

Approximate Zero-Variance Importance Sampling for Static Network Reliability Estimation

Pierre L'Ecuyer, Gerardo Rubino, Samira Saggadi and Bruno Tuffin

Abstract—We propose a new Monte Carlo method, based on dynamic importance sampling, to estimate the probability that a given set of nodes is connected in a graph (or network) where each link is failed with a given probability. The method generates the link states one by one, using a sampling strategy that approximates an ideal zero-variance importance sampling scheme. The approximation is based on minimal cuts in subgraphs. In an asymptotic rare-event regime where failure probability becomes very small, we prove that the relative error of our estimator remains bounded, and even converges to 0 under additional conditions, when the unreliability of individual links converges to 0. The empirical performance of the new sampling scheme is illustrated by examples.

Index Terms—Monte Carlo methods, network reliability, variance reduction.

ACRONYMS

BRE	bounded relative error
CLT	central-limit theorem
IS	importance sampling
MC	Monte Carlo
VRE	vanishing relative error

NOTATION

\mathcal{N}	set of nodes
\mathcal{L}	set of links $\{1, \dots, m\}$
m	number of links
\mathcal{G}	undirected graph $(\mathcal{N}, \mathcal{L})$
\mathcal{K}	set of target (destination) nodes
q_i	unreliability of link i
u	graph unreliability
X_i	(random) state of link i : 1 if operational, 0 otherwise
X	(X_1, \dots, X_m) : (random) configuration of the graph
x	$= (x_1, \dots, x_m)$: one realization of X

$\phi(X)$	structure function associated with the \mathcal{K} -terminal unreliability: $\phi(X) = 1$ if the nodes in \mathcal{K} are not connected, 0 otherwise
ϵ	rarity parameter
n	number of independent replications
$X^{(j)}$	j -th independent replication of X
\mathbb{P}	original probability law of the network
\mathbb{E}	expectation under \mathbb{P}
\tilde{q}_i	IS unreliability of link i
$\tilde{\mathbb{P}}$	IS probability law
$\tilde{\mathbb{E}}$	expectation under $\tilde{\mathbb{P}}$
$u_i(\cdot)$	$u_i(x_1, \dots, x_{i-1})$ is the unreliability of the graph when (x_1, \dots, x_{i-1}) are fixed
$\hat{u}_i(\cdot)$	$\hat{u}_i(x_1, \dots, x_{i-1})$ is an approximation of $u_i(x_1, \dots, x_{i-1})$
$L_i(x_i)$	$= x_i(1 - q_i)/(1 - \tilde{q}_i) + (1 - x_i)q_i/\tilde{q}_i$: likelihood ratio associated with step i
$L(x)$	$\prod_{i=1}^m L_i(x_i)$: global likelihood ratio
o	For two functions $f, g : (0, \infty) \rightarrow \mathbb{R}$, we say that $f(\epsilon) = o(g(\epsilon))$ if $\lim_{\epsilon \rightarrow 0^+} f(\epsilon)/g(\epsilon) = 0$
O	$f(\epsilon) = O(g(\epsilon))$ if $ f(\epsilon) \leq c_1 g(\epsilon)$ for some constant $c_1 > 0$ for all ϵ sufficiently small
\underline{O}	$f(\epsilon) = \underline{O}(g(\epsilon))$ if $ f(\epsilon) \geq c_2 g(\epsilon)$ for some constant $c_2 > 0$ for all ϵ sufficiently small
Θ	$f(\epsilon) = \Theta(g(\epsilon))$ if $f(\epsilon) = \underline{O}(\epsilon^d)$ and $f(\epsilon) = O(\epsilon^d)$

I. INTRODUCTION

THE STATIC network reliability problem, which consists in computing the probability that a given set of nodes in a graph is connected when each individual link is failed with a given probability, occurs in a wide range of applications [3], [6], [10]. Examples are easily found in telecommunications, where we may want to assess the probability that a selected group of nodes (which can be just a pair) can communicate, or in power supply systems, where we may want to estimate the risk that electricity is not provided to certain nodes, or in transportation systems, where links represent the roads and are subject to damages.

P. L'Ecuyer is with the Département d'Informatique et de Recherche Opérationnelle, Université de Montréal, C.P. 6128, Succ. Centre-Ville, Montréal, H3C 3J7, Canada. E-mail: lecuyer@iro.umontreal.ca.

G. Rubino, S. Saggadi and B. Tuffin is with INRIA Rennes Bretagne Atlantique, Campus Universitaire de Beaulieu, 35042 Rennes Cedex, FRANCE. E-mail: rubino,ssaggadi,btuffin@irisa.fr.

In all these settings, a disconnected set of nodes may have critical implications, either financially or security-wise, and an accurate reliability estimate is needed. For large graphs, exact computation of the *unreliability* (the probability of *system failure*, which happens when the considered set of nodes is not connected) is impractical, since the problem is NP-hard in general; for all known algorithms the required computational effort increases exponentially with the number of links [2], [9], [24].

Monte Carlo methods can estimate unreliability for very large graphs. In its crude form, the Monte Carlo algorithm samples n independent realizations of the graph, and estimates the unreliability u by the proportion of these n realizations for which the selected nodes are not connected. However, this simple approach becomes useless when the network is highly reliable (u is very close to 0), because one would need an excessively large value of n to obtain enough realizations where the nodes are not connected. For example, if $u = 10^{-10}$, we expect one such realization per ten billion runs on average and we need many more than that to estimate u even with only a single digit of accuracy. Such small values of u are commonplace in critical applications.

Special variance reduction techniques have been developed to address this type of *rare-event* estimation problem. Some of them are based on the *importance sampling* (IS) principle, which consists in changing the sampling probabilities of the links so that system failure is no longer a rare event, and multiplying the original estimator by an appropriate likelihood ratio to recover unbiasedness [1], [15], [18], [26]. Specific IS methods have been proposed for estimating the unreliability in static graphs; see [14], [23], or [6] and the references given there. The main difficulty with IS is to find an effective way of changing the sampling probabilities. If the probabilities of realizations that lead to failure are inflated unevenly, or if some of them are decreased too much, then the likelihood ratio may have a huge variance and this may result in a badly behaved estimator even if system failure is no longer a rare event [18]. This pitfall is hard to avoid in general and is often amplified when the failure probabilities approach 0.

Robustness of estimators in this type of situation is usually studied by examining the behavior of the relative error (the standard deviation divided by the mean) as a function of rarity [11], [12], [15], [17], [18]. Two characterizations examined in this paper are the widely-used *bounded relative error* (BRE) property, which holds when the relative error remains bounded regardless of the rarity [12], [15], and the stronger *vanishing relative error* (VRE) property, which says that the relative error converges to zero when the unreliability u goes to zero [17]. These properties are particularly relevant in situations where u

is very small: the relative width of a confidence interval on u based on the central-limit theorem (CLT) for a fixed sample size n when $u \rightarrow 0$ remains bounded if BRE holds and converges to zero if VRE holds [17].

The aim of this paper is to propose and study a new way of applying IS to address the static network reliability problem. The proposed IS strategy is based on a (theoretical) zero-variance IS scheme, which exists in principle whenever the original estimator cannot take negative values [1], [15], [18], [19]. We represent the sampling process of all links of the graph by a Markov chain that determines one new link at each step, and whose state is the states of all links that have already been sampled. We show how the transition probabilities of this Markov chain can be changed (at least in principle) to obtain a zero-variance IS estimator of u . Under this ideal IS scheme, the system always fails and the likelihood ratio is always equal to u , so the IS estimator is a constant, with probability 1. Unfortunately, implementing this scheme directly is not practically possible unless we can compute everything exactly in the first place. But it can be approximated, and this is what we do here.

We propose an approximation based on minimal cuts having (relatively) high failure probability in the subgraph that remains after removing the links already known to be failed, and enforcing the states of the links known to be operational, at each step of the Markov chain. These cuts are used to approximate the unreliability conditional on the current state, at each step. We prove that our estimators enjoy the BRE property, and that under an additional condition on the approximation of the conditional unreliability at each step, the VRE property is also verified. This type of asymptotic analysis as a function of a rarity parameter has been done extensively for queuing systems [4], [12] and for highly reliable Markovian systems [15], [20], [21], [25], [28]. For static network reliability, as far as we know, the only proposed estimator with the BRE property is the path-based estimator described in [7], [8], which was shown to have this property but under more restrictive conditions than those made here, and no estimator with the VRE property has been proposed so far. Our numerical results on various examples show that the method introduced here performs quite well even on large graphs with extremely small unreliabilities, and the VRE property can be observed empirically.

The paper is organized as follows. Section II presents the mathematical model. In Section III, we recall the *crude* Monte Carlo method, we show its inefficiency, and we explain how IS can be applied in general for this model. In Section IV, we construct an ideal (zero-variance) IS estimator. Its implementation would require the knowledge of the conditional unreliability (function) given the status

of arbitrary subsets of links, which is how course utopian, but an estimate of this conditional unreliability can be plugged into the formula in place of the exact one. We prove that this yields an estimator with BRE under appropriate conditions on the conditional reliability estimate, and with VRE under an additional condition. In Section V, we propose a specific (crude) approximation method for this conditional unreliability, based on the probability of a minimal cut. We prove that the resulting IS estimator always has the BRE property, and the VRE property under additional conditions that are often verified in our examples. In Section VI we illustrate numerically the efficiency of the method on several examples. Section VII provides a conclusion and discusses issues worthy of future work.

II. MODEL

We consider an undirected graph $\mathcal{G} = (\mathcal{N}, \mathcal{L})$ where \mathcal{N} is the set of nodes and $\mathcal{L} = \{1, \dots, m\}$ the set of links. We focus here on a static model, where time plays no explicit role. Nodes never fail but links are subject to independent failures, link $i \in \mathcal{L}$ failing with probability q_i , where $0 < q_i < 1$. A *configuration* [24] of the graph is given by the random vector

$$X = (X_1, \dots, X_m)$$

where for all $i \in \mathcal{L}$, $X_i = 1$ if link i works, and $X_i = 0$ if link i is failed. By retaining only the set of operational links $\mathcal{L}' \subseteq \mathcal{L}$, we end up with a (random) partial graph $\mathcal{G}' = (\mathcal{N}, \mathcal{L}')$ of \mathcal{G} . Knowing the vector X is equivalent to knowing \mathcal{G}' .

Our goal is to estimate the probability u that a given set of nodes \mathcal{K} are connected in the random graph \mathcal{G}' , or equivalently for the random configuration X . Formally, we define $\phi(X)$ by $\phi(X) = 1$ if the set of nodes \mathcal{K} is *not* connected in \mathcal{G}' , that is, they do not belong to the same connected component when the configuration is X , and $\phi(X) = 0$ otherwise. The expectation $u = \mathbb{E}[\phi(X)]$ is the \mathcal{K} -terminal unreliability. (For convenience, here we define $\phi(X)$ differently from the structure function that usually represents the reliability and is equal to $1 - \phi(X)$ [6]). The most frequent case is when estimating the *two-terminal* or *source-to-terminal unreliability*, where \mathcal{K} is comprised of only two nodes.

This unreliability metric can be written as

$$\begin{aligned} u &= \mathbb{E}[\phi(X)] = \sum_{x \in \{0,1\}^m} \phi(x) \mathbb{P}[X = x] \\ &= \sum_{x \in \{0,1\}^m} \phi(x) \prod_{i=1}^m (q_i(1 - x_i) + (1 - q_i)x_i) \end{aligned}$$

where $x = (x_1, \dots, x_m)$. The state space (the set of all possible configurations) is of cardinality 2^m . This means that computing u directly from this formula requires a time that increases exponentially with m , the number of links. As we said earlier, this problem is NP-hard in general [2], so for large graphs, approximation techniques are required, and Monte Carlo simulation becomes the method of choice when the graph is large enough.

We are interested mainly in the situation where the failure probabilities q_i of the individual links are very close to 0. We will capture this by studying the performance of our algorithm in an asymptotic regime where the q_i 's are polynomial functions of a rarity parameter $\epsilon \ll 1$, as done for example in [17], [22], [29]. More specifically, we shall assume that for each $i \in \mathcal{L}$, there are positive constants a_i and b_i (independent of ϵ) such that

$$q_i = a_i \epsilon^{b_i}. \quad (1)$$

Under this assumption, it is easy to verify (because the configuration space is finite and the probability of any configuration is a polynomial in ϵ) that the system unreliability (which depends on ϵ) is

$$u = u(\epsilon) = \Theta(\epsilon^c) \quad (2)$$

for some constant $c > 0$ (for similar remarks, see [9] or [24], for instance).

III. MONTE CARLO SIMULATION

The crude *Monte Carlo* (MC) method simulates n independent realizations of X , say $X^{(1)}, \dots, X^{(n)}$, and estimates u by the proportion of those realizations where the nodes in \mathcal{K} are not connected:

$$\bar{U}_{\text{MC}}^{(n)} = \frac{1}{n} \sum_{i=1}^n \phi(X^{(i)}).$$

When n is large enough, thanks to the central-limit theorem (CLT), we can in principle compute a confidence interval on u , say at confidence level α , by assuming that $\bar{U}_{\text{MC}}^{(n)}$ is approximately normally distributed. This gives the confidence interval

$$\left(\bar{U}_{\text{MC}}^{(n)} - c_\alpha S^{(n)} / \sqrt{n}, \bar{U}_{\text{MC}}^{(n)} + c_\alpha S^{(n)} / \sqrt{n} \right) \quad (3)$$

with confidence $1 - \alpha$, where c_α is the $1 - \alpha/2$ quantile of the standard normal distribution (with mean 0 and variance 1) and $(S^{(n)})^2 = \bar{U}_{\text{MC}}^{(n)}(1 - \bar{U}_{\text{MC}}^{(n)})n/(n-1)$ is the empirical variance of $\phi(X^{(1)}), \dots, \phi(X^{(n)})$.

The expected relative half-width of this confidence interval is

$$c_\alpha \frac{(\text{Var}[\phi(X)]/n)^{1/2}}{\mathbb{E}[\phi(X)]} = c_\alpha \left(\frac{1 - u}{un} \right)^{1/2},$$

which increases to ∞ when $u \rightarrow 0$ for fixed n . To keep it under control, we need $n = \underline{O}(1/u)$. This means that in the asymptotic regime defined by (1), we would need $n = n(\epsilon) = \underline{O}(\epsilon^{-c})$.

Our aim in this paper is to replace $\phi(X)$ by an alternative unbiased estimator Y with smaller variance. Numerous techniques have been designed to do that; we refer the reader to [6] for an overview. For fixed n and $\phi(X)$ replaced by Y , the half-width of the confidence interval above is proportional to the *relative error* of Y , defined as $\text{RE}[Y] = (\text{Var}[Y])^{1/2}/\mathbb{E}[Y]$. For $Y = \phi(X)$, the relative error is $\Theta(u^{-1/2})$. We are interested in having an estimator Y with *bounded relative error* (BRE), which means that $\text{RE}[Y]$ remains bounded when $u \rightarrow 0$, or in our framework when $\epsilon \rightarrow 0$. It would be even better if Y has the stronger property of *vanishing relative error* (VRE), which means that $\text{RE}[Y] \rightarrow 0$ when $\epsilon \rightarrow 0$. To our knowledge, no VRE estimator has been constructed so far for this problem.

IV. IMPORTANCE SAMPLING AND ZERO-VARIANCE APPROXIMATION

A. Importance Sampling

IS consists in replacing the probabilities of the 2^m possible configurations X by another set of probabilities [1], [15], [18]. If we denote the original probabilities by \mathbb{P} and the new ones by $\tilde{\mathbb{P}}$, we have

$$\begin{aligned} u = \mathbb{E}[\phi(X)] &= \sum_{x \in \{0,1\}^m} \phi(x) \mathbb{P}[X = x] \\ &= \sum_{x \in \{0,1\}^m} \phi(x) L(x) \tilde{\mathbb{P}}[X = x] \end{aligned}$$

where

$$L(x) = \mathbb{P}[X = x] / \tilde{\mathbb{P}}[X = x],$$

provided that $\tilde{\mathbb{P}}[X = x] > 0$ whenever $\phi(x) \mathbb{P}[X = x] > 0$. This $L(x)$ is the *likelihood ratio* of the old and new probabilities for outcome x . The unreliability can then be rewritten as

$$u = \tilde{\mathbb{E}}[\phi(X)L(X)] \quad (4)$$

where $\tilde{\mathbb{E}}[\cdot]$ is the expectation under $\tilde{\mathbb{P}}$. We thus have the following unbiased IS estimator of u :

$$\bar{U}_{\text{IS}}^{(n)} = \frac{1}{n} \sum_{j=1}^n \phi(X^{(j)}) L(X^{(j)}),$$

where $X^{(1)}, \dots, X^{(n)}$ are independent copies of X all distributed according to $\tilde{\mathbb{P}}$. A confidence interval on u can be computed as in (3), but with $\bar{U}_{\text{MC}}^{(n)}$ and $(S^{(n)})^2$ replaced by the sample mean and the sample variance of the $\phi(X^{(j)})L(X^{(j)})$.

It is well known that the optimal change of probabilities in this setting inflates all the probabilities by a factor proportional to $\phi(x)$ [1], [18]. This gives

$$\tilde{\mathbb{P}}[X = x] = \phi(x) \mathbb{P}[X = x] / u$$

for all $x \in \{0,1\}^m$. That is, the realizations where the system does not fail are given zero probability, and the probabilities of the other realizations are rescaled accordingly (divided by u). Under these new probabilities, with probability 1, the system fails and the likelihood ratio is u , so the estimator always takes the value u . Thus, this estimator has zero variance. However, implementing this IS scheme requires knowing u in the first place.

B. A Sequential Version of Zero-Variance Importance Sampling

We now reformulate the sampling of X as a Markov chain process and we define a zero-variance IS scheme for that process by adapting the techniques described in [16], [18], [19]. The link states X_1, \dots, X_m are generated successively in this order. At step i of the chain, we generate the coordinate X_i of X , over the set $\{0,1\}$. Under \mathbb{P} , these probabilities are $\mathbb{P}[X_i = 0] = 1 - \mathbb{P}[X_i = 1] = q_i$, and the X_i 's are independent. Under IS, we will change the probability q_i at each step and let it depend on the previously generated values X_1, \dots, X_{i-1} . We will see that by doing this optimally, one can obtain a zero-variance estimator.

We define

$$u_i(x_1, \dots, x_{i-1}) = \mathbb{E}[\phi(X) | X_1 = x_1, \dots, X_{i-1} = x_{i-1}],$$

which represents the unreliability of graph \mathcal{G}' conditional on the states of links 1 to $i-1$, which are already determined before step i . The unconditional unreliability of the graph can be written as $u = u_1(\emptyset)$. We will often use the fact that

$$\begin{aligned} u_i(x_1, \dots, x_{i-1}) &= q_i u_{i+1}(x_1, \dots, x_{i-1}, 0) \\ &\quad + (1 - q_i) u_{i+1}(x_1, \dots, x_{i-1}, 1). \end{aligned} \quad (5)$$

Suppose that for $i = 1, \dots, m$, we replace q_i by

$$\begin{aligned} \tilde{q}_i &\stackrel{\text{def}}{=} \tilde{\mathbb{P}}[X_i = 0 | X_1 = x_1, \dots, X_{i-1} = x_{i-1}] \\ &= \frac{q_i u_{i+1}(x_1, \dots, x_{i-1}, 0)}{u_i(x_1, \dots, x_{i-1})}. \end{aligned} \quad (6)$$

This gives

$$1 - \tilde{q}_i = \frac{(1 - q_i) u_{i+1}(x_1, \dots, x_{i-1}, 1)}{u_i(x_1, \dots, x_{i-1})}. \quad (7)$$

Using (6) and (7), we find that for a given realization $x = (x_1, \dots, x_m)$, the likelihood ratio for step i is

$$\begin{aligned} L_i(x_i) &= x_i \frac{1 - q_i}{1 - \tilde{q}_i} + (1 - x_i) \frac{q_i}{\tilde{q}_i} \\ &= \frac{u_i(x_1, \dots, x_{i-1})}{u_{i+1}(x_1, \dots, x_{i-1}, x_i)}. \end{aligned}$$

The overall likelihood ratio is

$$L(x) = \prod_{i=1}^m L_i(x_i)$$

and the final estimator is $\phi(X)L(X)$. The following theorem says that this sequential (state-dependent) IS scheme gives a zero-variance estimator. It is a consequence of the general results on zero-variance sampling for Markov chains, given in [16], [18], [19]. We give a direct proof from first principles.

Theorem 1: Under the sequential IS algorithm where the probabilities q_i are replaced by the conditional probabilities \tilde{q}_i defined in (6), with probability 1, we have $\phi(X) = 1$ and $L(X) = u$, so the IS estimator has zero variance.

Proof: We first show that $\tilde{\mathbb{P}}[\phi(X) = 1] = 1$. Let x be a configuration for which the system works, that is, $\phi(x) = 0$, and suppose that k is the index of the first coordinate such that $\phi(x_1, \dots, x_k, \dots) = 0$ whatever be the values of (x_{k+1}, \dots, x_m) , for the given (fixed) values of x_1, \dots, x_k . In this case, $u_{k+1}(x_1, \dots, x_k) = 0$, which implies from (6) that conditional on x_1, \dots, x_{k-1} , this particular x_k has zero probability. Therefore, every configuration x with $\phi(x) = 0$ has zero probability. We also have

$$\begin{aligned} L(x) &= \prod_{i=1}^m L_i(x_i) = \prod_{i=1}^m \frac{u_i(x_1, \dots, x_{i-1})}{u_{i+1}(x_1, \dots, x_{i-1}, x_i)} \\ &= \frac{u_1(\emptyset)}{u_{m+1}(x_1, \dots, x_m)} = \frac{u}{u_{m+1}(x_1, \dots, x_m)}. \end{aligned}$$

For any configuration x with $\phi(x) = 1$, we have $u_{m+1}(x_1, \dots, x_m) = \phi(x_1, \dots, x_m) = 1$, and therefore $L(x) = u$. This means that with probability 1, we have $\phi(X) = 1$ and $\phi(X)L(X) = u$. ■

Implementing this zero-variance IS scheme would require the perfect knowledge of all the functions u_i , and in particular the knowledge of $u_1(\emptyset) = u$, so it is not practical. In the next subsection we study the situation where each u_i is replaced by some approximation \hat{u}_i . One way to obtain such an approximation will be proposed in Section V.

C. Zero-Variance Approximation and Robustness Properties

The idea of zero-variance approximation is to replace the functions $u_i(\cdot)$ in (6) by some easy-to-compute approximations $\hat{u}_i(\cdot)$. This gives

$$\begin{aligned} \tilde{q}_i &= \tilde{\mathbb{P}}[X_i = 0] \\ &= \frac{q_i \hat{u}_{i+1}(x_1, \dots, x_{i-1}, 0)}{q_i \hat{u}_{i+1}(x_1, \dots, x_{i-1}, 0) + (1 - q_i) \hat{u}_{i+1}(x_1, \dots, x_{i-1}, 1)}. \end{aligned} \quad (8)$$

The intuition is that if $\hat{u}_{i+1}(\cdot)$ is not too far from $u_{i+1}(\cdot)$ for each i , the sampling will be done with probabilities that are not too far from the optimal ones, and the variance can then be reduced by large factors. Observe that the network unreliability u will not change if we modify the index of the edges of the graph, but the type of change of measure proposed here does depend on that index.

The next two theorems characterize that intuition in an asymptotic regime where $\epsilon \rightarrow 0$, while the graph remains fixed. They give sufficient conditions on the approximations $\hat{u}_i(\cdot)$ for the BRE or VRE property to hold.

Theorem 2: Suppose that for each i and $(x_1, \dots, x_i) \in \{0, 1\}^i$, $1 \leq i \leq m$, there is a constant $a_{i+1}(x_1, \dots, x_i)$ independent of ϵ such that

$$\begin{aligned} \hat{u}_{i+1}(x_1, \dots, x_i) &= a_{i+1}(x_1, \dots, x_i) u_{i+1}(x_1, \dots, x_i) \\ &\quad + o(u_{i+1}(x_1, \dots, x_i)). \end{aligned} \quad (9)$$

Then the estimator provided by the IS scheme with change of probabilities defined by (8) has the BRE property.

Proof: We need to show that under the IS scheme, $\text{Var}[\phi(X)L(X)] = O(u^2)$, or equivalently that $\tilde{\mathbb{E}}[\phi(X)L^2(X)] = O(u^2)$. From (5), we see that

$$(1 - q_i) u_{i+1}(x_1, \dots, x_{i-1}, 1) = O(u_i(x_1, \dots, x_{i-1}))$$

and

$$q_i u_{i+1}(x_1, \dots, x_{i-1}, 0) = O(u_i(x_1, \dots, x_{i-1})),$$

with at least one of these two terms being $\Theta(u_i(x_1, \dots, x_{i-1}))$. Using this and given that $a_{i+1}(x_1, \dots, 1)$ and $a_{i+1}(x_1, \dots, 0)$ are constants independent of ϵ , one can verify that there is a positive constant $b_i(x_1, \dots, x_{i-1})$ independent of ϵ and such that

$$\begin{aligned} &(1 - q_i) \hat{u}_{i+1}(x_1, \dots, x_{i-1}, 1) + q_i \hat{u}_{i+1}(x_1, \dots, x_{i-1}, 0) \\ &= (1 - q_i) a_{i+1}(x_1, \dots, 1) u_{i+1}(x_1, \dots, x_{i-1}, 1) \\ &\quad + q_i a_{i+1}(x_1, \dots, 0) u_{i+1}(x_1, \dots, x_{i-1}, 0) \\ &\quad + o(u_{i+1}(x_1, \dots, 1) + u_{i+1}(x_1, \dots, 0)) \\ &= b_i(x_1, \dots, x_{i-1}) u_i(x_1, \dots, x_{i-1}) + o(u_i(x_1, \dots, x_{i-1})). \end{aligned}$$

For example, if

$$(1 - q_i) u_{i+1}(x_1, \dots, x_{i-1}, 1) = o(u_i(x_1, \dots, x_{i-1})),$$

then

$$\begin{aligned} & a_{i+1}(x_1, \dots, 0)u_i(x_1, \dots, x_{i-1}) \\ &= q_i a_{i+1}(x_1, \dots, 0)u_{i+1}(x_1, \dots, x_{i-1}, 0) \\ & \quad + o(u_{i+1}(x_1, \dots, x_{i-1})), \end{aligned}$$

so we have $b_i(x_1, \dots, x_{i-1}) = a_{i+1}(x_1, \dots, 0)$. The other cases are similar.

The likelihood ratio for the sampling of link i is then

$$\begin{aligned} L_i(X_i) &= \frac{(1 - q_i)\widehat{u}_{i+1}(X_1, \dots, X_{i-1}, 1)}{\widehat{u}_{i+1}(X_1, \dots, X_{i-1}, X_i)} \\ & \quad + \frac{q_i\widehat{u}_{i+1}(x_1, \dots, x_{i-1}, 0)}{\widehat{u}_{i+1}(X_1, \dots, X_{i-1}, X_i)} \\ &= \frac{b_i(X_1, \dots, X_{i-1})}{a_{i+1}(X_1, \dots, X_i)} \frac{u_i(X_1, \dots, X_{i-1})}{u_{i+1}(X_1, \dots, X_i)} + o(1). \end{aligned}$$

Let

$$d = \max_{\{(x_1, \dots, x_i) : 1 \leq i \leq m\}} \frac{b_i(x_1, \dots, x_{i-1})}{a_{i+1}(x_1, \dots, x_i)}.$$

We have $d < \infty$, because this maximum is over a finite set, and the likelihood ratio for the whole sample path satisfies

$$\begin{aligned} L(X) &= \prod_{i=1}^m L_i(X_i) \\ &\leq \prod_{i=1}^m d \frac{u_i(X_1, \dots, X_{i-1})}{u_{i+1}(X_1, \dots, X_{i-1}, X_i)} + o(u) \\ &\leq d^m u + o(u). \end{aligned}$$

Therefore,

$$\widetilde{\mathbb{E}}[\phi(X)L^2(X)] \leq \widetilde{\mathbb{E}}[L^2(X)] \leq d^{2m}u^2 + o(u^2),$$

due to the finite sum in the expectation, which completes the proof. \blacksquare

In the next theorem, we additionally require that for all configurations x that are not rare under IS, at each step i of the method, the unreliability estimations lead to a sampling probability asymptotically equivalent to the zero-variance one, and we show that VRE holds under these conditions. We define

$$\mathcal{S}_1 = \{x \in \{0, 1\}^m : \phi(x) = 1 \text{ and } \widetilde{\mathbb{P}}[X = x] = \Theta(1)\}$$

and

$$\mathcal{S}_0 = \{x \in \{0, 1\}^m : \phi(x) = 1 \text{ and } \widetilde{\mathbb{P}}[X = x] = o(1)\}.$$

These are the sets of configurations where the system fails and which are no longer rare and still rare, respectively, under IS. Our additional conditions for VRE will involve only the configurations $x \in \mathcal{S}_1$.

Theorem 3: Let the assumptions of Theorem 2 hold, and suppose also that for all $x = (x_1, \dots, x_m) \in \mathcal{S}_1$ and for each i , one of the following three conditions is satisfied:

$$\begin{aligned} & \frac{\widehat{u}_{i+1}(x_1, \dots, x_{i-1}, 1)}{u_{i+1}(x_1, \dots, x_{i-1}, 1)} \\ &= \frac{\widehat{u}_{i+1}(x_1, \dots, x_{i-1}, 0)}{u_{i+1}(x_1, \dots, x_{i-1}, 0)} + o(1), \end{aligned} \quad (10)$$

or $x_i = 0$, $a_{i+1}(x_1, \dots, x_i) = 1$, and

$$(1 - q_i)\widehat{u}_{i+1}(x_1, \dots, x_{i-1}, 1) = o(q_i\widehat{u}_{i+1}(x_1, \dots, x_{i-1}, 0)) \quad (11)$$

or $x_i = 1$, $a_{i+1}(x_1, \dots, x_i) = 1$, and

$$q_i\widehat{u}_{i+1}(x_1, \dots, x_{i-1}, 0) = o((1 - q_i)\widehat{u}_{i+1}(x_1, \dots, x_{i-1}, 1)). \quad (12)$$

Then the estimator from the IS scheme defined via (8) has the VRE property. In particular, if the assumptions of Theorem 2 hold with $a_{i+1}(x_1, \dots, x_i) = 1$ for all $x \in \mathcal{S}_1$ and all i , then Condition (10) always hold.

Proof: We decompose the second moment in two terms as follows:

$$\begin{aligned} \widetilde{\mathbb{E}}[\phi(X)L^2(X)] &= \sum_{x \in \mathcal{S}_1} \phi(x)L^2(x)\widetilde{\mathbb{P}}[X = x] \\ & \quad + \sum_{x \in \mathcal{S}_0} \phi(x)L^2(x)\widetilde{\mathbb{P}}[X = x]. \end{aligned} \quad (13)$$

For the second term, following the same argument as in the proof of Theorem 2, we find that there exists a constant $d > 0$, independent of ϵ , such that $L^2(x) \leq d^{2m}u^2 + o(u^2)$, and, because \mathcal{S}_0 is finite,

$$\begin{aligned} & \sum_{x \in \mathcal{S}_0} \phi(x)L^2(x)\widetilde{\mathbb{P}}[X = x] \\ &\leq (d^{2m}u^2 + o(u^2)) \sum_{x \in \mathcal{S}_0} \phi(x)\widetilde{\mathbb{P}}[X = x]. \end{aligned}$$

But this last sum is $o(1)$, because it is a *finite* sum of $o(1)$ terms. Consequently,

$$\sum_{x \in \mathcal{S}_0} \phi(x)L^2(x)\widetilde{\mathbb{P}}[X = x] = o(u^2).$$

We now focus on the first term on the right side of (13). For the special case where the conditions of Theorem 2 hold with $a_{i+1}(x_1, \dots, x_i) = 1$ for all $x \in \mathcal{S}_1$ and all i , we can take $b_i(x_1, \dots, x_{i-1}) = 1$ and $d = 1$ in the proof of that theorem and it gives

$$\sum_{x \in \mathcal{S}_1} \phi(x)L^2(x)\widetilde{\mathbb{P}}[X = x] = u^2 + o(u^2). \quad (14)$$

For the more general case, when (10) holds, we have $b_i(x_1, \dots, x_{i-1}) = a_{i+1}(x_1, \dots, 1) = a_{i+1}(x_1, \dots, 0)$. When either (11) or (12) holds, we find from (9) and (5) that $b_i(x_1, \dots, x_{i-1}) = a_{i+1}(x_1, \dots, x_i) = 1$. So in all three cases, we can take $d = 1$, and we also obtain (14).

Combining the results for the two terms, we have

$$\widetilde{\text{Var}}[\phi(X)L(X)] \leq \widetilde{\mathbb{E}}[L^2(X)] - u^2 = o(u^2),$$

and this concludes the proof. \blacksquare

The next question is how to construct approximations $\widehat{u}_{i+1}(\cdot)$ for which these theorems apply. We make one proposal in the next section.

V. AN APPROXIMATION BASED ON MINIMAL CUTS

We introduce a method to approximate $\widehat{u}_{i+1}(\cdot)$ based on the probabilities of certain minimal cuts in subgraphs obtained at each step. We call this method the *mincut-maxprob* approximation. This approximation has the advantage of being relatively simple and easy to compute, and it always satisfies the condition of Theorem 2. It is not the only possibility; other approximations based on other types of cuts or on combinations of cuts might also work well.

We need some definitions and notation. A *cut* (or \mathcal{K} -cut) in the graph \mathcal{G} is a set of edges such that if we remove them from \mathcal{G} , the nodes in \mathcal{K} are not in the same connected component of the resulting graph. A *mincut* (minimal cut) is a cut that contains no other cut than itself. With any cut γ we associate the event $E(\gamma)$: “all links in γ are failed.” Observe that the system unreliability u can be expressed as

$$u = \mathbb{P}[\cup_{\text{all mincuts } \gamma} E(\gamma)], \quad (15)$$

because the system is non-operational if and only if at least one of its mincuts has all its links failed. If the edge unreliabilities have the form (1), then for any cut γ , $\mathbb{P}(E(\gamma)) = \prod_{i \in \gamma} q_i$ has also the same form. The inclusion-exclusion formula applied to u immediately leads then to (2).

A *mincut with maximal probability* is a mincut γ such that

$$\gamma = \arg \max_{\gamma'} \{ \mathbb{P}[E(\gamma')] : \gamma' \text{ is a mincut of } \mathcal{G} \}.$$

If we denote $p = \mathbb{P}(E(\gamma))$ and if H is the number of mincuts with maximal probability p , we have, using (15) and (2), $p = \Theta(\epsilon^c)$ and $u = Hp + o(\epsilon^c)$.

For each i , given x_1, \dots, x_i (assumed fixed), consider the graph $\mathcal{G}_i = \mathcal{G}_i(x_1, \dots, x_i)$ obtained from \mathcal{G} by removing all links $j \leq i$ for which $x_j = 0$ and forcing the links j such that $x_j = 1$ to be operational. Let \mathcal{C}_i be the set of mincuts in \mathcal{G}_i that contain no link $j \leq i$ for which $x_j = 1$. Let γ_i be a mincut of maximal probability in \mathcal{C}_i . That is,

$$\mathbb{P}[E(\gamma_i)] = \max \{ \mathbb{P}[E(\gamma)] : \gamma \in \mathcal{C}_i \}.$$

We propose to use $\mathbb{P}[E(\gamma_i)]$ as an approximation of $u_{i+1}(x_1, \dots, x_i)$. The idea is that when $E(\gamma_i)$ occurs, all

links in γ_i are failed so \mathcal{K} is necessarily disconnected, and γ_i is a cut of largest probability having this property. We disallow γ_i to contain links $j \leq i$ with $x_j = 1$ because we already know that these links are operational.

At step i of the Markov chain, we need $\widehat{u}_{i+1}(x_1, \dots, x_{i-1}, 0)$ and $\widehat{u}_{i+1}(x_1, \dots, x_{i-1}, 1)$ to compute \widetilde{q}_i via (8). For this, we find a mincut $\gamma_i^- = \gamma_i$ as above for x_1, \dots, x_{i-1} fixed and $x_i = 0$, and put $\widehat{u}_{i+1}(x_1, \dots, x_{i-1}, 0) = \mathbb{P}[E(\gamma_i^-)]$. We do the same with $x_i = 1$ to find a mincut $\gamma_i^+ = \gamma_i$ and put $\widehat{u}_{i+1}(x_1, \dots, x_{i-1}, 1) = \mathbb{P}[E(\gamma_i^+)]$. Note that these mincuts can be computed in polynomial time [27] (a typical way of doing this is by defining as the cost of link i the number $-\ln q_i$). Then we generate X_i using this probability \widetilde{q}_i and we compute the corresponding term of the likelihood ratio.

Of course, these computations add overhead to the algorithm compared with crude MC, but this computational burden can be more than compensated by the huge variance reduction, as our experiments will show. We now show that this proposed approximation always gives an estimator with BRE.

Theorem 4: The mincut-maxprob approximation described above always satisfies the condition of Theorem 2. Consequently it gives an IS scheme with the BRE property.

Proof: Recall that the unreliability of link i is $q_i = a_i \epsilon^{b_i}$. Fixing the state of links 1 to i to x_1, \dots, x_i , there exists a constant $c_{i+1}(x_1, \dots, x_i)$ independent of ϵ such that $u_{i+1}(x_1, \dots, x_i) = \Theta(\epsilon^{c_{i+1}(x_1, \dots, x_i)})$. This comes from

$$u_{i+1}(x_1, \dots, x_i) = \sum_{(x_{i+1}, \dots, x_m) \in \{0,1\}^{m-i}} \phi(x) \mathbb{P}[X = x],$$

which is a finite sum of polynomials in ϵ , therefore a polynomial too, and $c_{i+1}(x_1, \dots, x_i)$ is the minimum value of the exponents of this polynomial. Note that we also have

$$c_{i+1}(x_1, \dots, x_i) = \min \{ d : \mathbb{P}[E(\gamma)] = \Theta(\epsilon^d) \text{ for some mincut } \gamma \text{ of } \mathcal{G}_{i+1}(x_1, \dots, x_i) \}.$$

As a consequence, the mincut-maxprob approximation gives

$$\begin{aligned} \widehat{u}_{i+1}(x_1, \dots, x_i) &= \Theta(\epsilon^{c_{i+1}(x_1, \dots, x_i)}) \\ &= \Theta(u_{i+1}(x_1, \dots, x_i)). \end{aligned}$$

The assumptions of Theorem 2 are thus verified. \blacksquare

Theorem 5: Under an IS scheme based on the mincut-maxprob approximation, suppose that for each $x = (x_1, \dots, x_m) \in \mathcal{S}_1$ and for each $i \in \{1, \dots, m-1\}$, the graph $\mathcal{G}_i(x_1, \dots, x_i)$ contains only one mincut with probability $\Theta(u_{i+1}(x_1, \dots, x_i))$ and one of the following three conditions is satisfied:

(i) the graph $\mathcal{G}_i(x_1, \dots, x_{i-1}, 1 - x_i)$ contains only one mincut having probability $\Theta(u_{i+1}(x_1, \dots, x_{i-1}, 1 - x_i))$, or

(ii) $x_i = 0$ and

$$(1 - q_i)\widehat{u}_{i+1}(x_1, \dots, x_{i-1}, 1) = o(q_i\widehat{u}_{i+1}(x_1, \dots, x_{i-1}, 0)), \quad (16)$$

or

(iii) $x_i = 1$ and

$$q_i\widehat{u}_{i+1}(x_1, \dots, x_{i-1}, 0) = o((1 - q_i)\widehat{u}_{i+1}(x_1, \dots, x_{i-1}, 1)). \quad (17)$$

Then the assumptions of Theorem 3 are satisfied, and consequently we have an IS scheme with the VRE property.

Proof: Because there is only one mincut with probability $\Theta(u_i(x_1, \dots, x_{i-1}))$ in \mathcal{G}_{i-1} , the (finite) sum of probabilities of the other cuts is $o(u_i(x_1, \dots, x_{i-1}))$. Thus for all $i \in \{1, \dots, m-1\}$,

$$\frac{\widehat{u}_{i+1}(x_1, \dots, x_{i-1}, x_i)}{u_{i+1}(x_1, \dots, x_{i-1}, x_i)} = 1 + o(1).$$

Similarly the condition that there is also one mincut with probability $\Theta(u_{i+1}(x_1, \dots, 1 - x_i))$ corresponds to

$$\frac{\widehat{u}_{i+1}(x_1, \dots, x_{i-1}, 1 - x_i)}{u_{i+1}(x_1, \dots, x_{i-1}, 1 - x_i)} = 1 + o(1),$$

implying that (10) holds, while conditions (16) and (17) correspond exactly to (11) and (12). Then the result follows from Theorem 3. \blacksquare

Our next and final result gives sufficient VRE conditions for the special case where \mathcal{G} has homogeneous link reliabilities.

Theorem 6: Suppose that the graph \mathcal{G} is homogeneous, which means (without loss of generality) that $q_i = \epsilon$ for all i . Suppose also that we have an IS scheme based on the mincut-maxprob approximation, under which for each $x \in \mathcal{S}_1$ and each $i \in \{1, \dots, m-1\}$, one of the following two conditions holds:

(i) the number of mincuts of probability $\Theta(u_{i+1}(x_1, \dots, x_{i-1}, 0))$ in $\mathcal{G}_{i+1}(x_1, \dots, x_{i-1}, 0)$ is the same as the number of mincuts of probability $\Theta(u_{i+1}(x_1, \dots, x_{i-1}, 1))$ in $\mathcal{G}_{i+1}(x_1, \dots, x_{i-1}, 1)$, or

(ii) the graph $\mathcal{G}_i(x_1, \dots, x_i)$ contains only one mincut with probability $\Theta(u_{i+1}(x_1, \dots, x_i))$ and either $x_i = 0$ and

$$(1 - q_i)\widehat{u}_{i+1}(x_1, \dots, x_{i-1}, 1) = o(q_i\widehat{u}_{i+1}(x_1, \dots, x_{i-1}, 0)), \quad (18)$$

or $x_i = 1$ and

$$q_i\widehat{u}_{i+1}(x_1, \dots, x_{i-1}, 0) = o((1 - q_i)\widehat{u}_{i+1}(x_1, \dots, x_{i-1}, 1)). \quad (19)$$

Then the assumptions of Theorem 3 are satisfied, and consequently we have an IS scheme with the VRE property.

Proof: Let $f_{i+1}(x_1, \dots, x_i)$ be the number of mincuts with probability $\Theta(u_{i+1}(x_1, \dots, x_i))$. Under the first condition we have $f_{i+1}(x_1, \dots, x_{i-1}, 0) = f_{i+1}(x_1, \dots, x_{i-1}, 1)$. Similarly to the proof of Theorem 4, $u_{i+1}(x_1, \dots, x_i) = f_{i+1}(x_1, \dots, x_i)e^{c_{i+1}(x_1, \dots, x_i)} + o(u_{i+1}(x_1, \dots, x_i))$ while $\widehat{u}_{i+1}(x_1, \dots, x_i) = e^{c_{i+1}(x_1, \dots, x_i)}$. Then,

$$\begin{aligned} \frac{\widehat{u}_{i+1}(x_1, \dots, x_{i-1}, 1)}{u_{i+1}(x_1, \dots, x_{i-1}, 1)} &= \frac{1}{f_{i+1}(x_1, \dots, x_{i-1}, 0)} + o(1), \\ \frac{\widehat{u}_{i+1}(x_1, \dots, x_{i-1}, 0)}{u_{i+1}(x_1, \dots, x_{i-1}, 0)} &= \frac{1}{f_{i+1}(x_1, \dots, x_{i-1}, 1)} + o(1), \end{aligned}$$

which corresponds to (10). In the second set of conditions, (18) and (19) correspond to (11) and (12), respectively. The conditions of Theorem 3 are therefore verified. \blacksquare

We now give small examples to illustrate the mincut-maxprob approximation. In all our examples, the links are sampled by order of their number.

Example 1: Consider the tiny graph of Figure 1. The aim is to compute the probability that the gray nodes A and C are disconnected. The links are homogeneous with unreliability $q_i = \epsilon$ for $i = 1, 2, 3$, for some small real number ϵ . For this model, only three configurations lead

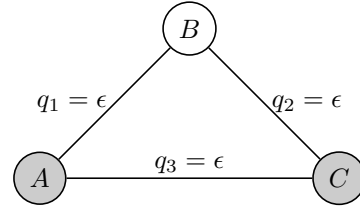


Fig. 1. A graph with three links

to disconnected nodes: $(0, 0, 0)$, $(0, 1, 0)$ and $(1, 0, 0)$ and the system unreliability is $u = \epsilon^3 + 2\epsilon^2(1 - \epsilon) = 2\epsilon^2 - \epsilon^3$. We can check that $u_2(0) = \epsilon$, $u_2(1) = \epsilon^2$, $u_3(0, 0) = \epsilon$, $u_3(0, 1) = \epsilon$, $u_3(1, 0) = \epsilon$ and $u_3(1, 1) = 0$. At step 1, we need to compute $\widehat{u}_2(0)$ and $\widehat{u}_2(1)$ to find \tilde{q}_1 and sample the first link. For $x_1 = 0$, the mincut γ_1^- contains only link 3 and its failure probability is $\widehat{u}_2(0) = \epsilon = u_2(0)$. For $x_1 = 1$, the mincut γ_1^+ contains links 2 and 3, and we have $\widehat{u}_2(1) = \epsilon^2 = u_2(1)$. This gives

$$\tilde{q}_1 = \frac{q_1\widehat{u}_2(0)}{q_1\widehat{u}_2(0) + (1 - q_1)\widehat{u}_2(1)} = \frac{\epsilon^2}{\epsilon^2 + (1 - \epsilon)\epsilon^2} = \frac{1}{2 - \epsilon}.$$

Continuing in this way, we get $\widehat{u}_3(0, 0) = \epsilon$, $\widehat{u}_3(0, 1) = \epsilon$, $\widehat{u}_3(1, 0) = \epsilon$, and $\widehat{u}_3(1, 1) = 0$, which are all equal to the corresponding values of $u_i(\cdot)$. The values of u_4 and \widehat{u}_4 are also both equal to those of ϕ . As a consequence, for this example, our approximate zero-variance estimator is exactly the zero-variance one.

Example 2: Consider the graph represented in Figure 2, where we want to compute the probability that the gray nodes A and D are disconnected. The links are again assumed homogeneous, with unreliabilities $q_i = \epsilon$ for $i = 1, \dots, 4$. There are 16 possible configurations of

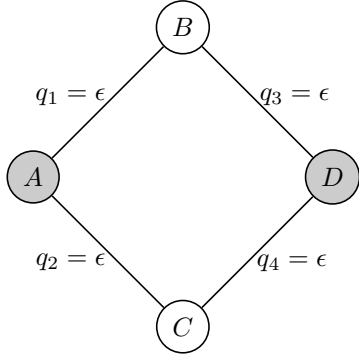


Fig. 2. A graph topology with four links and two nodes (A and D) that require to be connected

the 4 links, and one can verify that A and D are disconnected for 9 of them, and that the unreliability is $u = 4\epsilon^2 + o(\epsilon^2)$. The reader can verify in this case that $u_2(0) = 2\epsilon - \epsilon^2$, $u_2(1) = 2\epsilon^2 - \epsilon^3$, $u_3(0,0) = 1$, $u_3(0,1) = \epsilon$, $u_3(1,0) = \epsilon$, and $u_3(1,1) = \epsilon^2$, for instance, whereas the mincut-maxprob approximation gives $\hat{u}_2(0) = \epsilon$, $\hat{u}_2(1) = \epsilon^2$, $\hat{u}_3(0,0) = 1$, $\hat{u}_3(0,1) = \epsilon$, $\hat{u}_3(1,0) = \epsilon$, and $\hat{u}_3(1,1) = \epsilon^2$. The values of \hat{u}_4 are exactly the same as those of u_4 . Observe that $\hat{u}_2(0) \approx u_2(0)/2$ and $\hat{u}_2(1) \approx u_2(1)/2$, whereas $\hat{u}_i(\cdot) = u_i(\cdot)$ in the other cases. Thus, the conditions of Theorem 2 are satisfied, but not with $a_{i+1}(x_1, \dots, x_i) = 1$ for all i , because $a_2(0) = a_2(1) = 1/2$. However, the conditions of Theorem 3 are also satisfied, because $\hat{u}_2(0)/u_2(0) = \hat{u}_2(1)/u_2(1) + o(1)$, so we have VRE. In fact, one can also verify for this particular example that the IS scheme with our approximation gives exactly zero variance. The reason is that even when sampling the first link, the mincut-maxprob approximation yields

$$\begin{aligned} \tilde{q}_1 &= \frac{q_1 \hat{u}_2(0)}{q_1 \hat{u}_2(0) + (1 - q_1) \hat{u}_2(1)} \\ &= \frac{\epsilon^2}{\epsilon^2 + (1 - \epsilon)\epsilon^2} = \frac{1}{2 - \epsilon}, \end{aligned}$$

which is exactly the same as when using zero-variance IS:

$$\begin{aligned} \tilde{q}_1 &= \frac{q_1 u_2(0)}{q_1 u_2(0) + (1 - q_1) u_2(1)} \\ &= \frac{\epsilon(2\epsilon - \epsilon^2)}{\epsilon(2\epsilon - \epsilon^2) + (1 - \epsilon)(2\epsilon^2 - \epsilon^3)} = \frac{1}{2 - \epsilon}. \end{aligned}$$

The next two examples illustrate a situation where Theorems 5 applies and we have VRE, and another situation where only BRE holds.

Example 3: Consider the graph of Figure 3, where we again want to compute the probability that the gray nodes A and D are disconnected. Links are assumed homogeneous, with unreliability $q_i = \epsilon$ for $i = 1, \dots, 5$.

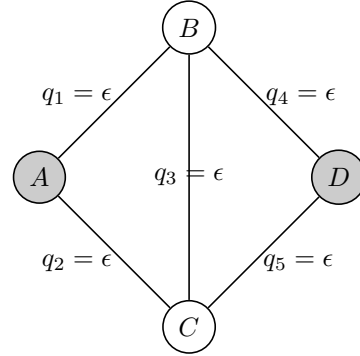


Fig. 3. Graph topology with five links and two nodes requiring to be connected

TABLE I
UNRELIABILITIES AND CORRESPONDING MINCUT-MAXPROB APPROXIMATIONS FOR EXAMPLE 3.

State (x_1, \dots, x_i)	$u_{i+1}(x_1, \dots, x_i)$	$\hat{u}_{i+1}(x_1, \dots, x_i)$
(0)	$\epsilon + o(\epsilon)$	ϵ
(1)	$\epsilon^2 + o(\epsilon^2)$	ϵ^2
(0,0)	1	1
(0,1)	$2\epsilon^2 - \epsilon^3$	ϵ^2
(1,0)	$2\epsilon^2 - \epsilon^3$	ϵ^2
(1,1)	ϵ^2	ϵ^2
(0,1,0)	ϵ	ϵ
(0,1,1)	ϵ^2	ϵ^2
(1,0,0)	ϵ	ϵ
(1,0,1)	ϵ^2	ϵ^2
(1,1,0)	ϵ^2	ϵ^2
(1,1,1)	ϵ^2	ϵ^2
(0,1,0,0)	ϵ	ϵ
(0,1,0,1)	ϵ	ϵ
(0,1,1,0)	ϵ	ϵ
(0,1,1,1)	0	0
(1,0,0,0)	1	1
(1,0,0,1)	0	0
(1,0,1,0)	ϵ	ϵ
(1,0,1,1)	0	0
(1,1,0,0)	ϵ	ϵ
(1,1,0,1)	0	0
(1,1,1,0)	ϵ	ϵ
(1,1,1,1)	0	0

Table I compares the unreliabilities and mincut-maxprob approximations at the different steps of the method. To simplify, we only list the states in which the final outcome (being connected or not) is not yet known. Table II, on the other hand, lists the 16 out of 32 configurations leading

TABLE II
PROBABILITIES OF THE 16 CONFIGURATIONS X THAT GIVE $\phi(X) = 1$ UNDER THE ORIGINAL AND MODIFIED PROBABILITIES, AND THE CORRESPONDING CONTRIBUTIONS TO THE SECOND MOMENT UNDER IS, FOR EXAMPLE 3.

Configuration	$\mathbb{P}[X = x]$	$\tilde{\mathbb{P}}[X = x]$	$L^2(x)\tilde{\mathbb{P}}[X = x]$
(0, 0, 0, 0, 0)	ϵ^5	$\frac{\epsilon^3}{(2-\epsilon)(1+\epsilon-\epsilon^2)}$	$\epsilon^7(2-\epsilon)(1+\epsilon-\epsilon^2)$
(0, 0, 0, 0, 1)	$\epsilon^4(1-\epsilon)$	$\frac{\epsilon^2(1-\epsilon)}{(2-\epsilon)(1+\epsilon-\epsilon^2)}$	$\epsilon^6(1-\epsilon)(2-\epsilon)(1+\epsilon-\epsilon^2)$
(0, 0, 0, 1, 0)	$\epsilon^4(1-\epsilon)$	$\frac{\epsilon^2(1-\epsilon)}{(2-\epsilon)(1+\epsilon-\epsilon^2)}$	$\epsilon^6(1-\epsilon)(2-\epsilon)(1+\epsilon-\epsilon^2)$
(0, 0, 0, 1, 1)	$\epsilon^3(1-\epsilon)^2$	$\frac{\epsilon(1-\epsilon)^2}{(2-\epsilon)(1+\epsilon-\epsilon^2)}$	$\epsilon^5(1-\epsilon)^2(2-\epsilon)(1+\epsilon-\epsilon^2)$
(0, 0, 1, 0, 0)	$\epsilon^4(1-\epsilon)$	$\frac{\epsilon^2(1-\epsilon)}{(2-\epsilon)(1+\epsilon-\epsilon^2)}$	$\epsilon^6(1-\epsilon)(2-\epsilon)(1+\epsilon-\epsilon^2)$
(0, 0, 1, 1, 0)	$\epsilon^3(1-\epsilon)^2$	$\frac{\epsilon(1-\epsilon)^2}{(2-\epsilon)(1+\epsilon-\epsilon^2)}$	$\epsilon^5(1-\epsilon)^2(2-\epsilon)(1+\epsilon-\epsilon^2)$
(0, 0, 1, 0, 1)	$\epsilon^3(1-\epsilon)^2$	$\frac{\epsilon(1-\epsilon)^2}{(2-\epsilon)(1+\epsilon-\epsilon^2)}$	$\epsilon^5(1-\epsilon)^2(2-\epsilon)(1+\epsilon-\epsilon^2)$
(0, 0, 1, 1, 1)	$\epsilon^2(1-\epsilon)^3$	$\frac{(1-\epsilon)^3}{(2-\epsilon)(1+\epsilon-\epsilon^2)}$	$\epsilon^4(1-\epsilon)^3(2-\epsilon)(1+\epsilon-\epsilon^2)$
(0, 1, 0, 0, 0)	$\epsilon^4(1-\epsilon)$	$\frac{\epsilon^2(1-\epsilon)}{(2-\epsilon)^2(1+\epsilon-\epsilon^2)}$	$\epsilon^6(1-\epsilon)(2-\epsilon)^2(1+\epsilon-\epsilon^2)$
(0, 1, 0, 1, 0)	$\epsilon^3(1-\epsilon)^2$	$\frac{\epsilon(1-\epsilon)^2}{(2-\epsilon)^2(1+\epsilon-\epsilon^2)}$	$\epsilon^5(1-\epsilon)^2(2-\epsilon)^2(1+\epsilon-\epsilon^2)$
(0, 1, 1, 0, 0)	$\epsilon^3(1-\epsilon)^2$	$\frac{\epsilon(1-\epsilon)^2}{(2-\epsilon)^2(1+\epsilon-\epsilon^2)}$	$\epsilon^5(1-\epsilon)^2(2-\epsilon)^2(1+\epsilon-\epsilon^2)$
(1, 0, 0, 0, 0)	$\epsilon^4(1-\epsilon)$	$\frac{\epsilon^2(1-\epsilon)}{(2-\epsilon)^2}$	$\epsilon^6(1-\epsilon)(2-\epsilon)^2$
(1, 0, 0, 0, 1)	$\epsilon^3(1-\epsilon)^2$	$\frac{\epsilon(1-\epsilon)^2}{(2-\epsilon)^2}$	$\epsilon^5(1-\epsilon)^2(2-\epsilon)^2$
(1, 0, 1, 0, 0)	$\epsilon^3(1-\epsilon)^2$	$\frac{\epsilon(1-\epsilon)^2}{(2-\epsilon)^2}$	$\epsilon^5(1-\epsilon)^2(2-\epsilon)^2$
(1, 1, 0, 0, 0)	$\epsilon^3(1-\epsilon)^2$	$\frac{\epsilon(1-\epsilon)^2}{(2-\epsilon)}$	$\epsilon^5(1-\epsilon)^2(2-\epsilon)$
(1, 1, 1, 0, 0)	$\epsilon^2(1-\epsilon)^3$	$\frac{(1-\epsilon)^3}{(2-\epsilon)}$	$\epsilon^4(1-\epsilon)^2(2-\epsilon)$

to disconnected nodes A and D , with their probabilities under the original probability law and under IS. Only two configurations, $(0, 0, 1, 1, 1)$ and $(1, 1, 1, 0, 0)$, have a probability $\Theta(1)$ under IS and it can be readily checked that the sufficient conditions of Theorem 6 are verified. To verify VRE directly, Table II also displays the contribution for each configuration to the variance of the overall estimator under IS. We find from this table and easy computations that $u = 2\epsilon^2 + 2\epsilon^3 - 5\epsilon^4 + 2\epsilon^5$ and

$$\sum_{x \in \{0,1\}^5} L^2(x)\phi(x)\tilde{\mathbb{P}}[X = x] = 4\epsilon^4 + o(\epsilon^4),$$

so the relative variance is

$$\frac{1}{u^2} \sum_{x \in \{0,1\}^5} \phi(x)L^2(x)\tilde{\mathbb{P}}[X = x] - 1 = o(1),$$

in accordance with Theorem 6.

Example 4: Consider a variation of Example 2 where \mathcal{K} is now the set of all nodes, as shown in Figure 4. The system works only if all nodes are connected. It is in a failed state if and only if at least two links are failed. This gives 11 failed configurations out of 16, and $u = 6\epsilon^2 + o(\epsilon^2)$. Table III gives the original and modified probabilities of those 11 configurations, as well as the contributions to the second moment under IS. Here, for instance, $u_3(0, 0) = \hat{u}_3(0, 0) = 1$, but $u_3(0, 1) = 2\epsilon + o(\epsilon)$

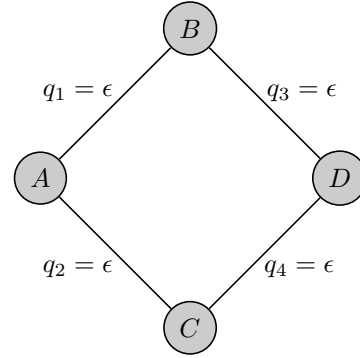


Fig. 4. A graph topology with four links and requiring full connectivity.

whereas $\hat{u}_3(0, 1) = \epsilon$. Thus if $x_1 = 0$, the unreliability of the second link under the zero-variance change of measure is $\epsilon u_3(0, 0) / (\epsilon u_3(0, 0) + (1-\epsilon)u_3(0, 1)) = 1/3 + o(1)$ and it is $1/2 + o(1)$ with the mincut-maxprob approximation. Thus we have BRE, thanks to Theorem 4, but the sufficient conditions of Theorems 5 or 6 are not verified, and this can be seen by looking for instance at configuration $(0, 0, 1, 1)$. One can directly see that VRE is not verified by directly computing the relative variance:

$$\frac{1}{u^2} \sum_{x \in \{0,1\}^m} \phi(x)L^2(x)\tilde{\mathbb{P}}[X = x] - 1 = \frac{1}{9} + o(1).$$

TABLE III
ORIGINAL AND MODIFIED PROBABILITIES OF THE 11 FAILED CONFIGURATIONS, AND ASSOCIATED CONTRIBUTION TO THE SECOND MOMENT UNDER IS, FOR EXAMPLE 4.

Configuration	$\mathbb{P}[X = x]$	$\tilde{\mathbb{P}}[X = x]$	$L^2(x)\tilde{\mathbb{P}}[X = x]$
(0, 0, 0, 0)	ϵ^4	$\frac{\epsilon^2}{(2-\epsilon)^2}$	$\epsilon^6(2-\epsilon)^2$
(0, 0, 0, 1)	$\epsilon^3(1-\epsilon)$	$\frac{\epsilon(1-\epsilon)}{(2-\epsilon)^2}$	$\epsilon^5(1-\epsilon)(2-\epsilon)^2$
(0, 0, 1, 0)	$\epsilon^3(1-\epsilon)$	$\frac{\epsilon(1-\epsilon)}{(2-\epsilon)^2}$	$\epsilon^5(1-\epsilon)(2-\epsilon)^2$
(0, 0, 1, 1)	$\epsilon^2(1-\epsilon)^2$	$\frac{(1-\epsilon)^2}{(2-\epsilon)^2}$	$\epsilon^4(1-\epsilon)^2(2-\epsilon)^2$
(0, 1, 0, 0)	$\epsilon^3(1-\epsilon)$	$\frac{\epsilon(1-\epsilon)}{(2-\epsilon)^3}$	$\epsilon^5(1-\epsilon)(2-\epsilon)^3$
(0, 1, 0, 1)	$\epsilon^2(1-\epsilon)^2$	$\frac{(1-\epsilon)^2}{(2-\epsilon)^3}$	$\epsilon^4(1-\epsilon)^2(2-\epsilon)^3$
(0, 1, 1, 0)	$\epsilon^2(1-\epsilon)^2$	$\frac{(1-\epsilon)^2}{(2-\epsilon)^3}$	$\epsilon^4(1-\epsilon)^2(2-\epsilon)^3$
(1, 0, 0, 0)	$\epsilon^3(1-\epsilon)$	$\frac{\epsilon(1-\epsilon)}{(2-\epsilon)^3}$	$\epsilon^5(1-\epsilon)(2-\epsilon)^3$
(1, 0, 0, 1)	$\epsilon^2(1-\epsilon)^2$	$\frac{(1-\epsilon)^2}{(2-\epsilon)^3}$	$\epsilon^4(1-\epsilon)^2(2-\epsilon)^3$
(1, 0, 1, 0)	$\epsilon^2(1-\epsilon)^2$	$\frac{(1-\epsilon)^2}{(2-\epsilon)^3}$	$\epsilon^4(1-\epsilon)^2(2-\epsilon)^3$
(1, 1, 0, 0)	$\epsilon^2(1-\epsilon)^2$	$\frac{(1-\epsilon)^2}{(2-\epsilon)^2}$	$\epsilon^4(1-\epsilon)^2(2-\epsilon)^2$

We emphasize that the BRE property is always verified, regardless of the order in which the links are sampled. For the VRE property, on the other hand, the sampling order may matter. Thus, the ordering may have an impact on the performance and this leads to the question: What is the best way to define that ordering? Intuitively, the general goal is to ensure that the approximation \tilde{q}_i , when sampling the i -th link, is as close as possible to the zero-variance probability. We do not yet have a precise algorithm (or a general heuristic rule) to find a good ordering *a priori* to achieve this goal.

VI. NUMERICAL ILLUSTRATIONS

We now report some simulation experiments with the mincut-maxprob approximation, for the toy examples of Section V and for larger ones. We also tried the method numerically for Examples 1 and 2, and the estimator was equal to u on each run, as expected.

Example 5: Table IV gives empirical results for Example 3, for sample size $n = 10^4$ and four different values of ϵ . It shows the estimate of u , a 95% confidence interval on u , the empirical standard deviation (per replication), and the empirical relative error. We observe a decrease of the relative error when ϵ decreases, in accordance with the VRE property.

Example 6: Table V reports similar results for Example 4. We saw in that example that the square relative error is $1/9 + o(1)$, so the relative error should converge to $1/3$ when $\epsilon \rightarrow 0$. This is exactly what we observe empirically.

Example 7: We now take a larger graph, made of 20 nodes and 30 links, with the dodecahedron topology as shown in Figure 5. This structure is often used as a benchmark for network reliability evaluation techniques

[6], [7], [13]. We consider the homogeneous case, where all links have the same unreliability ϵ , and we want to compute the probability that nodes A and B are disconnected. Links are ordered somewhat arbitrarily, according to their numbering in the figure. The empirical results, in Table VI, suggest that VRE holds. Our proposed IS scheme provides an extremely tight confidence interval on u when ϵ is very small. We know of no other method that provides VRE for this dodecahedron topology.

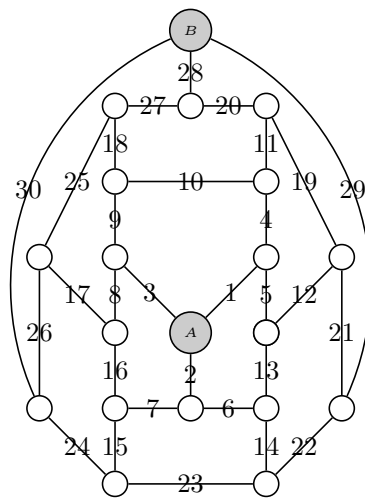


Fig. 5. Dodecahedron topology

Example 8: To experiment with an even larger network, of a size for which simulation is typically required, we construct a graph by juxtaposing three copies of the dodecahedron topology of Figure 5, forming a parallel system as shown in Figure 6. The resulting network has

TABLE IV
EMPIRICAL RESULTS FOR EXAMPLE 5, FOR $n = 10^4$ AND FOUR VALUES OF ϵ .

ϵ	Estimate	95% Confidence interval	Std deviation	Relative error
10^{-1}	2.150×10^{-2}	$(2.139 \times 10^{-2}, 2.160 \times 10^{-2})$	5.15×10^{-3}	0.24
10^{-2}	2.022×10^{-4}	$(2.018 \times 10^{-4}, 2.026 \times 10^{-4})$	2.07×10^{-5}	0.10
10^{-3}	2.0014×10^{-6}	$(2.0003 \times 10^{-6}, 2.0024 \times 10^{-6})$	5.28×10^{-8}	2.6×10^{-2}
10^{-4}	2.0002×10^{-8}	$(1.9998 \times 10^{-8}, 2.0006 \times 10^{-8})$	2.00×10^{-10}	1.0×10^{-2}

TABLE V
EMPIRICAL RESULTS FOR EXAMPLE 6, FOR $n = 10^4$ AND FOUR VALUES OF ϵ .

ϵ	Estimate	95% Confidence interval	Std deviation	Relative error
10^{-1}	5.271×10^{-2}	$(5.239 \times 10^{-2}, 5.302 \times 10^{-2})$	1.62×10^{-2}	0.31
10^{-2}	5.944×10^{-4}	$(5.905 \times 10^{-4}, 5.982 \times 10^{-4})$	1.96×10^{-4}	0.33
10^{-3}	6.015×10^{-6}	$(5.976 \times 10^{-6}, 6.054 \times 10^{-6})$	2.00×10^{-6}	0.33
10^{-4}	6.022×10^{-8}	$(5.982 \times 10^{-8}, 6.061 \times 10^{-8})$	2.00×10^{-8}	0.33

TABLE VI
EMPIRICAL RESULTS FOR EXAMPLE 7, FOR $n = 10^4$ AND FOUR VALUES OF ϵ .

ϵ	Estimate	95% Confidence interval	Std deviation	Relative error
10^{-1}	2.8960×10^{-3}	$(2.8276 \times 10^{-3}, 2.9645 \times 10^{-3})$	3.49×10^{-3}	1.2
10^{-2}	2.0678×10^{-6}	$(2.0611 \times 10^{-6}, 2.0744 \times 10^{-6})$	3.42×10^{-7}	0.17
10^{-3}	2.0076×10^{-9}	$(2.0053 \times 10^{-9}, 2.0099 \times 10^{-9})$	1.14×10^{-10}	0.057
10^{-4}	2.0007×10^{-12}	$(2.0000 \times 10^{-12}, 2.0014 \times 10^{-12})$	3.46×10^{-14}	0.017

90 links. The source and terminal nodes A and B of the three dodecahedron copies are merged, and the links are ordered in the same way as in Figure 5, starting with the first copy, then the second copy, and finally the third. We still consider the homogeneous case, where all links have the same unreliability ϵ , and we want to compute the probability that nodes A and B are disconnected. Note that the unreliability here is the cube of that of a single dodecahedron. The empirical results, in Table VII, suggest that VRE holds here too.

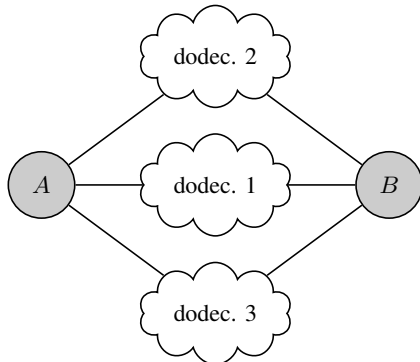


Fig. 6. Three dodecahedron topologies in parallel

Example 9: We consider a similar construction as in the previous example, except that the three dodecahedron copies are connected in series instead of in parallel, as shown in Figure 7. The source node is the source of the

first copy and the destination is the destination of the third one, while the destination of the first (respectively second) copy is the source of the second (respectively third). We still consider the homogeneous case, where all links have the same unreliability ϵ . Here the reliability (and not the unreliability) is the cube of that of a single dodecahedron topology. The empirical results in Table VIII suggest that only BRE holds.

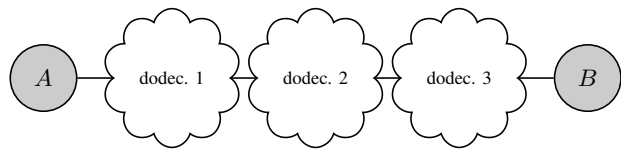


Fig. 7. Three dodecahedron topologies in series

The gain in accuracy (variance reduction) with the proposed method comes at the expense of increased computing times, because we need to find two mincuts with maximal probability at each step of the sampling process. This is done in polynomial time, using for instance the basic Ford-Fulkerson algorithm [27]. To give some idea of the added computational burden compared with crude Monte Carlo, we report here the average CPU time to simulate one realization of the graph with both methods, with our Java implementation on a MacBook with a 2.4 GHz Intel Core 2 Duo processor. For Example 3, the average simulation time is 9.4×10^{-6} seconds with crude Monte Carlo and 4.2×10^{-5} seconds with our mincut-

TABLE VII
EMPIRICAL RESULTS FOR EXAMPLE 8, FOR $n = 10^4$ AND THREE VALUES OF ϵ .

ϵ	Estimate	95% Confidence interval	Std deviation	Relative error
10^{-1}	2.3573×10^{-8}	$(2.2496 \times 10^{-8}, 2.4650 \times 10^{-8})$	5.49×10^{-8}	2.3
5×10^{-2}	2.5732×10^{-11}	$(2.5138 \times 10^{-11}, 2.6327 \times 10^{-11})$	3.03×10^{-11}	01.2
10^{-2}	8.7655×10^{-18}	$(8.7145 \times 10^{-18}, 8.8165 \times 10^{-18})$	2.60×10^{-18}	0.30

TABLE VIII
EMPIRICAL RESULTS FOR EXAMPLE 9, FOR $n = 10^4$ AND THREE VALUES OF ϵ .

ϵ	Estimate	95% Confidence interval	Std deviation	Relative error
10^{-2}	6.1919×10^{-6}	$(6.0190 \times 10^{-6}, 6.3648 \times 10^{-6})$	8.82×10^{-6}	1.4
10^{-3}	5.9730×10^{-9}	$(5.8213 \times 10^{-9}, 6.1146 \times 10^{-9})$	7.74×10^{-9}	1.3
10^{-4}	5.9246×10^{-12}	$(5.7766 \times 10^{-12}, 6.0727 \times 10^{-12})$	7.55×10^{-12}	1.3

maxprob IS algorithm. For Example 7 (the dodecahedron topology), the numbers are 4.8×10^{-5} seconds for crude Monte Carlo and 7.9×10^{-4} seconds for our method, so the time increases by a factor 16.5. For Example 8, this factor is about 52 while it is 26 for Example 9. These factors are not negligible, but they are modest in comparison with the (arbitrarily large) factors by which the variance can be reduced.

VII. CONCLUSIONS

We have described in this paper a new importance sampling scheme for the simulation of static reliability models. Our method tries to mimic the zero-variance change of measure, sampling sequentially the links and approaching the unreliability of the model given the status of previously sampled links by a rough estimation, here taken as the probability of a mincut. We proved that the BRE property is verified by our scheme when the unreliabilities of all links go to zero. This implies that compared with MC, the variance can be reduced by an arbitrarily large factor by taking ϵ small enough. Under additional conditions on the quality of the approximation, even VRE is verified. The efficiency of the method is illustrated on small examples to check the correctness of the implementation, and on a larger model often used in the area for illustration purposes. The VRE property is observed on benchmark examples, something not observed by the other rare-event simulation methods of the literature, emphasizing the interest of our proposal when links are very reliable.

This work will be pursued in several directions. We will first try to improve the approximation of the unreliability function, likely driving to a change of measure closer to the zero-variance one. The idea is to consider not only the mincut probability function, which is a lower bound of the unreliability, but other functions such as for instance an upper-bound taken as one minus the probability of a minpath, and to use a linear combination of those *basis*

functions to get a better approximation. This is similar to what is done in dynamic programming to approximate the Bellman value function. The coefficients of the linear combination will have to be learned. Another issue we would like to tackle is the combination of this importance sampling scheme with other variance reduction techniques such as those in [5], [7], where importance sampling is not considered. Adding an adapted zero-variance IS approximation to the techniques proposed there is likely to drastically reduce the variance and produce an extremely efficient estimator. A third direction for improvement is to combine our general procedure with graph reductions techniques that can be applied in polynomial time as a function of the model size, usually employed in the area, such as series-parallel reductions for example. These reductions should contribute to further diminish both the variance and the global cost in time of the method.

ACKNOWLEDGMENT

This work has been supported by an NSERC-Canada Discovery Grant and a Canada Research Chair to the first author, SIMERTEL région Bretagne project to the third author, and INRIA's associated team MOCQUASIN to all authors.

REFERENCES

- [1] S. Asmussen and P. W. Glynn, *Stochastic Simulation*. New York: Springer-Verlag, 2007.
- [2] M. O. Ball, "Computational complexity of network reliability analysis: An overview," *IEEE Transactions on Reliability*, vol. 35, no. 3, pp. 230–239, Aug. 1986.
- [3] R. Barlow and F. Proschan, *Statistical Theory of Reliability and Life Testing*. New York: Holt, Rinehart and Wilson, 1975.
- [4] J. Blanchet and M. Mandjes, "Rare event simulation for queues," in *Rare Event Simulation using Monte Carlo Methods*, G. Rubino and B. Tuffin, Eds. John Wiley & Sons, 2009, pp. 87–124, chapter 5.
- [5] H. Cancela and M. El Khadiri, "On the RVR simulation algorithm for network reliability evaluation," *IEEE Transactions on Reliability*, vol. 57, no. 12, pp. 207–212, 2003.

- [6] H. Cancela, M. El Khadiri, and G. Rubino, "Rare event analysis by monte carlo techniques in static models," in *Rare Event Simulation Using Monte Carlo Methods*, G. Rubino and B. Tuffin, Eds. Wiley, 2009, pp. 145–170, chapter 7.
- [7] H. Cancela, P. L'Ecuyer, M. Lee, G. Rubino, and B. Tuffin, "Analysis and improvements of path-based methods for Monte Carlo reliability evaluation of static models," in *Simulation Methods for Reliability and Availability of Complex Systems*, S. M. J. Faulin, A. A. Juan and E. Ramirez-Marquez, Eds. Springer Verlag, 2009.
- [8] H. Cancela, G. Rubino, and B. Tuffin, "New measures of robustness in rare event simulation," in *Proceedings of the 2005 Winter Simulation Conference*, M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines, Eds. IEEE Press, 2005, pp. 519–527.
- [9] C. J. Colbourn, *The Combinatorics of Network Reliability*. New York: Oxford University Press, 1987.
- [10] I. B. Gertsbakh and Y. Shpungin, *Models of Network Reliability*. Boca Raton, FL: CRC Press, 2009.
- [11] P. Glynn, G. Rubino, and B. Tuffin, "Robustness properties and confidence interval reliability issues," in *Rare Event Simulation using Monte Carlo Methods*, G. Rubino and B. Tuffin, Eds. John Wiley & Sons, 2009, pp. 63–84, chapter 4.
- [12] P. Heidelberger, "Fast simulation of rare events in queueing and reliability models," *ACM Transactions on Modeling and Computer Simulation*, vol. 5, no. 1, pp. 43–85, 1995.
- [13] K. P. Hui, N. Bean, M. Kraetzl, and D. P. Kroese, "The tree cut and merge algorithm for estimation of network reliability," *Prob. in Eng. and Inf. Sci.*, vol. 17, no. 1, pp. 24–45, 2003.
- [14] C. H. Jun and S. Ross, "System reliability by simulation: random hazards versus importance sampling," *Prob. in Eng. and Inf. Sci.*, vol. 6, pp. 119–311, 1992.
- [15] S. Juneja and P. Shahabuddin, "Rare event simulation techniques: An introduction and recent advances," in *Simulation*, ser. Handbooks in Operations Research and Management Science, S. G. Henderson and B. L. Nelson, Eds. Amsterdam, The Netherlands: Elsevier, 2006, pp. 291–350, chapter 11.
- [16] C. Kollman, K. Baggerly, D. Cox, and R. Picard, "Adaptive importance sampling on discrete Markov chains," *Annals of Applied Probability*, vol. 9, no. 2, pp. 391–412, 1999.
- [17] P. L'Ecuyer, J. H. Blanchet, B. Tuffin, and P. W. Glynn, "Asymptotic robustness of estimators in rare-event simulation," *ACM Transactions on Modeling and Computer Simulation*, vol. 20, no. 1, p. Article 6, 2010.
- [18] P. L'Ecuyer, M. Mandjes, and B. Tuffin, "Importance sampling and rare event simulation," in *Rare Event Simulation Using Monte Carlo Methods*, G. Rubino and B. Tuffin, Eds. Wiley, 2009, pp. 17–38, chapter 2.
- [19] P. L'Ecuyer and B. Tuffin, "Approximate zero-variance simulation," in *Proceedings of the 2008 Winter Simulation Conference*. IEEE Press, 2008, pp. 170–181.
- [20] —, "Approximating zero-variance importance sampling in a reliability setting," *Annals of Operations Research*, 2009, to appear.
- [21] M. K. Nakayama, "General Conditions for Bounded Relative Error in Simulations of Highly Reliable Markovian Systems," *Advances in Applied Probability*, vol. 28, pp. 687–727, 1996.
- [22] —, "General conditions for bounded relative error in simulations of highly reliable Markovian systems," *Advances in Applied Probability*, vol. 28, pp. 687–727, 1996.
- [23] S. Ross, "A new simulation estimator of system reliability," *J. of Applied Mathematics and Stochastic Analysis*, vol. 3, pp. 331–336, 1994.
- [24] G. Rubino, "Network reliability evaluation," in *State-of-the art in performance modeling and simulation*, K. Bagchi and J. Walrand, Eds. Gordon and Breach Books, 1998.
- [25] G. Rubino and B. Tuffin, "Markovian models for dependability analysis," in *Rare Event Simulation using Monte Carlo Methods*, G. Rubino and B. Tuffin, Eds. John Wiley & Sons, 2009, pp. 125–144, chapter 6.
- [26] G. Rubino and B. Tuffin, Eds., *Rare Event Simulation using Monte Carlo Methods*. John Wiley & Sons, 2009.
- [27] R. Sedgewick and M. Schidlowsky, *Algorithms in Java, Part 5: Graph Algorithms*. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc., 2003.
- [28] P. Shahabuddin, "Importance Sampling for the Simulation of Highly Reliable Markovian Systems," *Management Science*, vol. 40, no. 3, pp. 333–352, March 1994.
- [29] —, "Importance sampling for the simulation of highly reliable Markovian systems," *Management Science*, vol. 40, no. 3, pp. 333–352, 1994.

Pierre L'Ecuyer is a 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. He is a member of the CIRRELT and GERAD research centers. 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, Statistics and Computing*, *Management Science*, *International Transactions in Operational Research*, *The Open Applied Mathematics Journal*, and *Cryptography and Communications*. He obtained the *E. W. R. Steacie fellowship* in 1995–97, a *Killam fellowship* in 2001–03, and became an *INFORMS Fellow* in 2006. His recent research articles are available on-line from his web page: <http://www.iro.umontreal.ca/~lecuyer>.

Gerardo Rubino is a senior researcher at INRIA (the French National Institute for Research in Computer Science and Control) where he is the leader of the DIONYSOS (Dependability, Interoperability and performance analysis of networkS) team. His research interests are in the quantitative analysis of computer and communication systems, mainly using probabilistic models. He also works on the quantitative evaluation of perceptual quality of multimedia communications over the Internet. He recently co-edited the book *Rare event simulation using Monte Carlo methods* published by John Wiley & Sons in 2009. He is a member of the IFIP WG 7.3.

Samira Saggadi is a PhD student at INRIA, Rennes, France. Her research interests are the design of new rare event simulation techniques and their application to dependability and/or telecommunication problems.

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 eight 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 and telecommunication-related economical models. He is currently Associate Editor for *INFORMS Journal on Computing*, *ACM Transactions on Modeling and Computer Simulation* and *Mathematical Methods of Operations Research*. He has written or co-written two books devoted to simulation: *Rare event simulation using Monte Carlo methods* published by John Wiley & Sons in 2009, and *La simulation de Monte Carlo* (in French), published by Hermes Editions in 2010. His web page is http://www.irisa.fr/dionysos/pages_perso/tuffin/Tuffin_en.htm.