ , we define  $v_3(y)$  as the direct estimate of  $\mu(y)$  and we compute  $\alpha(y)$  such that  $(v_0(y))^{\alpha(y)} = v_3(y)$ . For each state  $y \in \mathcal{U} \setminus \mathcal{E}$ , we may interpolate exponentially as follows: We can select a state  $y' \in \mathcal{E}$  such that  $y' \leq y$ , and define  $v_3(y) = (v_0(y))^{\alpha(y')}$  (the selection is arbitrary and could be problem-dependent). Or we may select *all* states  $y' \in \mathcal{E}$  such that  $y \neq y' \leq y$ , and define  $v_4(y)$  as the average of the corresponding values of  $(v_0(y))^{\alpha(y')}$ . Yet another possibility is to estimate a *single* exponent  $\alpha$  for all states  $y$ , based on all available information, e.g., by least-squares regression. As an extreme case, taking  $\mathcal{E} = \mathcal{U}$  means that we have a direct estimate of  $\mu(y)$  for all states  $y \in \mathcal{U}$ , and no interpolation is needed. Then, we are back to an estimator similar to that of Ahamed et al. (2006), depending on how the estimation is done (these authors change the measure dynamically at each step, at the same time as they update their estimates of  $\mu(y)$ ). Our proposal is a matter of compromise between this extreme case and just taking  $v_0$  (the other extreme).

All these possibilities would deserve further analysis and empirical comparison in realistic examples. In the next section, we compare some of them with the best known IS heuristics, on a parameterized example.

Various heuristics have actually been proposed to select the IS change of measure for the type of HRMS setting considered here (Shahabuddin 1994b, Nakayama 1996, Cancela et al. 2002, Tuffin 2004). These methods do not directly attempt to approximate the zero-variance IS. Their aim is

to increase the probabilities of the transition failures, so that we hit  $\mathcal{F}$  before  $\mathbf{0}$  more frequently. The best known are probably *simple failure biasing* (SFB) and *balanced failure biasing* (BFB), developed by Shahabuddin (1994b). For each state  $y \in \mathcal{U} \setminus \{\mathbf{0}\}$ , SFB and BFB increase the probability that the next transition is a failure to a fixed constant  $0 < \rho < 1$ . Within the subset of failure transitions, the individual probabilities are taken proportional to the original ones for SFB and uniform over all transitions for BFB. Within the set of repair transitions, they are taken proportional to the original ones for both schemes. Under certain conditions on the system, BFB has been proved to yield an IS estimator with the BRE property. The standard MC estimator does not enjoy this property.

Alexopoulos and Shultes (2001) have proposed a variation called SBLR, in which the change of measure is made in a way that over any cycle in the sequence of states visited during the simulation, the cumulated likelihood ratio remains bounded when the failure probabilities converge to zero. Some improvements are also proposed that use structural information by identifying events on shortest paths to failure, and pushing more toward those events. We do not consider these versions here.

## NUMERICAL EXAMPLE

Consider a system with  $c = 3$  component types, with  $n_1 = n_2 = n_3$ . Any failed component has an exponentially distributed repair time with rate 1, and an exponentially distributed time to failure with rate  $\lambda_i$  for components of type  $i$ , where  $\lambda_1 = \varepsilon$ ,  $\lambda_2 = 1.5\varepsilon$ , and  $\lambda_3 = 2\varepsilon^2$ , for some parameter  $\varepsilon$ . The system is down whenever fewer than two components of any type are operational. The system starts in state  $\mathbf{0}$  in which all components are operational, and we want to estimate the probability  $\mu(\mathbf{0})$  that it reaches the set  $\mathcal{F}$  of failure states before returning to state  $\mathbf{0}$ . We will experiment with different values of the two parameters  $(n_i, \varepsilon)$ , given in the first two columns of Table 1.

For this example, we define  $v_0(y)$  as follows. For each component type  $i$ , we consider the path that goes from  $y$  to  $\mathcal{F}$  whose only transitions are failures of components of type  $i$ . There are three such paths for each state  $y \in \mathcal{U}$ , and their probabilities  $p(\pi)$  are very easy to compute. We simply define  $v_0(y)$  as the sum of their probabilities.

The methods we try and compare are BFB, SBLR, and our proposed IS method based on zero-variance approximation, using the function  $v_0$  just described, and its modifications  $v_1$  and  $v_2$  defined earlier. We will denote these methods by  $\text{IS}(v_0)$ ,  $\text{IS}(v_1)$ ,  $\text{IS}(v_2)$ , respectively. For each parameter set, Table 1 gives our best estimate of  $\mu(\mathbf{0})$ , obtained from a

very large number of simulation runs with our IS strategies (these numbers are accurate at least for the digits given in the table), the rough-cut approximation  $v_0(\mathbf{0})$  of  $\mu(\mathbf{0})$ , and the estimate obtained from  $n = 2^{20}$  (approximately one million) independent simulation runs, by each of the five methods. Table 2 shows the empirical variances computed from those  $n$  runs, for each method. The table also contains the squared mean, so that one can quickly get an idea of the relative variance (the squared relative error, i.e., the variance divided by the square mean).

In those tables, the entries in parentheses (and in red) are empirical means and variances that clearly underestimate their exact counterparts by a large factor. The red entries for the empirical means are actually even lower (by a large factor) than the crude lower bound  $v_0(\mathbf{0})$ . When we have serious underestimation for the mean, then we have it for the variance as well. The usual explanation is that certain types of paths that have an important relative contribution to the mean are given a too small probability by the IS heuristic, and never occur in the sample, so their contribution is totally missed by the estimator. This reduces both the empirical mean and empirical variance (but not the true variance). If one of these paths would occur, it would potentially have a huge contribution, due to a large likelihood ratio. Interestingly, the entry in brackets (in green) underestimates the exact mean by a significant factor. All these problematic entries are for the BFB and SBLR heuristics, and things generally worsen when  $\mu(\mathbf{0})$  gets smaller and  $n_i$  gets larger. These heuristics were designed to cope with very low failure rates for the components (very small  $\varepsilon$ ), but are not doing very well when the paths to  $\mathcal{F}$  have a large number of transitions (i.e., when  $n_i$  is large, in our example).

With our proposed approach, using  $v_0$  as an approximation already gives much better results than BFB and SBLR, in the sense that we at least get the right order of magnitude for all parameter values. The adjustment  $v_1$  does not provide much improvement over  $v_0$  in this example, whereas  $v_2$  does provide a significant improvement. This seems to confirm the idea that the exponential correction should take into account the distance to failure. This appears to be especially true when  $n_i$  is small, in which case we need much less corrections for the states that are very close to failure than for the initial state (compare the results of  $v_2$  with the other ones when  $n_i = 3$ ).

With any of the three variants ( $v_0$ ,  $v_1$ , and  $v_2$ ), we are able to estimate very small probabilities (smaller than  $10^{-55}$  in our example) quite accurately with a reasonably small number of simulation runs. For example, for  $n_i = 12$  and  $\varepsilon = 0.001$ , if we take the average over  $n$  simulation runs with  $v_2$ , the relative error is  $\sqrt{9.4 \times 10^{-56}} / (3.9 \times 10^{-28} \sqrt{n}) \approx 0.786 / \sqrt{n}$ , so we only need  $n = (78.6/x)^2$  to get  $x\%$  relative

Table 1: Parameter sets and estimates of  $\mu(\mathbf{0})$  with each method.

$n_i$	$\varepsilon$	$\mu(\mathbf{0})$	$\nu_0(\mathbf{0})$	BFB	SBLR	IS( $\nu_0$ )	IS( $\nu_1$ )	IS( $\nu_2$ )
3	0.001	$2.6 \times 10^{-3}$	$1.3 \times 10^{-3}$	$2.7 \times 10^{-3}$	$2.6 \times 10^{-3}$	$2.6 \times 10^{-3}$	$2.6 \times 10^{-3}$	$2.6 \times 10^{-3}$
6	0.01	$1.8 \times 10^{-7}$	$3.4 \times 10^{-8}$	$1.9 \times 10^{-7}$	$[9.9 \times 10^{-7}]$	$1.8 \times 10^{-7}$	$1.8 \times 10^{-7}$	$1.8 \times 10^{-7}$
6	0.001	$1.7 \times 10^{-11}$	$3.4 \times 10^{-12}$	$1.8 \times 10^{-11}$	$(1.8 \times 10^{-16})$	$1.7 \times 10^{-11}$	$1.7 \times 10^{-11}$	$1.7 \times 10^{-11}$
12	0.1	$6.0 \times 10^{-8}$	$3.2 \times 10^{-9}$	$4.8 \times 10^{-8}$	$1.3 \times 10^{-8}$	$6.0 \times 10^{-8}$	$6.2 \times 10^{-8}$	$6.7 \times 10^{-8}$
12	0.001	$3.9 \times 10^{-28}$	$3.5 \times 10^{-29}$	$(1.8 \times 10^{-40})$	$(2.9 \times 10^{-45})$	$3.9 \times 10^{-28}$	$3.9 \times 10^{-28}$	$3.9 \times 10^{-28}$

Table 2: Empirical variances based on  $n = 2^{20}$  independent runs.

$n_i$	$\varepsilon$	$(\mu(\mathbf{0}))^2$	$\alpha$	BFB	SBLR	IS( $\nu_0$ )	IS( $\nu_1$ )	IS( $\nu_2$ )
3	0.001	$6.8 \times 10^{-6}$	0.906	$1.8 \times 10^{-2}$	$8.0 \times 10^{-3}$	$6.5 \times 10^{-4}$	$2.7 \times 10^{-3}$	$9.3 \times 10^{-9}$
6	0.01		0.903	$6.3 \times 10^{-11}$	$(4.5 \times 10^{-16})$	$2.0 \times 10^{-14}$	$1.2 \times 10^{-14}$	$7.7 \times 10^{-15}$
6	0.001		0.939	$8.8 \times 10^{-19}$	$(2.0 \times 10^{-26})$	$1.2 \times 10^{-23}$	$1.1 \times 10^{-23}$	$7.6 \times 10^{-24}$
12	0.1		0.851	$8.1 \times 10^{-10}$	$1.7 \times 10^{-10}$	$1.6 \times 10^{-10}$	$2.9 \times 10^{-10}$	$1.5 \times 10^{-11}$
12	0.001		0.963	$(3.2 \times 10^{-74})$	$(3.5 \times 10^{-84})$	$1.4 \times 10^{-55}$	$9.3 \times 10^{-56}$	$9.4 \times 10^{-56}$

error. For example,  $n = 61$  suffice to obtain 10% relative error. For  $n_i = 12$  and  $\varepsilon = 0.1$ , on the other hand, the relative error with  $\nu_2$  is  $\sqrt{1.5 \times 10^{-11}} / (6.0 \times 10^{-8} \sqrt{n}) \approx 645 / \sqrt{n}$ , so we need  $n \approx (64500/x)^2$  to get  $x\%$  relative error.

Another observation is that  $\alpha$  (used for  $\nu_1$ ) increases, and apparently converges to 1, when  $\varepsilon \rightarrow 0$  for fixed  $n_i$ , or when  $n_i$  increases while  $\varepsilon$  is fixed. The intuitive explanation is that for this example, the dominant (most probable) paths leading to  $\mathcal{F}$  are the direct ones, that involve failures of components of only one type, and these paths are all included in the computation of  $\nu_1$ .

## CONCLUSION

Zero-variance simulation is an utopian ideal that can be achieved only in very simple situations where the quantities of interest can be computed exactly without doing any simulation. However, it can be approximated to a reasonable extent in several interesting situations, and can provide very large variance reduction factors in a practical way. Our numerical examples in this paper illustrate this: our proposed approach yields a low-variance estimator in a setting where all other previously proposed algorithms break their teeth. The method relies on a reasonable approximation of the function  $\mu$  and this would generally depend on the model. Further studies with specific classes of models in different

application areas are needed and should confirm the huge under-explored potential of this approach.

## ACKNOWLEDGMENTS

This research has been supported by Grant OGP-0110050 and a Canada Research Chair to the first author, as well as EuroFGI Network of Excellence and INRIA's cooperative research initiative RARE to the second author. The paper was written while the first author was at IRISA, in Rennes, France.

## REFERENCES

- Ahamed I.; Borkar V.S.; and Juneja S., 2006. *Adaptive Importance Sampling for Markov Chains using Stochastic Approximation*. *Operations Research*, 54, no. 3, 489–504.
- Alexopoulos C. and Shultes B.C., 2001. *Estimating Reliability Measures for Highly-Dependable Markov Systems, Using Balanced Likelihood Ratios*. *IEEE Transactions on reliability*, 50, no. 3, 265–280.
- Booth T.E., 1982. *Automatic Importance Estimation in Forward Monte Carlo Calculations*. *Transactions of the American Nuclear Society*, 41, 308–309.
- Booth T.E., 1987. *Generalized Zero-Variance Solutions and Intelligent Random Numbers*. In *Proceedings of the 1987*

- Winter Simulation Conference. IEEE Press, 445–451.
- Bucklew J.A., 2004. *Introduction to Rare Event Simulation*. Springer-Verlag, New York.
- Cancela H.; Rubino G.; and Tuffin B., 2002. *MTTF Estimation by Monte Carlo Methods Using Markov Models*. *Monte Carlo Methods and Applications*, 8, no. 4, 312–341.
- Fishman G.S., 1996. *Monte Carlo: Concepts, Algorithms, and Applications*. Springer Series in Operations Research. Springer-Verlag, New York.
- Glynn P.W. and Iglehart D.L., 1989. *Importance Sampling for Stochastic Simulations*. *Management Science*, 35, 1367–1392.
- Glynn P.W. and Whitt W., 1989. *Indirect Estimation Via  $L = \lambda w$* . *Operations Research*, 37, 82–103.
- Heidelberger P., 1995. *Fast Simulation of Rare Events in Queueing and Reliability Models*. *ACM Transactions on Modeling and Computer Simulation*, 5, no. 1, 43–85.
- Juneja S. and Shahabuddin P., 2006. *Rare Event Simulation Techniques: An Introduction and Recent Advances*. In S.G. Henderson and B.L. Nelson (Eds.), *Simulation*, Elsevier, Amsterdam, The Netherlands, Handbooks in Operations Research and Management Science. 291–350. Chapter 11.
- Kollman C.; Baggerly K.; Cox D.; and Picard R., 1999. *Adaptive importance sampling on discrete Markov chains*. *Annals of Applied Probability*, 9, no. 2, 391–412.
- Law A.M., 2007. *Simulation Modeling and Analysis*. McGraw-Hill, New York, fourth ed.
- L'Ecuyer P. and Tuffin B., 2006. *Splitting with Weight Windows to Control the Likelihood Ratio in Importance Sampling*. In *Proceedings of ValueTools 2006: International Conference on Performance Evaluation Methodologies and Tools*. ACM Publications, Pisa, Italy, 7 pages.
- Nakayama M.K., 1996. *General Conditions for Bounded Relative Error in Simulations of Highly Reliable Markovian Systems*. *Advances in Applied Probability*, 28, 687–727.
- Nakayama M.K. and Shahabuddin P., 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, no. 3, 333–352.
- Tuffin B., 2004. *On Numerical Problems in Simulations of Highly Reliable Markovian Systems*. In *Proceedings of the 1st International Conference on Quantitative Evaluation of Systems (QEST)*. IEEE CS Press, University of Twente, Enschede, The Netherlands, 156–164.

## AUTHOR BIOGRAPHIES

**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 he was elected INFORMS Fellow in 2006. His recent research articles are available on-line from his web page: <http://www.iro.umontreal.ca/~lecuyer>.

**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*.