

Approximating Zero-Variance Importance Sampling in a Reliability Setting

Pierre L'Ecuyer · Bruno Tuffin

This draft: January 21, 2009

Abstract We consider a class of Markov chain models that includes the highly reliable Markovian systems (HRMS) often used to represent the evolution of multicomponent systems in reliability settings. We are interested in the design of efficient importance sampling (IS) schemes to estimate the reliability of such systems by simulation. For these models, there is in fact a zero-variance IS scheme that can be written exactly in terms of a value function that gives the expected cost-to-go (the exact reliability, in our case) from any state of the chain. This IS scheme is impractical to implement exactly, but it can be approximated by approximating this value function. We examine how this can be effectively used to estimate the reliability of a highly-reliable multicomponent system with Markovian behavior. In our implementation, we start with a simple crude approximation of the value function, we use it in a first-order IS scheme to obtain a better approximation at a few selected states, then we 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. We also perform an asymptotic analysis in which the HRMS model is parameterized in a standard way by a rarity parameter ε , so that the relative error (or relative variance) of the crude Monte Carlo estimator is unbounded when $\varepsilon \rightarrow 0$. We show that with our approximation, the IS estimator has bounded relative error (BRE) under very mild conditions, and vanishing relative error (VRE), which means that the relative error converges to 0 when $\varepsilon \rightarrow 0$, under slightly stronger conditions.

Keywords Monte Carlo · rare events · importance sampling

1 Introduction

Estimating dependability measures for a highly-reliable multicomponent system is an important problem in many areas of applications such as telecommunications, computer sys-

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

B. Tuffin
INRIA Rennes – Bretagne Atlantique
Campus Universitaire de Beaulieu, 35042 Rennes Cedex, FRANCE

tems, aircraft design, air traffic control, power utilities, and many others. Dependability measures of interest include (among others) the *mean time to failure* (MTTF), defined as the expected time until the first failure of the system given that all its components are initially operational, the *mean time between failures* (MTBF), defined as the inverse of the average number of system failures per unit of time in the long run (over an infinite horizon), and the system *availability*, defined as the fraction of the time when the system is operational, in the long run [8,23]. Dependability is typically improved by introducing redundancy in the components. To estimate these measures, the system is often modeled as a continuous-time Markov chain (CTMC), by assuming that component lifetimes and repair times are exponentially distributed.

However, with the exception of very simple situations, the state space of the CTMC is usually so large that analytic and numerical methods are impractical, and one must rely on simulation to estimate the dependability measures of interest. Moreover, standard (crude) Monte Carlo method is much too inefficient, because system failures occur too rarely to provide a meaningful estimator within a reasonable amount of time. In this context, importance sampling (IS) is the standard way of making the important rare events (the system failures, in this case) occur more frequently, in order to recover an unbiased estimator with smaller variance [10, 14, 22, 26]. The main difficulty is to find a good, robust, and easily implementable IS strategy. Popular IS heuristics previously proposed for this class of problems include simple failure biasing (SFB), balanced failure biasing (BFB), the general biasing scheme (GBS), failure-distance biasing (FDB), and the balance likelihood ratio (BLR) method [2, 6, 7, 12, 13, 23, 26, 27, 29]. Here, we study a different approach in which we approximate the zero-variance change of measure via a simple approximation of the value function (in this case, the conditional probability that the rare event occurs, as a function of the current state).

In agreement with previous literature, we adopt a CTMC model with finite state space, whose jumps (transitions) correspond to component failures and repairs. This model is regenerative, and we can define the regeneration times as times when the chain returns to the initial state, in which all the components are operational. The transitions that correspond to component failures have very small rates, but not those that correspond to repairs. As a result, the system tends to return to its initial state very often and rarely reaches the failed state.

In this setting, the MTTF can be written as a ratio of two expectations, where the numerator is the expected time until the system reaches either the failed state or the initial state, and the denominator is the probability μ_0 that the system reaches the failed state before returning to its initial state, both under the assumption that the system starts from its initial state [10, 23, 25]. The numerator is easy to estimate by standard Monte Carlo, but not the denominator, because it is the probability of a rare event. The MTBF and the availability can also be written as ratios of expectations where one of the two expectations in the ratio involves the same rare event as in μ_0 . IS schemes that work well for μ_0 typically work fine also for these expectations [10, 12]. In this paper, we focus on low-variance estimation of μ_0 .

Since the occurrence of the rare event of interest here does not depend on the sojourn times in the states visited by the CTMC, it suffices to simulate the embedded discrete-time Markov chain (DTMC), and this is what we will do.

The remainder is organized as follows. In Section 2, we specify the HRMS model considered in this paper and we briefly summarize the IS heuristics proposed earlier for this model. In Section 3, we define a zero-variance IS sampling scheme to estimate μ_0 , and we propose practical heuristics to approximate this scheme in practice. Zero-variance estimators for general Markov chains have been studied earlier, in [5, 14, 15, 20], for example. In

asymptotic analysis of IS for highly-dependable systems, the failure probabilities are often parameterized by some parameter ε so that these probabilities converge to 0 when $\varepsilon \rightarrow 0$ while the repair probabilities remain bounded, and one studies the asymptotic behavior of the estimator (e.g., its relative error) when $\varepsilon \rightarrow 0$. In Section 4, we recall the basic definitions for this type of analysis, and prove some asymptotic properties of the IS estimators proposed in Section 3. We show that our IS estimators have the BRE property under mild conditions, and the VRE property under slightly stronger conditions. In numerical experiments reported in Section 5, the new IS methods outperform the IS heuristics previously proposed for this class of problems. A conclusion follows in Section 6. A subset of this paper (mostly parts of Sections 3 and 5) was presented at the 2007 European Simulation Conference [19].

2 Markovian Model of a Highly Reliable System

We consider a model of highly reliable Markovian system (HRMS) as in [6, 25–27], among others. The system has c types of components, with n_i identical components of type i , for $i = 1, \dots, c$. We assume that each component is either in a failed state or in an operational state, and that the system evolves as a CTMC whose *state* is a vector $y = (y^{(1)}, \dots, y^{(c)})$, where $y^{(i)}$ represents the number of *failed* components of type i . This implies in particular that the failure and repair rates depend only on y , but they may depend on the entire state y . This CTMC has a finite state space \mathcal{Y} of cardinality $(n_1 + 1) \cdots (n_c + 1)$. Let $\lambda(y, y')$ denote its jump rate from state y to state y' . A jump corresponds either to the (simultaneous) failure of one or more components, or to the repair of one or more failed components. This model covers the notion of failure propagation, where the failure of a component may trigger the (almost) simultaneous failure of other components. It also covers deferred and group repairs [12], where repairs can occur only when there are enough failed components. For example, for a certain type of component, all components of that type might be repaired as a group (simultaneously) at a given rate when at least three of them are failed, and the repair rate for these components is zero otherwise. However, our results in this paper are obtained under additional assumptions that preclude these deferred and group repairs.

One could of course define more general CTMC models where additional information must be incorporated in the state. This happens, for example, if the failed components are repaired according to some priority rules that take into account their failure time (then the state space remains finite but the state must contain the relevant ordering information), or if some of the lifetime or repair distributions are no longer exponential (then the state space is infinite, with continuous coordinates). Our proposed methodology could be extended in principle to these more general models, but it may become much more complicated to implement. The IS methodology studied here is only for the estimation of μ_0 in the CTMC model.

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 $y \geq y' \in \mathcal{Y}$, then $y' \in \mathcal{U}$) that contains the initial state $\mathbf{0} = (0, \dots, 0)$ in which all the components are operational. When we return to this state where all the components are operational, we must recognize that we are not in the same situation as initially, because we now know that we have returned to $\mathbf{0}$ before reaching the failed state. In this second case, we will call the state $\mathbf{0}'$, to make the distinction. In other words, we split the state $\mathbf{0}$ in two states: From now on, $\mathbf{0}$ will refer only to the initial state at the beginning of the simulation, and $\mathbf{0}'$ will refer to the state where all the components are operational at any other stage of the simulation. Formally, this increases the cardinality of the state space by 1. We also denote $\mathcal{U}_0 = \mathcal{U} \setminus \{\mathbf{0}'\}$.

Let $\{Y_j, j \geq 0\}$ be the DTMC embedded in the CTMC. That is, the CTMC starts in state Y_0 and enters state Y_j at its j th jump, for $j = 1, 2, 3, \dots$. The transition probabilities for this DTMC are

$$p(y, y') = \mathbb{P}[Y_j = y' \mid Y_{j-1} = y] = \frac{\lambda(y, y')}{\sum_{y'' \in \mathcal{Y}} \lambda(y, y')}$$

for all $y, y' \in \mathcal{Y}$. We assume that state $\mathbf{0}'$ can be reached (directly or indirectly) from any state $y \in \mathcal{U}_0$. Let $\tau_{\mathcal{F}} = \inf\{j \geq 0 : Y_j \in \mathcal{F}\}$, $\tau_{\mathbf{0}'} = \inf\{j \geq 0 : Y_j = \mathbf{0}'\}$, and $\tau = \min(\tau_{\mathcal{F}}, \tau_{\mathbf{0}'})$. Our assumption that $\mathbf{0}'$ can be reached from any $y \in \mathcal{U}$ implies that $\mathbb{E}[\tau] < \infty$. For all $y \in \mathcal{Y}$, define

$$\mu(y) = \mathbb{P}[\tau_{\mathcal{F}} < \tau_{\mathbf{0}'} \mid Y_0 = y].$$

Note that $\mu(y) = 1$ when $y \in \mathcal{F}$, $\mu(\mathbf{0}') = 0$, and $\mu(\mathbf{0}) = \mu_0$ is the probability we want to estimate. To avoid trivialities, we suppose that $0 < \mu_0 < 1$.

The standard (crude) Monte Carlo method [3] estimates μ_0 by simulating n independent copies of the DTMC under its original probability law, and averages the n copies of the random variable $X = \mathbb{I}[\tau_{\mathcal{F}} < \tau_{\mathbf{0}'}]$, where \mathbb{I} is the indicator function. The variance of X is $\mu_0(1 - \mu_0)$. A confidence interval on μ_0 can be 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 . In our rare-event setting, however, only rare realizations of X are nonzero. The crude Monte Carlo estimator has a relative error (the standard deviation divided by the mean) of $\sqrt{\mu_0(1 - \mu_0)/n}/\mu_0$, which is approximately $(n\mu_0)^{-1/2}$ when μ_0 is small, and becomes very large when μ_0 is very small.

IS consists in replacing the transition probabilities $p(y, y')$ by new probabilities $q(y, y')$ that satisfy $q(y, y') > 0$ whenever $p(y, y')\mu(y') > 0$. We use \mathbb{E}_{is} and Var_{is} to denote the expectation and variance operators under these new probabilities. The estimator X is replaced by

$$X_{\text{is}} = X \prod_{j=1}^{\tau} \frac{p(Y_{j-1}, Y_j)}{q(Y_{j-1}, Y_j)},$$

where the last product is the *likelihood ratio* associated with the sample path. This estimator is unbiased [9]:

$$\mathbb{E}_{\text{is}}[X_{\text{is}} \mid Y_0 = y] = \mathbb{E}[X \mid Y_0 = y] = \mu(y).$$

IS schemes such as SFB, BFB, RBS, FDB, and BLR, mentioned earlier, provide specific ways of selecting the probabilities $q(y, y')$.

BFB [26] defines

$$q(y, y') = \begin{cases} \frac{1}{|F(y)|} & \text{if } y' \in F(y) \text{ and } p_{\text{R}}(y) = 0; \\ \rho \frac{1}{|F(y)|} & \text{if } y' \in F(y) \text{ and } p_{\text{R}}(y) > 0; \\ (1 - \rho) \frac{p(y, y')}{p_{\text{R}}(y)} & \text{if } y' \in R(y); \\ 0 & \text{otherwise,} \end{cases}$$

where $F(y)$ is the set of states y' directly reachable from y (in a single transition) by a failure event, $R(y)$ is the set of states directly reachable from y (in a single transition) by a repair event, $p_{\text{R}}(y) = \mathbb{P}[Y_j \in R(y) \mid Y_{j-1} = y]$, and $\rho \in (0, 1)$ is a constant, usually taken as 0.5. This ρ represents the fraction of probability devoted to failure transitions at each step. This probability is divided equally between all failure transitions having nonzero probability.

SFB [27] is similar, except that $1/|F(y)|$ is replaced by $p(y, y')/\sum_{y' \in F(y)} p(y, y')$ in the above equation, i.e., the failure probability is allocated to the transitions in proportion to their original probabilities.

GBS and other refinements were developed in [12, 13] for situations where there are high-probability cycles. The main idea is to make sure that the probabilities are not reduced too much along these high-probability cycles, because these cycles would otherwise contribute huge values (with low probabilities) to the likelihood ratio, thus increasing the variance. High-probability cycles are common in models with deferred or group repairs, for example. In our asymptotic analysis in Section 4, Assumption 1 will disallow high-probability cycles.

Other schemes have been proposed, trying to take advantage of some knowledge or learning of the model structure. For instance, we may decompose the set of component types in those have already experienced a failure and those who have not, and give fixed probabilities to each of those transitions (to favor those that seem to drive us closer to a failed state). In each case, we can balance the probabilities in each subset or take them proportional to the original ones in the spirit of SFB and BFB [6]. In particular, FDB [7] changes the probabilities by taking into account the minimal number of transitions to failure from each state y' to which we can jump.

In BLR [2], the probabilities are changed in a way that over any cycle, the cumulated likelihood ratio remains bounded when the failure probabilities converge to zero. Variants are also defined that use structural information by identifying events on shortest paths to failure, and pushing more toward those events.

3 Approximate Zero-Variance Sampling

Here we propose and study an alternative heuristic based on the approximation of the following zero-variance IS sampling scheme. Suppose we apply IS to our DTMC with

$$q(y, y') = q^*(y, y') \stackrel{\text{def}}{=} \begin{cases} p(y, y')\mu(y')/\mu(y) & \text{if } 0 < \mu(y) < 1, \\ p(y, y') & \text{otherwise.} \end{cases} \quad (1)$$

Note that $\sum_{y' \in \mathcal{Y}} q(y, y') = 1$ for each y .

Proposition 1 *With the probabilities (1), we have $\text{Var}_{\text{is}}[X_{\text{is}}] = 0$ and $\mathbb{E}_{\text{is}}[\tau] < \infty$ for any initial state $Y_0 = y \in \mathcal{U}_0$.*

Proof Under these probabilities, $\mathbb{P}[X = 0] = 0$, because $q^*(y, \mathbf{0}') = \mu(\mathbf{0}') = 0$ for $y \in \mathcal{U}_0$, so $X = 1$ with probability 1. If $X = 1$, then $\mu(Y_\tau) = 1$, and therefore

$$X_{\text{is}} = X \prod_{j=1}^{\tau} \frac{p(Y_{j-1}, Y_j)}{q(Y_{j-1}, Y_j)} = \prod_{j=1}^{\tau} \frac{\mu(Y_{j-1})}{\mu(Y_j)} = \frac{\mu(Y_0)}{\mu(Y_\tau)} = \mu_0,$$

a constant, so its variance is zero.

For the second property, observe that a state $y \neq \mathbf{0}'$ with $\mu(y) = 0$ cannot be reached under the new probabilities. Then, one can see that there is a constant $\delta > 0$ such that from any visited state $y \neq \mathbf{0}'$, there is a path of probability at least δ leading to \mathcal{F} . The result then follows from a standard geometric trials argument. \square

This zero-variance scheme cannot be implemented exactly, because the function μ is unknown, but it can be replaced by an approximation ν that is easy to compute during the simulation. That is, for each simulation run, we take an approximation ν of the function μ

and plug it in (1) in place of μ to define the change of probabilities for the IS estimator. This gives

$$q(y, y') = \begin{cases} p(y, y')v(y')/\tilde{v}(y) & \text{if } 0 < \tilde{v}(y) < 1, \\ p(y, y') & \text{otherwise,} \end{cases} \quad (2)$$

where $\tilde{v}(y) = \sum_{y' \in \mathcal{Y}} p(y, y')v(y')$. We will assume henceforth that v matches μ at $\mathbf{0}'$ and in \mathcal{F} , where its value is known: $v(\mathbf{0}') = 0$ and $v(y) = 1$ for $y \in \mathcal{F}$.

Several types of *adaptive IS* methods that learn a good function v iteratively have been proposed and studied in the literature; see [1, 4, 14, 15, 20, 24, 28], for example. The general idea of these adaptive methods is to use the realizations of the previous sample paths, where each sample path provides one realization of X , to construct an approximation v . Under some conditions, this approximation can sometimes be shown to converge to μ as $n \rightarrow \infty$, where n is the number of realizations of X . There are even situations where one can prove that the variance converges to zero exponentially in n [4, 15].

In some of these methods, either v or the IS probabilities q are restricted to be in a preselected parametric class of functions parameterized by a multivariate parameter θ , and the current value of θ used for each sample path is selected based of what has been learned from the previous sample paths. Some methods directly parameterize the IS probabilities with θ and then learn adaptively the θ that minimizes the variance. Another class of methods defines v as a linear combination of a fixed set of basis functions, and estimates the best coefficients in the linear combination (the parameters) by least squares, using data from the previous sample paths [14]. The main difficulty with these methods is the choice of basis functions, and they often require a significant amount of work and storage to update θ .

In this paper, we shall consider simpler, more direct, approximations that can be computed at little cost and require practically no storage.

As a rough-cut approximation of the function μ in our model, we start with the following. For any state $y \in \mathcal{U}_0$, let $\Pi(y)$ 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 $y_j \in \mathcal{U}_0$ for $j = 0, \dots, k-1$, $p(y_{j-1}, y_j) > 0$ for $j = 1, \dots, k$, and $y_k \in \mathcal{F}$. Each path $\pi \in \Pi(y)$ has original probability

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

and we have that $\mu(y) = \sum_{\pi \in \Pi(y)} p(\pi)$. However, the latter sum is normally 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

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

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). In general, the worst-case time complexity for computing this shortest path is $\Theta(|\mathcal{U}| \log |\mathcal{U}|)$, which is certainly better than the $\Theta(|\mathcal{U}|^3)$ work required to solve the linear system that would provide all the probabilities $\mu(y)$ for $y \in \mathcal{U}$, but still expensive for large systems. Moreover, the shortest path would be needed from each state y visited during the simulation. Of course, one could save work by storing in a hash table any shortest path computed so far, including the shortest paths from other states y' computed simultaneously while computing the shortest path from a given state y . Whenever we would

need $v_0(y)$ for some state y , we would first check the hash table to see if it has already been computed. Fortunately, all this overhead can often be bypassed by exploiting the model's structure, sometimes even for large and complicated systems, as we shall illustrate in our examples. In many cases, the shortest path can be computed almost at no cost from any y .

This v_0 would do well in the cases where a single path dominates the sum. 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 a small set of disjoint 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. For example, in some cases, it makes sense to compute the probability of a path that leads to \mathcal{F} from failures of a single type of component, do this for each component type, and add these probabilities. This computation can be done very quickly. We will use it in our numerical illustrations, for systems that fail whenever fewer than a given number of components of any given type are operational. When \mathcal{F} has a different type of structure, other inexpensive approximations can often be used instead. For example, if the system fails when the *total* number of failed components (regardless of their type) exceeds a given threshold, we can approximate $\mu(y)$ by approximating the probability of reaching the failed state by a sequence of failure transitions only (no repair). We will give an illustration of this in Section 5.

Our general definition of v_0 for the rest of the paper is then as follows: Select a small set of disjoint paths $\Pi_0(y) \subseteq \Pi(y)$ and define

$$v_0(y) = \sum_{\pi \in \Pi_0(y)} p(\pi).$$

In [18], we had good luck with the following simple type of additional 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 α 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 μ for $y \in \mathcal{F}$ and matches its estimate at $y = \mathbf{0}$. In between, it uses an exponential interpolation, motivated by the crude idea that if there is a single component type, y denotes the number of failed components, and if we assume that the probability that the next event is a failure does not depend on y , then the model turns essentially into a gambler's ruin problem and $\mu(y)$ decreases (approximately) as an exponential function of $-y$ [20].

This idea can be refined. One possibility is to replace α 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 α 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 α 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 [1], 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 Section 5, we compare some of them with the best known IS heuristics, on a few examples.

4 Asymptotic analysis

Asymptotic analysis of IS estimators in the context of HRMS systems, to characterize their behavior when the failure rates converge to zero in certain ways while the rest remains fixed, is usually done by parameterizing the transition rates of the CTMC in a polynomial form as follows [14,21,26]:

$$\lambda(y, y') = \lambda(y, y', \varepsilon) = a(y, y') \varepsilon^{b(y, y')}$$

for some state-dependent constants $a(y, y') \geq 0$ and $b(y, y') \geq 0$ (that do not depend on ε). We usually have $b(y, y') > 0$ for failure transitions and $b(y, y') = 0$ for repair transitions. We then look at what happens when $\varepsilon \rightarrow 0$. Thus, the failure rates become smaller and smaller when $\varepsilon \rightarrow 0$, but the repair rates remain in $\Theta(1)$, and $\mu_0 = \mu_0(\varepsilon) > 0$ converges to 0. This parameterization is transmitted to the transition probabilities $p(y, y')$ of the DTMC. The rationale is that studying the asymptotic properties in this type of setting should give a good idea of what happens for a given model, provided that we use the right constants $b(y, y')$ to somehow mimic the relationships between the different failure rates.

Suppose we have an estimator $X = X(\varepsilon)$ taking its value in $[0, \infty)$, such that $\mathbb{E}[X(\varepsilon)] = \mu_0(\varepsilon)$ and $\text{Var}[X(\varepsilon)] = \sigma_0^2(\varepsilon)$ for each $\varepsilon > 0$.

Definition 1 (a) The estimator $X(\varepsilon)$ has *bounded relative error* (BRE) [11,26], or equivalently *bounded relative variance*, if

$$\limsup_{\varepsilon \rightarrow 0} \frac{\sigma_0(\varepsilon)}{\mu_0(\varepsilon)} < \infty. \quad (3)$$

(b) It has *vanishing relative error* (VRE) [16] if

$$\limsup_{\varepsilon \rightarrow 0} \frac{\sigma_0(\varepsilon)}{\mu_0(\varepsilon)} = 0, \quad (4)$$

or equivalently if

$$\limsup_{\varepsilon \rightarrow 0} \frac{\mathbb{E}[X^2(\varepsilon)]}{\mu_0^2(\varepsilon)} = 1. \quad (5)$$

BRE essentially means that for a given number of simulation runs, the relative width of a confidence interval based on the central-limit theorem remains bounded when $\varepsilon \rightarrow 0$. Under the assumption that the DTMC has no high-probability cycle, BFB provides an IS estimator with BRE, whereas SFB does not [21]. On the other hand, BRE does not necessarily mean that the estimator is practically efficient, because there could be a large hidden constant. It has been recognized that BFB can waste computing time by giving certain low-probability paths more weight than necessary, and this can degrade performance especially when the system has a high level of redundancy [2, 23].

VRE is obviously much stronger than BRE and it is related to zero-variance simulation in the sense that VRE implies that $X(\varepsilon)$ is generated from a probability law that converges in the L_∞ norm to the zero-variance IS scheme [16]. In our setting, this means that VRE implies that

$$\lim_{\varepsilon \rightarrow 0} \sup_{y, y' \in \mathcal{Y}} |q(y, y') - q^*(y, y')| = 0.$$

VRE means that the estimation becomes *easier* when $\varepsilon \rightarrow 0$, which is the opposite of what normally happens with crude Monte Carlo.

In what follows, we provide sufficient conditions for VRE, and then for BRE, for IS sampling based on zero-variance approximation with some function v , applied to the HRMS model of Section 2. The key ingredient is the quality of the approximation of μ by v .

But first, we start by giving a simple example showing that the zero-variance IS scheme is not (asymptotically) balanced in general, in the sense that from a given state, the optimal transition probabilities for different failure transitions can be of different orders (different powers of ε). When this happens, BFB is likely to perform poorly compared with a better (unbalanced) heuristic, even if the asymptotic BRE property holds. Again, this poor performance of BFB (and of balanced schemes in general) in some situations has been observed earlier, for example in [23].

Example 1 Suppose that $c = 2$, $n_1 = n_2 = 2$, and that the system is operational when at least two components (of any kind) are operational. Let the transitions probabilities of the corresponding DTMC be those given in Figure 1, where the states in \mathcal{F} are shaded, and state $y = (y^{(1)}, y^{(2)})$ means (as usual) that $y^{(1)}$ components of type 1 and $y^{(2)}$ components of type 2 are down.

For this small example, we can easily compute $\mu(y)$ for all states y by solving the following system of linear equations (the balance equations of the DTMC):

$$\begin{cases} \mu(1, 1) &= 2\varepsilon^2 + (1 - 2\varepsilon^2)\mu(0, 1)/2 + (1 - 2\varepsilon^2)\mu(1, 0)/2 \\ \mu(0, 2) &= \varepsilon + (1 - \varepsilon)\mu(0, 1) \\ \mu(2, 0) &= \varepsilon + (1 - \varepsilon)\mu(1, 0) \\ \mu(0, 1) &= \varepsilon^3\mu(0, 2) + \varepsilon\mu(1, 1) \\ \mu(1, 0) &= \varepsilon^3\mu(2, 0) + \varepsilon\mu(1, 1) \\ \mu(0, 0) &= \mu(0, 1)/2 + \mu(1, 0)/2 \end{cases}$$

This gives

$$\begin{cases} \mu(1, 1) = 2\varepsilon^2 + (1 - \varepsilon)\frac{2\varepsilon^3 + \varepsilon^4}{1 - \varepsilon + \varepsilon^3 + \varepsilon^4} = \Theta(\varepsilon^2) \\ \mu(2, 0) = \mu(0, 2) = \varepsilon + (1 - \varepsilon)\frac{2\varepsilon^3 + \varepsilon^4}{1 - \varepsilon + \varepsilon^3 + \varepsilon^4} = \Theta(\varepsilon) \\ \mu(0, 1) = \mu(1, 0) = \mu(0, 0) = \frac{2\varepsilon^3 + \varepsilon^4}{1 - \varepsilon + \varepsilon^3 + \varepsilon^4} = \Theta(\varepsilon^3). \end{cases}$$




Fig. 1 Transition probabilities in Example 1

Now, if we look for instance at the two failure transitions from state $(0,1)$, under the zero-variance IS, we have

$$q((0,1), (0,2)) = \frac{\varepsilon^3 \mu(0,2)}{\mu(0,1)} = \Theta(\varepsilon),$$

whereas

$$q((0,1), (1,1)) = \frac{\varepsilon \mu(1,1)}{\mu(0,1)} = \Theta(1),$$

which means that the change of measure is unbalanced between these two failure transitions. The explanation is that all the paths from $(0,0)$ to \mathcal{F} that include the transition $(0,1) \rightarrow (0,2)$ have much smaller probability than those that include the transition $(0,1) \rightarrow (1,1)$ ($O(\varepsilon^4)$ compared with $\Theta(\varepsilon^3)$), so we should waste much less time simulating the former than the latter; i.e., an optimal IS scheme must take $q((0,1), (0,2))$ significantly smaller than $q((0,1), (1,1))$. On the other hand, taking $q((0,1), (0,2))$ too small can be as bad (or worse) than taking it too large. If it is too small, the likelihood ratio will become very large when this transition occurs, and this would increase the variance significantly. \square

The next theorem gives a sufficient condition for VRE. It uses the following definition. For all $y \in \mathcal{U}_0$, there is a set of paths $\Pi_d(y) \subset \Pi(y)$ such that

$$p(\pi) = p(\pi, \varepsilon) = a(\pi) \varepsilon^{b(y)} + o(\varepsilon^{b(y)})$$

for all $\pi \in \Pi_d(y)$, for some constants $a(\pi) > 0$ and $b(y) > 0$ that do not depend on ε , and $p(\pi, \varepsilon) = o(\varepsilon^{b(y)})$ for all $\pi \notin \Pi_d(y)$. The paths in this set $\Pi_d(y)$ are called the *dominant paths from y to \mathcal{F}* . They account asymptotically for all the probability of reaching \mathcal{F} , in the sense that

$$\lim_{\varepsilon \rightarrow 0} \frac{1}{\mu_0(y)} \sum_{\pi \in \Pi_d(y)} p(\pi, \varepsilon) = 1. \quad (6)$$

For $y = \mathbf{0}$, they are just called the dominant paths. For our next results, we make the following assumption.

Assumption 1 All cycles that belong to some path $\pi \in \Pi(y)$ have probability $O(\varepsilon^\delta)$, for some constant $\delta > 0$.

This assumption implies that Π_d cannot contain paths having a cycle, and therefore Π_d must be finite. It also rules out several models with deferred or group repairs, for which the Markov chain has high-probability cycles, as mentioned earlier.

Theorem 1 Under Assumption 1, if $\lim_{\varepsilon \rightarrow 0} v(y)/\mu(y) = 1$ for all $y \in \mathcal{Y}$, then we have VRE. This happens, in particular, if $v(y)$ is defined as the sum of probabilities of all dominant paths from y to \mathcal{F} .

Proof Because \mathcal{Y} is finite and $v(y) = \mu(y)(1 + o(1))$ for all y , there are positive constants δ_1 and r , independent of ε and y , such that

$$|v(y)/\mu(y) - 1| \leq \delta_1 \varepsilon^r$$

for ε sufficiently small. Therefore, there exist $\delta > 0$, independent of ε and y , such that

$$\frac{\tilde{v}(y)}{v(y')} = \sum_{y'' \in \mathcal{Y}} \frac{p(y, y'')v(y'')}{v(y')} \leq \sum_{y'' \in \mathcal{Y}} p(y, y'') \frac{\mu(y'')}{\mu(y')} (1 + \delta \varepsilon^r) = \frac{\mu(y)}{\mu(y')} (1 + \delta \varepsilon^r)$$

for ε sufficiently small. Thus, the IS estimator satisfies

$$X_{\text{is}} \leq \prod_{j=1}^{\tau} \frac{\tilde{v}(Y_{j-1})}{v(Y_j)} \leq \mu_0 (1 + \delta \varepsilon^r)^\tau,$$

giving

$$\mathbb{E}_{\text{is}}[X_{\text{is}}^2] \leq \mu_0^2 \mathbb{E}_{\text{is}}[(1 + \delta \varepsilon^r)^{2\tau}]. \quad (7)$$

Define $m(y)$ as the maximal length (in number of transitions) of a path $\pi \in \Pi_d(y)$. For $y \in \mathcal{F}$, we have $m(y) = 0$. Let $m = \max_{y \in \mathcal{U}_0} m(y)$, the maximal length taken over the whole set of states. Note that this m is finite thanks to our assumption that no dominant path can have a cycle. Similarly, define $p_c(y)$ as the probability, under IS, of using a dominant path from y to \mathcal{F} given that we are in state $y \in \mathcal{U}_0$, and let $p_0 = \min_{y \in \mathcal{U}_0} p_c(y)$. Also for all integers $i \geq 0$, under IS,

$$\mathbb{P}[\tau \geq mi] \leq (1 - p_0)^i,$$

because this probability does not exceed the probability of not reaching \mathcal{F} in i independent trials, starting from y . Consequently, the random variable τ/m is stochastically bounded by a geometric random variable Y with parameter p_0 , whose generating function $G_Y(s) = \mathbb{E}[s^Y] = p_0 s / [1 - (1 - p_0)s]$ is finite for $(1 - p_0)|s| < 1$. Therefore,

$$\mathbb{E}_{\text{is}}[X_{\text{is}}^2] \leq \mu_0^2 \mathbb{E}[(1 + \delta \varepsilon^r)^{2m} \tau/m] \leq \mu_0^2 \frac{p_0 (1 + \delta \varepsilon^r)^{2m}}{1 - (1 - p_0)(1 + \delta \varepsilon^r)^{2m}} \quad (8)$$

provided that $(1 - p_0)(1 + \delta \varepsilon^r)^{2m} < 1$, where the last inequality bounds the moment generating function of τ/m by that of the geometric random variable.

It remains to show that $p_0 \rightarrow 1$ when $\varepsilon \rightarrow 0$, and then by plugging this in (8), we would obtain that $\mathbb{E}_{\text{is}}[X_{\text{is}}^2] \rightarrow \mu_0^2$ as $\varepsilon \rightarrow 0$, leading to VRE and completing the proof.

Under the assumptions of the theorem, the transition probabilities under IS satisfy

$$q(y, y') = \frac{p(y, y')v(y')}{\sum_{z \in \mathcal{Y}} p(y, z)v(z)} = \frac{p(y, y')\mu(y') + o(1)}{\sum_{z \in \mathcal{Y}} p(y, z)\mu(z) + o(1)} = \frac{p(y, y')\mu(y')}{\mu(y)} + o(1).$$

So under IS, for $y \in \mathcal{O}_0$, any path $\pi = (y = y_0 \rightarrow y_1 \rightarrow \dots \rightarrow y_k) \in \Pi(y)$ has probability

$$p_{\text{is}}(\pi) = \prod_{j=1}^k \frac{p(y_{j-1}, y_j) \mu(y_j)}{\mu(y_{j-1})} + o(1) = \frac{p(\pi)}{\mu(y)} + o(1).$$

Thus, the dominant paths have probability $\Theta(1)$ and the non-dominant paths have probability $o(1)$. This implies that $p_c(y) \rightarrow 1$, and then that $p_0 \rightarrow 1$, when $\varepsilon \rightarrow 0$. \square

We now give an example showing that VRE is sometimes satisfied even when the conditions of Theorem 1 are not satisfied. The example satisfies the weaker condition that $v(y)$ is the probability of one dominant path, which has the same order of magnitude as the sum of probabilities over all dominant paths. One might conjecture from this example that estimating $\mu(y)$ by a $v(y)$ having the right order of magnitude in terms of ε would be enough for VRE. We give another example showing that this is not true. That is, it shows that VRE may not hold if we do not split the probability correctly, at each step, among *all* the transitions that belong to a dominant path.

Example 2 In the example of Figure 1, suppose we take $v(y) = v_0(y)$, defined as the probability of the (single) most probable path from y to \mathcal{F} . This gives $v(0,0) = \varepsilon^3/2$, $v(0,1) = v(1,0) = \varepsilon^3$, $v(2,0) = v(0,2) = \varepsilon$, and $v(1,1) = \varepsilon^2$. Using this v in (2) leads to the transition probabilities $q(y, y')$ given in Table 1.

Table 1 Transition probabilities $q(y, y')$ for Example 2.

Origin state	Destination state	Probability
(0,0)	(0,1)	1/2
(0,0)	(1,0)	1/2
(0,1)	(0,2)	$\varepsilon/(1+\varepsilon)$
(0,1)	(1,1)	$1/(1+\varepsilon)$
(1,0)	(2,0)	$\varepsilon/(1+\varepsilon)$
(1,0)	(1,1)	$1/(1+\varepsilon)$
(1,1)	(0,1), (1,0)	$(1/2 - \varepsilon^2)\varepsilon/(2 + (1 - 2\varepsilon^2)\varepsilon)$
(1,1)	(1,2), (2,1)	$1/(2 + (1 - 2\varepsilon^2)\varepsilon)$
(0,2)	(0,1)	$\varepsilon^2(1-\varepsilon)/(1+\varepsilon^2(1-\varepsilon))$
(0,2)	(1,2)	$1/(1+\varepsilon^2(1-\varepsilon))$
(2,0)	(1,0)	$\varepsilon^2(1-\varepsilon)/(1+\varepsilon^2(1-\varepsilon))$
(2,0)	(2,1)	$1/(1+\varepsilon^2(1-\varepsilon))$

Let $\beta(y) = \text{Var}_{\text{is}}[X_{\text{is}}]$ when $Y_0 = y$, under these probabilities. By conditioning on the first transition, we find that these variances satisfy the system of equations

$$\beta(y) = \sum_{y' \in \mathcal{O}} q(y, y') \left(\frac{p(y, y')}{q(y, y')} \mu(y') \right)^2 - (\mu(y))^2 + \sum_{y' \in \mathcal{O}} q(y, y') \left(\frac{p(y, y')}{q(y, y')} \right)^2 \beta(y')$$

for all $y \in \mathcal{O}$. Solving this system gives, in particular,

$$\beta(0,0) = \frac{3 + 12\varepsilon + 12\varepsilon^2 + o(\varepsilon^2)}{1 - 4\varepsilon + 2\varepsilon^2 + o(\varepsilon^2)} \varepsilon^7 = \Theta(\varepsilon^7).$$

Hence, $\beta(0,0)/\mu^2(0,0) = \Theta(\varepsilon)$, meaning that we have VRE. On the other hand, the conditions of Theorem 1 are not satisfied in this case, because $v(1,1) = \varepsilon^2$ whereas $\mu(1,1) =$

$2\varepsilon^2 + o(\varepsilon^2)$. In fact we somehow got lucky to get VRE. For instance, $v(1, 1)$ wrongly estimates $\mu(1, 1)$ by a multiplicative constant, but the corresponding transition failure probabilities nevertheless have the good first order approximation, because an error by the same multiplicative constant appears in both the numerator and the denominator of the most important transition probabilities, so these errors cancel out. This error cancellation occurs at all nodes in this particular example! However, this does not occur in general, for other examples. \square

Suppose that for a given initial state y , there are different dominant paths, using different initial transitions, say (y, y_1) and (y, y_2) , for which $v(y_1) = \mu(y_1)(c_1 + o(1))$ and $v(y_2) = \mu(y_2)(c_2 + o(1))$, but with different multiplicative constants c_1 and c_2 . Using this v , the new probability

$$q(y, y_1) = \frac{p(y, y_1)v(y_1)}{p(y, y_1)v(y_1) + p(y, y_2)v(y_2)}(1 + o(1))$$

does not correspond asymptotically to the zero-variance probability

$$\frac{p(y, y_1)\mu(y_1)}{\mu(y)}(1 + o(1)).$$

The next example illustrates this.

Example 3 Consider again a system with $c = 2$ and $n_1 = n_2 = 2$, and where failure occurs whenever any two components are failed. The transition probabilities are depicted Figure 2. For each state y , define $v(y)$ as the probability of the most probable path from y to \mathcal{F} . This




Fig. 2 Transition probabilities for Example 3.

gives $v(0, 0) = \varepsilon/2$, $v(0, 1) = \varepsilon$, and $v(1, 0) = \varepsilon$, whereas the exact values are $\mu(0, 0) = (3\varepsilon + \varepsilon^2)/2$, $v(0, 1) = 2\varepsilon$, and $v(1, 0) = \varepsilon + \varepsilon^2$. Using this function v in (2) gives the transition probabilities $q(y, y')$ of Figure 3. It leads to $q((0, 0), (1, 0)) = 1/2$, whereas for the zero-variance change of measure, we should have

$$q((0, 0), (1, 0)) = \frac{(\varepsilon + \varepsilon^2)/2}{(\varepsilon + \varepsilon^2)/2 + \varepsilon} = \frac{1}{3}.$$

This asymptotic difference prevents VRE to hold. On the other hand, it is not difficult to verify that BRE holds: Just by enumerating all paths (there are not so many here), one can check that $\mu(0, 0) = (3\varepsilon + \varepsilon^2)/2$ and $\beta(0, 0) = \varepsilon^2/4 - \varepsilon^3/2 + \varepsilon^4/4$. \square




Fig. 3 Transitions probabilities based on v for Example 3.

In general, we have the following sufficient condition for BRE, which is weaker than the condition for VRE.

Theorem 2 *If Assumption 1 holds and $v(y) = \Theta(\mu(y))$, then we have BRE. This condition holds in particular if we take $v(y)$ as the sum of probabilities of any nonempty subset $\Pi_0(y)$ of $\Pi_d(y)$ (it could be a single path, or more).*

Proof For the first part, the proof mimics that of the previous theorem. There are positive constants δ'_1, δ'_2 and r' independent of ε and y , and c_y independent of ε , such that

$$c_y \mu(y) (1 - \delta'_1 \varepsilon^{r'}) \leq v(y) \leq c_y \mu(y) (1 + \delta'_2 \varepsilon^{r'})$$

for ε sufficiently small. Therefore, there exists δ', δ'' independent of ε and y such that

$$\frac{\tilde{v}(y)}{v(y')} \leq \frac{\sum_{y'' \in \mathcal{Y}} p(y, y'') c_{y''} \mu(y'')}{c_{y'} \mu(y')} (1 + \delta' \varepsilon^{r'})$$

for all y, y' , for ε sufficiently small. Let $c_M = \max_{y \in \mathcal{Y}} c_y$ and $c_m = \min_{y \in \mathcal{Y}} c_y$. Then

$$\frac{\tilde{v}(y)}{v(y')} \leq (1 + \delta' \varepsilon^{r'}) \frac{c_M \mu(y)}{c_m \mu(y')},$$

and for $Y_0 = y$,

$$\mathbb{E}_{\text{is}}[X_{\text{is}}^2] \leq \mu_0^2 \mathbb{E} \left[\left(\frac{c_M}{c_m} (1 + \delta \varepsilon^r) \right)^{2\tau} \right].$$

For the rest of the proof, we use a similar argument as in Theorem 1. We have that $q(y, y') = \Theta(p(y, y') \mu(y') / \mu(y))$, and so any path $\pi \in \Pi(y)$ has probability $p_{\text{is}}(\pi) = \Theta(p(\pi) / \mu(y))$ under IS. It follows that the dominant paths have probability $\Theta(1)$, the non-dominant paths have probability $o(1)$, and that $p_0 \rightarrow 1$ when $\varepsilon \rightarrow 0$. From this, by bounding again the moment generating function of τ/m , we obtain that $\mathbb{E}_{\text{is}}[X_{\text{is}}^2] / \mu_0^2$ remains bounded when $\varepsilon \rightarrow 0$.

For the second part, we have $v(y) = \sum_{\pi \in \Pi_0(y)} p(\pi)$. If $v(y) = o(\mu(y))$, then it means that no dominant path is taken into account in the computation of $v(y)$, which contradicts the fact that $\Pi_0(y)$ contains at least one dominant path. \square

5 Numerical examples

Example 4 We consider a system with $c = 3$ component types, with $n_1 = n_2 = n_3$. Each component has an exponentially distributed time to failure with rate λ_i for components of type i , where $\lambda_1 = \varepsilon$, $\lambda_2 = 1.5\varepsilon$, and $\lambda_3 = 2\varepsilon^2$, for some parameter ε . Any failed component has an exponentially distributed repair time with rate 1. Times to failure and repair times are all independent. The system is down whenever fewer than two components of any one type are operational. We want to estimate μ_0 , as explained earlier. We will experiment with different values of n_i and ε .

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. For example, for the initial state $y = \mathbf{0}$ and component type 1, the sample path π corresponds to $n_1 - 1$ failures of components of type 1 in succession, and its probability is

$$p(\pi) = \prod_{j=0}^{n_1-2} \frac{(n_1 - j)\lambda_1}{(n_1 - j)\lambda_1 + n_2\lambda_2 + n_3\lambda_3 + j}.$$

Again, this choice of v_0 is appropriate for the specific structure of \mathcal{F} considered here.

The methods we try and compare are BFB, a version of BLR proposed in [2] named simple BLR (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 $ZVA(v_0)$, $ZVA(v_1)$, $ZVA(v_2)$, respectively. For each parameter set that we have selected, Table 2 gives (in the third column) our best estimate of μ_0 , obtained from a very large number of simulation runs with our ZVA strategies (these numbers are accurate at least for the digits given in the table), the rough-cut approximation $v_0(\mathbf{0})$ of μ_0 , and the estimate obtained from $n = 2^{20}$ (approximately one million) independent simulation runs, by each of the five methods. A simulation run is defined as a sample path of the DTMC starting in state $\mathbf{0}$ and running until we reach the stopping time τ . Here it would be feasible to compute μ_0 numerically by building and solving a large system of linear equations giving $\mu(y)$ for all states y , because the number of states for $n_i = 12$ is only 2198, but we did not do it because it was much simpler (and accurate enough for our purpose) to use simulation. Table 3 shows the empirical variances computed from those n runs, for each method. One can easily get an idea of the square relative error (the variance divided by the square mean) by dividing these values by μ_0^2 , also given in the table.

Table 2 Parameter sets and estimates of $\mu(\mathbf{0})$ with each method, for Example 4.

n_i	ε	μ_0	$v_0(\mathbf{0})$	BFB	SBLR	$ZVA(v_0)$	$ZVA(v_1)$	$ZVA(v_2)$
3	0.001	2.6×10^{-3}	1.3×10^{-3}	2.7×10^{-3}	2.6×10^{-3}	2.6×10^{-3}	2.6×10^{-3}	2.6×10^{-3}
6	0.01	1.8×10^{-7}	3.4×10^{-8}	1.9×10^{-7}	(9.9×10^{-11})	1.8×10^{-7}	1.8×10^{-7}	1.8×10^{-7}
6	0.001	1.7×10^{-11}	3.4×10^{-12}	1.8×10^{-11}	(1.8×10^{-16})	1.7×10^{-11}	1.7×10^{-11}	1.7×10^{-11}
12	0.1	6.0×10^{-8}	3.2×10^{-9}	4.8×10^{-8}	1.3×10^{-8}	6.0×10^{-8}	6.2×10^{-8}	6.7×10^{-8}
12	0.001	3.9×10^{-28}	3.5×10^{-29}	(1.8×10^{-40})	(2.9×10^{-45})	3.9×10^{-28}	3.9×10^{-28}	3.9×10^{-28}

Table 3 Empirical variances based on $n = 2^{20}$ independent runs, for Example 4.

n_i	ε	α	μ_0^2	BFB	SBLR	ZVA(v_0)	ZVA(v_1)	ZVA(v_2)	RE(v_2)
3	0.001	0.906	6.8×10^{-6}	6.2×10^{-5}	8.0×10^{-3}	2.2×10^{-8}	6.6×10^{-9}	9.3×10^{-9}	0.04
6	0.01	0.903	3.2×10^{-14}	6.3×10^{-11}	(4.5×10^{-16})	2.0×10^{-14}	1.2×10^{-14}	7.7×10^{-15}	0.48
6	0.001	0.939	3.0×10^{-22}	8.8×10^{-19}	(2.0×10^{-26})	1.2×10^{-23}	1.1×10^{-23}	7.6×10^{-24}	0.16
12	0.1	0.851	3.6×10^{-15}	8.1×10^{-10}	1.7×10^{-10}	1.6×10^{-10}	2.9×10^{-10}	1.5×10^{-11}	64.5
12	0.001	0.963	1.5×10^{-55}	(3.2×10^{-74})	(3.5×10^{-84})	1.4×10^{-55}	9.3×10^{-56}	9.4×10^{-56}	0.78

In those tables, the entries in parentheses are empirical means and variances that clearly underestimate their exact counterparts by a large factor. For the empirical means, those entries 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. All these problematic entries are for the BFB and SBLR heuristics, and things generally worsen when μ_0 gets smaller and n_i gets larger. These heuristics were designed to cope with very low failure rates for the components (very small ε), but are not doing very 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 more significant improvement for many cases. 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 error. For example, $n = 61$ suffices to obtain 10% relative error. For $n_i = 12$ and $\varepsilon = 0.1$, on the other hand, the relative error with v_2 is $\sqrt{1.5 \times 10^{-11}} / (6.0 \times 10^{-8} \sqrt{n}) \approx 64.5 / \sqrt{n}$, so we need $n \approx (6450/x)^2$ to get $x\%$ relative error. The relative errors with v_2 are given in the table, under RE(v_2). Note that with $n_i = 12$, the relative error with our heuristic is much smaller with $\varepsilon = 0.001$ than with $\varepsilon = 0.1$. This is most likely due to the fact that our approximation v becomes too crude when ε increases.

Another observation is that α (used for v_1) increases, and apparently converges to 1, when $\varepsilon \rightarrow 0$ for fixed n_i , or when n_i increases while ε 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 v_1 .

Of course, comparing only the variances and neglecting the computing costs might be unfair. A standard practice is to look at the work-normalized variance, which is the variance

Table 4 Total computing times (in seconds) for $n = 2^{20}$ runs, for Example 4.

n_i	ε	BFB	SBLR	ZVA(v_0)	ZVA(v_1)	ZVA(v_2)
3	0.001	2.2	2.8	3.6	6.7	6.7
6	0.01	17.4	12	37	54	54
6	0.001	17	11	36	47	47
12	0.1	47	102	447	652	525
12	0.001	47	25	89	114	115

multiplied by the (expected) time required to compute the estimator. On the other hand, the computing time often depends very much on the implementation details, the computer, the programming language, and even the compiler. For this reason, we report the computing times separately. Table 4 reports the CPU times needed to make $n = 2^{20}$ independent simulation runs (this excludes the pilot runs made to estimate α used in v_1 and v_2). The simulations were performed in Java using SSJ [17] on a computer with a 2.00 GHz Intel Pentium processor. The CPU times are generally larger with our new methods than with BFB and SBLR, by a factor of about 3 or 4 in most cases and up to about 15 in the worst case (with ZVA(v_1) when $n_i = 12$ and $\varepsilon = 0.1$). This higher cost is more than compensated by the large gains in accuracy; in fact, BFB and SBLR only provide misleading mean and variance estimates in many cases. We note that the larger CPU times for ZVA come not only from the time required to compute the function $v(\cdot)$ at each step, but also (sometimes more importantly) from the fact that the sample paths are longer on average, because they reach \mathcal{F} more often. \square

Example 5 We now consider a simplified version of an example taken from [25], to which the BFB methodology is supposed to fit well. The system is comprised of two sets of processors with two processors per set, two sets of disk controllers with two controllers per set, and six clusters of disks with four disks per cluster. The failure rates for processors, controllers, and disks are 5×10^{-5} , 2×10^{-5} and 2×10^{-5} , respectively. The repair rate is 1 for each type of component. In each disk cluster, data is replicated, which means that the failure of a single disk does not provoke system's failure. The system is operational if all data is accessible from both processor types, meaning that at least one processor of each type, one controller of each set, and three disks of each cluster are operational. This can be modeled by $c = 10$ different types of components, to differentiate between the different sets of the same kind of component, and each n_i is 2 or 4. We use again 2^{20} independent sample paths for the estimations. Table 5 displays the empirical variance and the CPU times for the 2^{20} runs, for each method. The exact value is $\mu_0 \approx 5.6 \times 10^{-5}$ and the exponent in v_1 is $\alpha \approx 0.949$. We see that the ZVA methods work very nicely for this example as well, for all three choices of v . They outperform both BFB and SBLR. \square

Table 5 Empirical variances and CPU times for Example 5.

Method	BFB	SBLR	ZVA(v_0)	ZVA(v_1)	ZVA(v_2)
Variance	5.8×10^{-8}	1.3×10^{-4}	2.3×10^{-12}	1.0×10^{-12}	1.2×10^{-12}
CPU time	8.5	9.0	27	39	40

In the next example, we replace the approximation v_0 by a different one, to illustrate the idea that a good choice of v really depends on the structure of the model.

Example 6 We consider a system that fails when the total number of failed components (regardless of their type) reaches a given number (threshold). This kind of system makes sense if we classify the components in types according to their failure rate, and we need a minimal total number. In this case, we can easily find a good approximation for the probability that there is the required number of failures (in succession) before any repair occurs, and use it for v . A simple approximation is: the probability that the first transition is a failure (which is 1), multiplied by the probability that the second transition is a failure given that the first failed component is one with the largest failure rate, multiplied by the probability that the third transition is a failure given that the first two failed components have the largest failure rates, and so on. We call v_3 this approximation.

For a numerical illustration, we take a system comprised of 20 types of components numbered from 0 to 19, with 4 components of each type. All repair rates are assumed to be 1, but component's failure rates differ: type- i components have failure rate $\lambda_i = (1 + i/10)\varepsilon$ for $0 \leq i \leq 9$ and $\lambda_i = i\varepsilon^2/10$ for $10 \leq i \leq 19$, where $\varepsilon = 10^{-3}$. The system is failed whenever a total of 7 components are failed. Table 6 compares the results (estimates, variance, CPU times, and average number of steps per run) for BFB, SBLR, and ZVA(v_3). Again, the results obtained by ZVA are very accurate; the variance is reduced with respect to BFB by a factor 6×10^6 . The increase in CPU time is on the other hand very limited. To investigate more closely where this additional time comes from, the last line of Table 6 reports the average number of steps per run. It shows that SBLR has shorter paths on average, due to the fact that small paths from $\mathbf{0}$ to $\mathbf{0}'$ are more likely to happen. With ZVA(v_3), no sample path finishes at $\mathbf{0}'$, and the paths are longer on average for this reason, but the increase with respect to BFB is very small. Thus, in this example, the larger CPU times for ZVA come mostly from the required time to compute the approximation v_3 .

Table 6 Empirical estimates variancesCPU timesand average steps per run for Example 6.

Method	BFB	SBLR	ZVA(v_3)
Estimate	3.1×10^{-11}	(3.5×10^{-14})	3.0×10^{-11}
Variance	8.5×10^{-18}	(5.1×10^{-25})	1.3×10^{-24}
CPU time	11	19	97
Steps per run	7.0	4.7	7.2

6 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. The method relies on a reasonable approximation of the function μ and this would generally depend on the model. We have shown in this paper how very simple approximations of the zero-variance change of measure, in a reliability setting, can bring significant improvements. In our numerical examples, the proposed approach yields a low-variance estimator where all other previously

proposed algorithms miss the target. Further studies with larger and more complex systems are needed to find the practical limits of the method. We anticipate that difficulties may show up for large and complicated systems for which there is a large number of dominant paths whose sum of probabilities is hard to approximate. Additional heuristics might be needed for this type of situation.

Acknowledgements

This research has been supported by Grant OGP-0110050 and a Canada Research Chair to the first author, EuroNF Network of Excellence to the second author, and INRIA's associated team MOCQUASIN to both authors.

References

1. Ahamed, I., Borkar, V.S., Juneja, S.: Adaptive importance sampling for Markov chains using stochastic approximation. *Operations Research* **54**(3), 489–504 (2006)
2. Alexopoulos, C., Shultes, B.C.: Estimating reliability measures for highly-dependable Markov systems, using balanced likelihood ratios. *IEEE Transactions on reliability* **50**(3), 265–280 (2001)
3. Asmussen, S., Glynn, P.W.: *Stochastic Simulation*. Springer-Verlag, New York (2007)
4. Booth, T.E.: Exponential convergence for Monte Carlo particle transport? *Transactions of the American Nuclear Society* **50**, 267–268 (1985)
5. Booth, T.E.: Generalized zero-variance solutions and intelligent random numbers. In: *Proceedings of the 1987 Winter Simulation Conference*, pp. 445–451. IEEE Press (1987)
6. Cancela, H., Rubino, G., Tuffin, B.: MTTF estimation by Monte Carlo methods using Markov models. *Monte Carlo Methods and Applications* **8**(4), 312–341 (2002)
7. Carrasco, J.A.: Failure distance-based simulation of repairable fault tolerant systems. In: *Proceedings of the 5th International Conference on Modeling Techniques and Tools for Computer Performance Evaluation*, pp. 351–365. Elsevier, Amsterdam (1992)
8. Gertsbakh, I.B.: Asymptotic methods in reliability theory: A review. *Advances in Applied Probability* **16**, 147–175 (1984)
9. Glynn, P.W., Iglehart, D.L.: Importance sampling for stochastic simulations. *Management Science* **35**, 1367–1392 (1989)
10. Goyal, A., Shahabuddin, P., Heidelberger, P., Nicola, V.F., Glynn, P.W.: A unified framework for simulating Markovian models of highly reliable systems. *IEEE Transactions on Computers* **C-41**, 36–51 (1992)
11. Heidelberger, P.: Fast simulation of rare events in queueing and reliability models. *ACM Transactions on Modeling and Computer Simulation* **5**(1), 43–85 (1995)
12. Juneja, S., Shahabuddin, P.: Fast simulation of Markov chains with small transition probabilities. *Management Science* **47**(4), 547–562 (2001)
13. Juneja, S., Shahabuddin, P.: Splitting-based importance sampling algorithms for fast simulation of Markov reliability models with general repair policies. *IEEE Transactions on Reliability* **50**(3), 235–245 (2001)
14. Juneja, S., Shahabuddin, P.: Rare event simulation techniques: An introduction and recent advances. In: S.G. Henderson, B.L. Nelson (eds.) *Simulation, Handbooks in Operations Research and Management Science*, pp. 291–350. Elsevier, Amsterdam, The Netherlands (2006). Chapter 11
15. Kollman, C., Baggerly, K., Cox, D., Picard, R.: Adaptive importance sampling on discrete Markov chains. *Annals of Applied Probability* **9**(2), 391–412 (1999)
16. L'Ecuyer, P., Blanchet, J.H., Tuffin, B., Glynn, P.W.: Asymptotic robustness of estimators in rare-event simulation. *ACM Transactions on Modeling and Computer Simulation* (2008). To appear
17. L'Ecuyer, P., Buist, E.: Simulation in Java with SSJ. In: M.E. Kuhl, N.M. Steiger, F.B. Armstrong, J.A. Joines (eds.) *Proceedings of the 2005 Winter Simulation Conference*, pp. 611–620. IEEE Press, Piscataway, NJ (2005)
18. L'Ecuyer, P., Tuffin, B.: 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*, p. 7 pages. ACM Publications, Pisa, Italy (2006)

19. L'Ecuyer, P., Tuffin, B.: Effective approximation of zero-variance simulation in a reliability setting. In: Proceedings of the 2007 European Simulation and Modeling Conference, pp. 48–54. EUROSIS, Ghent, Belgium (2007)
20. L'Ecuyer, P., Tuffin, B.: Approximate zero-variance simulation. In: Proceedings of the 2008 Winter Simulation Conference, pp. 170–181. IEEE Press (2008)
21. Nakayama, M.K.: General conditions for bounded relative error in simulations of highly reliable Markovian systems. *Advances in Applied Probability* **28**, 687–727 (1996)
22. Nakayama, M.K., Shahabuddin, P.: 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 (2004)
23. Nicola, V.F., Shahabuddin, P., Nakayama, M.K.: Techniques for fast simulation models of highly dependable systems. *IEEE Transactions on Reliability* **50**(3), 246–264 (2001)
24. Rubinstein, R.Y., Shapiro, A.: *Discrete Event Systems: Sensitivity Analysis and Stochastic Optimization by the Score Function Method*. Wiley, New York (1993)
25. Shahabuddin, P.: Fast transient simulation of Markovian models of highly dependable systems. *Performance Evaluation* **20**, 267–286 (1994)
26. Shahabuddin, P.: Importance sampling for the simulation of highly reliable Markovian systems. *Management Science* **40**(3), 333–352 (1994)
27. Shahabuddin, P., Nicola, V.F., Heidelberger, P., Goyal, A., Glynn, P.W.: Variance reduction in mean time to failure simulations. In: Proceedings of the 1988 Winter Simulation Conference, pp. 491–499. IEEE Press (1988)
28. Spanier, J.: A new multistage procedure for systematic variance reduction in Monte Carlo. *SIAM Journal on Numerical Analysis* **8**, 548–554 (1971)
29. Tuffin, B.: On numerical problems in simulations of highly reliable Markovian systems. In: Proceedings of the 1st International Conference on Quantitative Evaluation of SysTems (QEST), pp. 156–164. IEEE CS Press, University of Twente, Enschede, The Netherlands (2004)