

COMBINATION OF CONDITIONAL MONTE CARLO AND APPROXIMATE ZERO-VARIANCE
IMPORTANCE SAMPLING FOR NETWORK RELIABILITY ESTIMATION

Hector Cancela

Universidad de la República
Montevideo, URUGUAY

Gerardo Rubino

INRIA Rennes Bretagne Atlantique
Campus Universitaire de Beaulieu
35042 Rennes Cedex, FRANCE

Pierre L'Ecuyer

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

Bruno Tuffin

INRIA Rennes Bretagne Atlantique
Campus Universitaire de Beaulieu
35042 Rennes Cedex, FRANCE

ABSTRACT

We study the combination of two efficient rare event Monte Carlo simulation techniques for the estimation of the connectivity probability of a given set of nodes in a graph when links can fail: approximate zero-variance importance sampling and a conditional Monte Carlo method which conditions on the event that a prespecified set of disjoint minpaths linking the set of nodes fails. Those two methods have been applied separately. Here we show how their combination can be defined and implemented, we derive asymptotic robustness properties of the resulting estimator when reliabilities of individual links go arbitrarily close to one, and we illustrate numerically the efficiency gain that can be obtained.

1 INTRODUCTION

In the design of telecommunication networks, an historically important topic from the reliability point of view has been the study of connectivity properties of the network topology. This study can be accomplished by building a probabilistic model of the network focusing on the possible failure of its components, from which we can compute the probability of still supporting the specified communications when taking these failures into account. Today, these problems are of high importance with the technological arrival of different types of wireless architectures. In these contexts, network components (and particularly, links) are more prone to fail because of changes in the environment, changes which are difficult to control. *Network reliability* is the branch of Operations Research where these problems are studied. In this area, exact computations are usually hard (i.e., highly computationally expensive), and Monte Carlo techniques provide the most powerful tools for quantitative evaluations of the systems under consideration.

In this paper, we consider the most representative model in network reliability. One can think of the system as a communication network, but there are several other applications of that same model. The model is represented by a non-oriented graph $\mathcal{G} = (\mathcal{N}, \mathcal{L})$ where \mathcal{N} is the set of m nodes, and $\mathcal{L} = \{1, \dots, \ell\}$ is the set of links connecting nodes. Nodes are assumed perfect, in the sense that they never fail. On the other hand, links may fail, with probability q_i for link i ($1 \leq i \leq \ell$), and failures are assumed to occur *independently* across links. Our goal is to compute the probability $q(\mathcal{G})$ that a given subset \mathcal{K} of the set of nodes is *not* connected, that is, that these nodes are not in the same connected component of \mathcal{G} after removing the failed links. In this context, the graph together with the probabilistic model is often called a *network*, and $q(\mathcal{G})$ is the *network unreliability*. The most frequent case is when estimating the *two-terminal* or *source-to-terminal unreliability*, where \mathcal{K} is comprised of only two nodes.

Formally, define $X_i = 1$ if link i works, and $X_i = 0$ otherwise. The random *configuration* of the graph is the vector $X = (X_1, \dots, X_\ell)$. Let us denote by $\phi(\cdot)$ the following structure function of the network: $\phi(\cdot)$ is defined on the set of all configurations $\{0, 1\}^\ell$ and takes values in $\{0, 1\}$; for each configuration $x = (x_1, \dots, x_\ell)$, we have $\phi(x) = 1$ if the set \mathcal{K} is *not* included in the same connected component in the graph that contains only the links i

for which $x_i = 1$, and the value 0 otherwise. Observe that, for convenience, our definition of $\phi(\cdot)$ differs from the structure function that is commonly adopted, and which corresponds to $1 - \phi(\cdot)$.

The network unreliability is then

$$q(\mathcal{G}) = \mathbb{E}[\phi(X)] = \sum_{x \in \{0,1\}^\ell} \phi(x) \mathbb{P}[X = x] = \sum_{x \in \{0,1\}^\ell} \phi(x) \prod_{i=1}^{\ell} (q_i(1 - x_i) + (1 - q_i)x_i).$$

The number of terms to be considered for computing this probability has 2^ℓ elements, that is, its size increases exponentially with the number of links. Actually, the \mathcal{K} -connectivity problem is known to be NP-hard in general (Ball 1986). As a consequence Monte Carlo simulation becomes a relevant tool when looking at large graphs.

Standard Monte Carlo consists in generating n independent copies of X (that is of the random graph), say $X^{(1)}, \dots, X^{(n)}$, and to take as an (unbiased) approximation of $q(\mathcal{G})$ the average value $\sum_{j=1}^n \phi(X^{(j)})/n$. From the law of large numbers, this approximation converges almost surely to $q(\mathcal{G})$ as $n \rightarrow \infty$. One can compute a confidence interval with confidence $1 - \alpha$, centered at the approximation, of half-width $c_\alpha \sigma / \sqrt{n}$, where c_α is the $1 - \alpha/2$ quantile of the standard normal distribution (with mean 0 and variance 1) and σ is the standard deviation of $\phi(X)$. From the Central Limit Theorem, this confidence interval is approximately valid when $nq(\mathcal{G})$ is large enough (Asmussen and Glynn 2007).

The standard Monte Carlo is based on the direct use of $\phi(X)$, which is a Bernoulli random variable, and therefore $\sigma^2 = q(\mathcal{G})(1 - q(\mathcal{G}))$. Then one can note that the relative half-width of the confidence interval, given by this half-width divided by the expected value we are computing, i.e., $n^{-1/2} c_\alpha \sigma / q(\mathcal{G})$, is $c_\alpha \sqrt{1 - q(\mathcal{G})} / \sqrt{nq(\mathcal{G})} \rightarrow \infty$ as $q(\mathcal{G}) \rightarrow 0$. In other words, the rarer the event, the larger the sample size required to get a confidence interval with a fixed relative accuracy.

This highlights the need for designing more subtle simulation techniques dealing with rare events. For more about rare event simulation, the reader can look at Rubino and Tuffin (2009). The goal is to use another estimator Y of $q(\mathcal{G})$ with variance $\text{Var}[Y]$ for which a relative accuracy is not (or at least less) sensitive to rarity. In general, define the relative error as $\text{RE}[Y] = (\text{Var}[Y])^{1/2} / \mathbb{E}[Y]$. An estimator Y will be said to verify bounded relative error (BRE) if $\text{RE}[Y]$ remains bounded as $q(\mathcal{G}) \rightarrow 0$. In that case, the sample size needed to get a specified relative accuracy is bounded whatever the rarity of the event. An even better property is the so-called vanishing relative error (VRE), meaning that $\text{RE}[Y] \rightarrow 0$ when $\varepsilon \rightarrow 0$ (L'Ecuyer et al. 2010a).

A well-known variance reduction technique is importance sampling (IS) (Asmussen and Glynn 2007). It consists in changing the probability law \mathbb{P} (and expectation \mathbb{E}) driving the system to a new one $\tilde{\mathbb{P}}$ (and associated expectation $\tilde{\mathbb{E}}$), but simulating a new random variable $\phi(X)L(X)$ instead of $\phi(X)$ to keep an unbiased estimator, with $L(x) = d\tilde{\mathbb{P}}(x)/d\mathbb{P}(x)$ is the associated likelihood ratio. Indeed, $\mathbb{E}[\phi(X)] = \int \phi(x) d\mathbb{P}(x) = \int \phi(x)L(x) d\tilde{\mathbb{P}}(x) = \tilde{\mathbb{E}}[\phi(X)L(X)]$. By an appropriate change of probability law, the variance $\text{Var}[\phi(X)L(X)]$ can be significantly reduced with respect to $\text{Var}[\phi(X)]$.

The rest of the paper is organized as follows. Section 2 details the simulation method studied in this paper. It combines conditional Monte Carlo and importance sampling. Section 3 analyzes its robustness properties as the reliability of individual links goes to 1. Section 4 illustrates the gain that can be obtained on benchmark examples, compared with standard Monte Carlo, and also compared with conditional Monte Carlo or importance sampling alone.

2 DESCRIPTION OF THE METHOD

The method we are going to describe and study is based on the use of minpaths. A minpath is a minimal set of links whose correct operation ensures that the global system is operational, i.e., that all terminals in \mathcal{K} are connected. In the general \mathcal{K} -terminal network reliability case, a minpath will topologically correspond to a \mathcal{K} -tree (a tree included in the given graph, whose terminal set is \mathcal{K}). In the case of source-to-terminal network reliability, that is, when $|\mathcal{K}| = 2$, a minpath topologically corresponds to an elementary path between the terminals. Although the numerical illustrations will be given in the case of the source-to-terminal network reliability, the general description given below corresponds to the general case.

Given the graph \mathcal{G} and the terminal set \mathcal{K} , select a set of h disjoint minpaths $\mathcal{P} = \{P_1, \dots, P_h\}$ connecting each the set \mathcal{K} . For the selection of the set \mathcal{P} , different methods for finding trees in graphs can be used. For the source-to-terminal case, see Remark 1 below or Cancela et al. (2009). Define by p_j ($1 \leq j \leq h$) the probability

that all links of minpath P_j work. Due to the independence of failures, $p_j = \prod_{i \in P_j} (1 - q_i)$. Let also m_j be the number of links in P_j .

Remark now that obtaining a configuration where the set \mathcal{X} of nodes is not connected requires that at least one link on each minpath P_j is failed. Let C be the event “every minpath in \mathcal{P} has at least one link that does not work.” This event has probability $p = \prod_{j=1}^h (1 - p_j)$. Instead of generating independent copies of the configuration X of the graph under its original distribution, we will generate them conditional on C , and then multiply the estimator by $p = \mathbb{P}[C]$ to recover an unbiased estimator. Let Y denote the Bernoulli random variable $\phi(X)$ conditional on C . We have

$$\mathbb{E}[\phi(X)] = \mathbb{E}[\phi(X) | C] \mathbb{P}[C] = p \mathbb{E}[Y].$$

We can estimate $q(\mathcal{G})$ by the average of n independent copies of pY . This kind of conditional Monte Carlo procedure is known to always reduce the variance with respect to the standard estimator (Asmussen and Glynn 2007).

It remains to see how we can efficiently generate X conditional on C . Suppose that the links on each minpath are ordered (in any way), and let $i_{j,1}, \dots, i_{j,m_j}$ be the ordered list of links on minpath P_j . Define the random variable W_j as the index of the first failed link in this list, for each P_j . The idea is to first sample these W_j directly, and then sample the other links. Note that $\mathbb{P}[W_j = w]$ is simply the probability that links numbered 1 to $w - 1$ on P_j are working and link w is not, conditionally on the fact that at least one link in the minpath is failed. We then have, $\forall w \in \{1, \dots, m_j\}$ (Cancela et al. 2009),

$$\mathbb{P}[W_j = w] = \frac{(1 - q_{i_{j,1}})(1 - q_{i_{j,2}}) \cdots (1 - q_{i_{j,w-1}})q_{i_{j,w}}}{1 - (1 - q_{i_{j,1}})(1 - q_{i_{j,2}}) \cdots (1 - q_{i_{j,m_j}})}$$

All other links $i \notin \cup_{j=1}^h \{i_{j,1}, \dots, i_{j,m_j}\}$ (i.e., that belong to none of the paths P_j) can then be sampled from their *a priori* original Bernoulli distribution, with a probability q_i to be failed, since no additional information is available from the event C . This gives a (conditional) configuration from which one can check if the nodes in \mathcal{X} are connected or not (i.e., compute Y).

This describes the conditional Monte Carlo procedure, which we will now combine with IS. To do that, instead of sampling the links with the probabilities just described, we will use modified probabilities prescribed by an IS procedure that mimics the zero-variance change of measure. A general description of zero-variance IS estimators can be found in L'Ecuyer and Tuffin (2008), and a specific description for the static reliability problem (but not using the conditional expectation that minpaths are failed) in L'Ecuyer et al. (2010b).

Recall that minpaths in \mathcal{P} are indexed (ordered) from 1 to h . Each minpath P_j also has its links ordered from 1 to m_j . Finally, the rest of the links are also indexed in some order. Define $\mathcal{G}_{j:w_1, \dots, w_j}$ as the graph such for each minpath P_k ($1 \leq k \leq j$) the first failed link is the w_k -th, that is, all links 1 to $w_k - 1$ on every minpath P_k are considered as working while the w_k -th is failed, the status of all other links of \mathcal{G} being still considered random. The principle of our IS procedure is the following:

- First sample the first broken link on each minpath P_j knowing the status of the first broken link of minpaths P_1, \dots, P_{j-1} according to

$$\tilde{\mathbb{P}}[W_j = w] = \mathbb{P}[W_j = w] \frac{q(\mathcal{G}_{j:w_1, \dots, w_{j-1}, w})}{\sum_{k=1}^{m_j} \mathbb{P}[W_j = k] q(\mathcal{G}_{j:w_1, \dots, w_{j-1}, k})}. \tag{1}$$

- Second, the link states not sampled yet are determined *successively* with the following IS rule, according to the predefined order of links (but where link states already sampled on minpaths are considered fixed). Define $\mathcal{G}_{i,0,\text{hist}}$ (respectively $\mathcal{G}_{i,1,\text{hist}}$) as the graph for which all the states of each minpath P_j up to W_j have been sampled as described just above, as well as all links $k \notin \cup_{j=1}^h \{i_{j,1}, \dots, i_{j,m_j}\}$ such that $k < i$, and link i is selected as failed (respectively as working). Here “hist” stands for “history” (to simplify the notation). Link i is sampled as failed with probability

$$\tilde{q}_i = \frac{q_i q(\mathcal{G}_{i,0,\text{hist}})}{q_i q(\mathcal{G}_{i,0,\text{hist}}) + (1 - q_i) q(\mathcal{G}_{i,1,\text{hist}})}. \tag{2}$$

The corresponding likelihood ratio is given by

$$L(X) = \prod_{j=1}^h \frac{\mathbb{P}[W_j = w_j]}{\tilde{\mathbb{P}}[W_j = w_j]} \prod_{1 \leq i \leq \ell, i \notin \cup_{j=1}^h \{i_{j,1}, \dots, i_{j,w_j}\}} \frac{q_i(1 - X_i) + (1 - q_i)X_i}{\tilde{q}_i(1 - X_i) + (1 - \tilde{q}_i)X_i}. \quad (3)$$

Define

$$Y' = \phi(X)L(X).$$

Theorem 1. *The estimator Y' of $\mathbb{E}[Y]$ is unbiased and has variance $\text{Var}[Y'] = 0$.*

Proof. We are going to show that Y' always produces the value $\mathbb{E}[Y] = q(\mathcal{G})$. As a consequence, we will have $\tilde{\mathbb{E}}[Y'] = \mathbb{E}[Y]$ and $\text{Var}[Y'] = 0$. For all $j \in \{0, \dots, h\}$, define $Y'(w_1, \dots, w_j)$ as a random variable having the same distribution as Y conditional on $W_1 = w_1, \dots, W_j = w_j$. Observe that $Y'(w_1, \dots, w_h)$ is the estimator obtained by applying the IS scheme (2) to sample successively the states of the links of the graph $\mathcal{G}_{h;w_1, \dots, w_h}$. From the results of L'Ecuyer et al. (2010b), we always have $Y'(w_1, \dots, w_h) = q(\mathcal{G}_{h;w_1, \dots, w_h})$. From (1), for $j \leq h$,

$$\begin{aligned} Y'(w_1, \dots, w_{j-1}) &= Y'(w_1, \dots, w_j) \frac{\mathbb{P}[W_j = w_j]}{\tilde{\mathbb{P}}[W_j = w_j]} \\ &= Y'(w_1, \dots, w_j) \frac{1}{q(\mathcal{G}_{j;w_1, \dots, w_{j-1}, w_j})} \sum_{k=1}^{m_j} \mathbb{P}[W_j = k] q(\mathcal{G}_{j;w_1, \dots, w_{j-1}, k}). \end{aligned}$$

But by induction (starting from $j = k$ down to $j = 1$), if $Y'(w_1, \dots, w_j) = q(\mathcal{G}_{j;w_1, \dots, w_j})$, using also that, by conditioning over the value of W_j ,

$$\sum_{k=1}^{m_j} \mathbb{P}[W_j = k] q(\mathcal{G}_{j;w_1, \dots, w_{j-1}, k}) = q(\mathcal{G}_{j-1;w_1, \dots, w_{j-1}}), \quad (4)$$

we get $Y'(w_1, \dots, w_{j-1}) = q(\mathcal{G}_{j-1;w_1, \dots, w_{j-1}}) \forall j$, and therefore $Y' = q(\mathcal{G})$. \square

Unfortunately, this estimator cannot be implemented exactly, because it requires the knowledge of $q(\mathcal{G}_{j;w_1, \dots, w_{j-1}, w})$, $q(\mathcal{G}_{i,0,\text{hist}})$ and $q(\mathcal{G}_{i,1,\text{hist}})$. If those values were known, there would be no need to use simulation, because $q(\mathcal{G})$ could be computed immediately. Instead, we propose here to use an approximation $\hat{q}(\cdot)$ of $q(\cdot)$, and to plug this approximation in expressions (1) and (2) in place of $q(\cdot)$. Our approximation follows the same principle as in L'Ecuyer et al. (2010b): we define it as the maximal probability of a mincut of the graph, and refer to it as a *mincut-maxprob* approximation. Recall that a *cut* (or \mathcal{H} -cut) in the graph \mathcal{G} is a set of links such that if we remove them from \mathcal{G} , the nodes in \mathcal{H} are not in the same connected component of the resulting graph, and that a *mincut* (minimal cut) is a cut that contains no other cut than itself. The probability of a cut is defined as the probability that all links of the cut fail. This approximation of $q(\cdot)$ has the advantage of being relatively easy to compute, in polynomial time, using for instance the basic Ford-Fulkerson algorithm (Sedgewick and Schidlowsky 2003).

Remark 1. In Cancela et al. (2009), we have shown that for the conditional Monte Carlo approach, in the source-to-terminal homogeneous case (that is when all links have the same reliability), there always exists a construction of the set of minpaths (actually in this case paths) \mathcal{P} yielding an “efficient” estimator, but that this is not possible in general in the heterogeneous case. The idea is to also use here the Ford-Fulkerson algorithm which runs in polynomial time in terms of the number of nodes and links. We prove in the next section that thanks to the IS change of measure, when sampling the first failed link on a minpath, robustness properties can be obtained in general (when reliabilities of individual links go to one).

3 ROBUSTNESS ANALYSIS

The goal being to design a method that provides accurate results in the case of rare events, we are interested in how our algorithms behave in asymptotic regimes where the links have failure probabilities that converge to 0. To define such a regime, we introduce a rarity parameter $\varepsilon \ll 1$, as done for example in L'Ecuyer et al. (2010a),

Nakayama (1996), Shahabuddin (1994), and we assume that for each link i , there are positive constants a_i and b_i (independent of ε) such that

$$q_i = a_i \varepsilon^{b_i}. \quad (5)$$

Thus we are investigating situations where the failure probabilities q_i of the individual links are very close to 0. 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

$$q(\mathcal{G}) = q(\mathcal{G}, \varepsilon) = \Theta(\varepsilon^c) \quad (6)$$

for some constant $c > 0$ (see Colbourn (1987), for instance) and $q(\mathcal{G}) \rightarrow 0$ as $\varepsilon \rightarrow 0$. The BRE and VRE properties can then be reformulated in terms of the relative error as $\varepsilon \rightarrow 0$ instead of as $q(\mathcal{G}) \rightarrow 0$.

We have the following robustness property for our combined scheme.

Proposition 1. *If $\hat{q}(\mathcal{G}') = \Theta(q(\mathcal{G}'))$ as $\varepsilon \rightarrow 0$ whatever the graph \mathcal{G}' , then our method yields BRE.*

Proof. The proof is similar to that in L'Ecuyer et al. (2010b), except that we need to separate the following two steps: (i) sampling the first failed link on each minpath and (ii) sampling the state of the remaining links one after the other. Let $a_{j;w_1,\dots,w_j}$ be the constant independent of ε (and strictly larger than 0) such that

$$\hat{q}(\mathcal{G}_{j;w_1,\dots,w_j}) = a_{j;w_1,\dots,w_j} q(\mathcal{G}_{j;w_1,\dots,w_j}) + o(q(\mathcal{G}_{j;w_1,\dots,w_j})).$$

From (4), all the terms $\mathbb{P}[W_j = k] q(\mathcal{G}_{j;w_1,\dots,w_{j-1},k})$ are $O(q(\mathcal{G}_{j-1;w_1,\dots,w_{j-1}}))$, with at least one term $\Theta(q(\mathcal{G}_{j-1;w_1,\dots,w_{j-1}}))$; this means that there exists a constant $b_{j-1;w_1,\dots,w_{j-1}} > 0$ independent of ε verifying

$$\sum_{k=1}^{m_j} \mathbb{P}[W_j = k] \hat{q}(\mathcal{G}_{j;w_1,\dots,w_{j-1},k}) = b_{j-1;w_1,\dots,w_{j-1}} q(\mathcal{G}_{j-1;w_1,\dots,w_{j-1}}) + o(q(\mathcal{G}_{j-1;w_1,\dots,w_{j-1}}));$$

similar arguments were developed in the proof of Theorem 2 of L'Ecuyer et al. (2010b). When sampling the first failed link on the j -th minpath ($1 \leq j \leq h$), the likelihood ratio is

$$\begin{aligned} L_j(W_j) &= \frac{\sum_{k=1}^{m_j} \mathbb{P}[W_j = k] \hat{q}(\mathcal{G}_{j;w_1,\dots,w_{j-1},k})}{\hat{q}(\mathcal{G}_{j;w_1,\dots,w_{j-1},W_j})} \\ &\leq \frac{b_{j-1;w_1,\dots,w_{j-1}} q(\mathcal{G}_{j-1;w_1,\dots,w_{j-1}})}{a_{j;w_1,\dots,w_j} q(\mathcal{G}_{j;w_1,\dots,w_j})} + o(1) \\ &\leq d_1 \frac{q(\mathcal{G}_{j-1;w_1,\dots,w_{j-1}})}{q(\mathcal{G}_{j;w_1,\dots,w_j})} + o(1) \end{aligned}$$

with $d_1 = \max b_{j-1;w_1,\dots,w_{j-1}} / a_{j;w_1,\dots,w_j}$, the maximum being taken over all j and the whole (finite) set of values (w_1, \dots, w_h) (note that as the $a_{j;w_1,\dots,w_j}$ are larger than 0, d_1 is well-defined).

Similarly, from the proof of Theorem 2 of L'Ecuyer et al. (2010b), there exists a constant d_2 independent of ε such that the likelihood ratio, when sampling the state of link i given that all the states of each minpath P_j up to W_j have already been sampled, as well as all links $k \notin \cup_{j=1}^h \{i_{j,1}, \dots, i_{j,W_j}\}$, $k < i$, verifies

$$L_i(X_i) \leq d_2 \frac{q(\mathcal{G}_{i,X_i,\text{hist}})}{q(\mathcal{G}_{i-1,X_{i-1},\text{hist}})} + o(1).$$

We then have from (3),

$$L(X) \leq \prod_{j=1}^h \left(d_1 \frac{q(\mathcal{G}_{j-1;w_1,\dots,w_{j-1}})}{q(\mathcal{G}_{j;w_1,\dots,w_j})} + o(1) \right) \prod_{1 \leq i \leq \ell, i \notin \cup_{j=1}^h \{i_{j,1}, \dots, i_{j,W_j}\}} \left(d_2 \frac{q(\mathcal{G}_{i,X_i,\text{hist}})}{q(\mathcal{G}_{i-1,X_{i-1},\text{hist}})} + o(1) \right) \leq d^\ell q(\mathcal{G}) + o(q(\mathcal{G}))$$

with $d = \max(d_1, d_2, 1)$.

Due to the finite number of terms in the sum (in the expectation formula), we have

$$\tilde{\mathbb{E}}[(Y')^2] \leq \tilde{\mathbb{E}}[(L(X))^2] \leq d^{2\ell}(q(\mathcal{G}))^2 + o((q(\mathcal{G}))^2),$$

meaning that BRE property is verified. \square

Corollary 2. *With the mincut-maxprob approximation, $\hat{q}(\mathcal{G}') = \Theta(q(\mathcal{G}'))$ as $\varepsilon \rightarrow 0$ whatever be the graph \mathcal{G}' . As a consequence, we always get the BRE property with our method.*

Proof. This corollary is a direct consequence of Proposition 1. The fact that $\hat{q}(\mathcal{G}') = \Theta(q(\mathcal{G}'))$ is proved in L'Ecuyer et al. (2010b). It basically comes from the fact that if $q(\mathcal{G}') = \Theta(\varepsilon^{c'})$ for some $c' > 0$, the mincut-maxprob approximation is also $\Theta(\varepsilon^{c'})$ otherwise, there would be a contradiction with its definition). \square

There are situations where not only the BRE property but also the VRE property is verified using the mincut-maxprob approximation. The next proposition provides sufficient (but unnecessary) conditions for that. Define $\mathcal{S}_1 = \{x \in \{0, 1\}^\ell : \phi(x) = 1 \text{ and } \tilde{\mathbb{P}}[X = x] = \Theta(1)\}$ and $\mathcal{S}_0 = \{x \in \{0, 1\}^\ell : \phi(x) = 1 \text{ and } \tilde{\mathbb{P}}[X = x] = o(1)\}$, the sets of configurations under which the system fails and which are no longer rare and still rare, respectively, under IS.

Proposition 2. *Suppose that the assumptions of Proposition 1 hold and that for each configuration of \mathcal{S}_1 , we have:*

- (i) *When sampling the first failed link W_j on minpath P_j , for every j and every w such that*

$$\mathbb{P}[W_j = w] \hat{q}(\mathcal{G}_{j;w_1, \dots, w_{j-1}, w}) = \Theta\left(\sum_{k=1}^{m_j} \mathbb{P}[W_j = k] \hat{q}(\mathcal{G}_{j;w_1, \dots, w_{j-1}, k})\right),$$

we have

$$\hat{q}(\mathcal{G}_{j;w_1, \dots, w_{j-1}, w}) / q(\mathcal{G}_{j;w_1, \dots, w_{j-1}, w}) \rightarrow c_j(w_1, \dots, w_{j-1})$$

as $\varepsilon \rightarrow 0$, where $c_j(w_1, \dots, w_{j-1})$ is independent of w , and $c_j(w_1, \dots, w_{j-1}) = 1$ if there is only one such w .

- (ii) *When sampling the remaining links, for all $x = (x_1, \dots, x_\ell) \in \mathcal{S}_1$, one of the following three conditions of Theorem 5 of L'Ecuyer et al. (2010b) is verified:*

either

$$\frac{\hat{q}(\mathcal{G}_{i,1,\text{hist}})}{q(\mathcal{G}_{i,1,\text{hist}})} = \frac{\hat{q}(\mathcal{G}_{i,0,\text{hist}})}{q(\mathcal{G}_{i,0,\text{hist}})} + o(1),$$

or $x_i = 0$, $\hat{q}(\mathcal{G}_{i,0,\text{hist}}) = q(\mathcal{G}_{i,0,\text{hist}}) + (q(\mathcal{G}_{i,0,\text{hist}}))$, and $(1 - q_i)\hat{q}(\mathcal{G}_{i,1,\text{hist}}) = o(q_i\hat{q}(\mathcal{G}_{i,0,\text{hist}}))$

or $x_i = 1$, $\hat{q}(\mathcal{G}_{i,1,\text{hist}}) = q(\mathcal{G}_{i,1,\text{hist}}) + (q(\mathcal{G}_{i,1,\text{hist}}))$, and $q_i\hat{q}(\mathcal{G}_{i,0,\text{hist}}) = o((1 - q_i)\hat{q}(\mathcal{G}_{i,1,\text{hist}}))$.

Then, the VRE property is verified.

Proof. The second moment is decomposed in two terms using the sets of configurations \mathcal{S}_1 and \mathcal{S}_0 :

$$\tilde{\mathbb{E}}[\phi(X)L^2(X)] = \sum_{x \in \mathcal{S}_1} \phi(x)L^2(x)\tilde{\mathbb{P}}[X = x] + \sum_{x \in \mathcal{S}_0} \phi(x)L^2(x)\tilde{\mathbb{P}}[X = x]. \quad (7)$$

As in the proof of Proposition 1, there exists a constant $d > 0$, independent of ε , such that $L^2(x) \leq d^{2m}(q(\mathcal{G}))^2 + o((q(\mathcal{G}))^2)$, and, because \mathcal{S}_0 is finite,

$$\sum_{x \in \mathcal{S}_0} \phi(x)L^2(x)\tilde{\mathbb{P}}[X = x] \leq (d^{2m}(q(\mathcal{G}))^2 + o((q(\mathcal{G}))^2)) \sum_{x \in \mathcal{S}_0} \phi(x)\tilde{\mathbb{P}}[X = x],$$

with this last sum being $o(1)$, as a finite sum of terms that are all $o(1)$. Consequently,

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

For the first term on the right side of (7), on \mathcal{S}_1 it has been proved in L'Ecuyer et al. (2010b) that $d_2 = 1$ in the proof of Proposition 1, under our assumption (ii). For $x \in \mathcal{S}_1$, it is easy to check that the corresponding value w of W_j is necessarily such that $\mathbb{P}[W_j = w]\hat{q}(\mathcal{G}_{j;w_1,\dots,w_{j-1},w}) = \Theta(\sum_{k=1}^{m_j} \mathbb{P}[W_j = k]\hat{q}(\mathcal{G}_{j;w_1,\dots,w_{j-1},k}))$, otherwise $\tilde{\mathbb{P}}[W_j = k] = o(1)$ and so would be $\tilde{\mathbb{P}}[X = x]$. Let $\mathcal{M}_j = \{w \in \{1, \dots, m_j\} : \mathbb{P}[W_j = w]\hat{q}(\mathcal{G}_{j;w_1,\dots,w_{j-1},w}) = \Theta(\sum_{k=1}^{m_j} \mathbb{P}[W_j = k]\hat{q}(\mathcal{G}_{j;w_1,\dots,w_{j-1},k}))\}$ be the set of indices for which $[W_j = k]$ is not rare under IS. We then have from the first bullet

$$\begin{aligned} L_j(W_j) &= \frac{\sum_{k=1}^{m_j} \mathbb{P}[W_j = k]\hat{q}(\mathcal{G}_{j;W_1,\dots,W_{j-1},k})}{\hat{q}(\mathcal{G}_{j;W_1,\dots,W_{j-1},W_j})} \\ &= \frac{\sum_{k \in \mathcal{M}_j} \mathbb{P}[W_j = k]\hat{q}(\mathcal{G}_{j;W_1,\dots,W_{j-1},k}) + o(1)}{\hat{q}(\mathcal{G}_{j;W_1,\dots,W_{j-1},W_j})} \\ &= \frac{\sum_{k \in \mathcal{M}_j} \mathbb{P}[W_j = k]c_j(W_1, \dots, W_{j-1})q(\mathcal{G}_{j;W_1,\dots,W_{j-1},k}) + o(1)}{c_j(W_1, \dots, W_{j-1})q(\mathcal{G}_{j;W_1,\dots,W_{j-1},W_j}) + o(1)} \\ &= \frac{\sum_{k \in \mathcal{M}_j} \mathbb{P}[W_j = k]q(\mathcal{G}_{j;W_1,\dots,W_{j-1},k}) + o(1)}{q(\mathcal{G}_{j;W_1,\dots,W_{j-1},W_j}) + o(1)} \\ &= \frac{q(\mathcal{G}_{j-1;W_1,\dots,W_{j-1}})}{q(\mathcal{G}_{j;W_1,\dots,W_j})} + o(1). \end{aligned}$$

This gives $d_1 = 1$ in the proof of Proposition 1 and

$$\sum_{x \in \mathcal{S}_1} \phi(x)L^2(x)\tilde{\mathbb{P}}[X = x] = (q(\mathcal{G}))^2 + o((q(\mathcal{G}))^2). \tag{8}$$

Combining the results for the two terms, we have that $\tilde{\text{Var}}[\phi(X)L(X)] \leq \tilde{\mathbb{E}}[L^2(X)] - (q(\mathcal{G}))^2 = o((q(\mathcal{G}))^2)$. \square

Remark 2. Sufficient conditions for getting the BRE property for the above conditional Monte Carlo method were derived in Cancela, Rubino, and Tuffin (2005). When additionally using zero-variance approximation, we obtain the stronger property that BRE is *always* verified, and even VRE can be obtained in some cases.

4 NUMERICAL ILLUSTRATIONS

Example 1. Consider the small example of Figure 1, where we want to compute the probability that the gray nodes A and D are disconnected. The links are assumed homogeneous, with unreliabilities $q_i = \varepsilon$ for $i = 1, \dots, 5$. We consider $\mathcal{P} = \{\{1, 4\}, \{2, 5\}\}$. Here $p = (1 - (1 - \varepsilon)^2)^2 = (2\varepsilon - \varepsilon^2)^2$. For this small model, one can easily list the 32 possible configurations, identify those that lead to disconnected nodes A and D (there are 16), compute their original probabilities and their probabilities under IS, compute their contribution to the second moment of random variable Y , and then compute the exact unreliability and exact variances. The exact value is $q(\mathcal{G}) = 2\varepsilon^2 + 2\varepsilon^3 - 5\varepsilon^4 + 2\varepsilon^5$. For the conditional Monte Carlo (CMC) method, $\sigma^2 = 4\varepsilon^4 - 8\varepsilon^5 - 10\varepsilon^6 + 42\varepsilon^7 - 46\varepsilon^8 + 22\varepsilon^9 - 4\varepsilon^{10}$, while for the combination method, $\hat{\sigma}^2 = 4\varepsilon^5 - 15\varepsilon^6 + 20\varepsilon^7 - 9\varepsilon^8 - 3\varepsilon^9 + 4\varepsilon^{10} - \varepsilon^{11}$. We see that VRE is verified for the Combination method, whereas only BRE is verified for CMC. One can check from L'Ecuyer et al. (2010b) that the Combination method and ZVA give *exactly* the same IS probability whatever the configuration. As a consequence, the variance is also exactly the same.

Example 2. We now consider an example with the topology depicted in Figure 2, where the source s and the terminal node t are connected through h disjoint paths $P_1 \dots, P_h$, having no node in common excepted s and t . In such a model, it is immediate that $p = \mathbb{P}(C) = q$, since, if every path in $\mathcal{P} = \{P_1 \dots, P_h\}$ has at least one link down, there is no way to joint t from s , and reciprocally, if s and t are unconnected, then necessarily at least




Figure 1: Graph topology with five links and two nodes requiring to be connected

one link in any of the h paths has a failed link. This means that using the CMC method, by definition, the sum of the probabilities of the failing configurations under the CMC sampling procedure is 1 (remember that we are conditioning with respect to the event C), and then, the variance of the CMC estimator is

$$\sigma^2 = p^2 \sum_{x \in \{0,1\}^5} \phi(x) \mathbb{P}[X = x] - (q(\mathcal{G}))^2 = p^2 - (q(\mathcal{G}))^2 = 0.$$

In this type of topology, the combination method will sample exactly as the CMC one, leading to a zero-variance estimator as well.




Figure 2: Graph topology made of three direct and independent paths

Let us now illustrate numerically the power of the method on three examples. The first two, the Arpanet and the dodecahedron topologies, are benchmark topologies used in the literature (Cancela, El Khadiri, and Rubino 2009). The third one is a topology made of direct paths with few arcs connecting the paths. In the tables, the methods are abbreviated by CMC for the conditional Monte Carlo method based on minpaths but without IS, ZVA for the zero-variance IS approximation of L'Ecuyer et al. (2010b) where no conditional expectation is used, i.e., the method for which (2) is used successively for all links, and “Combination” when both methods are combined as described earlier. This will help weighting the relative importance of IS and conditioning. In all illustrations, we consider the homogeneous case, where all links have the same unreliability ϵ , and we want to compute the probability that nodes s and t are disconnected. In the general non-homogeneous case, ZVA and “Combination” are known to yield at least BRE, which is not the case for CMC; we focus on the homogeneous case to illustrate the potential gain in the less advantageous situation.

Example 3. The first topology, presented in Figure 3, is the topology of a version of the Arpanet (“Advanced Research Projects Agency Network”), the first packet-based network in the USA, at the beginning of its life-time. Links are ordered as displayed in Figure 3. We compare in Table 1 the results obtained for our method with the




Figure 3: A version of the Arpanet topology.

cases where only zero-variance approximation is used and only conditional Monte Carlo based on the selected paths is used, this for various values of ϵ . We consider two disjoint paths, the first one made of links $\{1, 3, 6, 9, 11, 12\}$, and the second one made of links $\{2, 20, 19, 21, 22, 23, 24, 25, 26\}$.

Table 1: Empirical results for the Arpanet topology, for $n = 10^4$ and four values of ϵ .

Method	ϵ	Estimate	95% Confidence interval	Std deviation	Relative error
CMC	10^{-1}	9.406×10^{-2}	$(9.142 \times 10^{-2}, 9.670 \times 10^{-2})$	1.347×10^{-1}	1.43
CMC	10^{-3}	5.875×10^{-6}	$(5.546 \times 10^{-6}, 6.203 \times 10^{-6})$	1.675×10^{-5}	2.85
CMC	10^{-5}	5.837×10^{-10}	$(5.508 \times 10^{-10}, 6.166 \times 10^{-10})$	1.677×10^{-9}	2.87
CMC	10^{-6}	5.837×10^{-12}	$(5.509 \times 10^{-12}, 6.166 \times 10^{-12})$	1.677×10^{-11}	2.87
ZVA	10^{-1}	9.295×10^{-2}	$(8.469 \times 10^{-2}, 1.0121 \times 10^{-1})$	4.214×10^{-1}	4.53
ZVA	10^{-3}	5.956×10^{-6}	$(5.794 \times 10^{-6}, 6.118 \times 10^{-6})$	8.260×10^{-6}	1.39
ZVA	10^{-5}	5.897×10^{-10}	$(5.748 \times 10^{-10}, 6.045 \times 10^{-10})$	7.575×10^{-10}	1.28
ZVA	10^{-6}	5.897×10^{-12}	$(5.748 \times 10^{-12}, 6.045 \times 10^{-12})$	7.575×10^{-12}	1.28
Combination	10^{-1}	9.245×10^{-2}	$(8.780 \times 10^{-2}, 9.710 \times 10^{-2})$	2.372×10^{-1}	2.57
Combination	10^{-3}	6.093×10^{-6}	$(6.007 \times 10^{-6}, 6.180 \times 10^{-6})$	4.408×10^{-6}	0.723
Combination	10^{-5}	6.065×10^{-10}	$(5.987 \times 10^{-10}, 6.144 \times 10^{-10})$	4.000×10^{-10}	0.659
Combination	10^{-6}	6.065×10^{-12}	$(5.986 \times 10^{-12}, 6.143 \times 10^{-12})$	4.000×10^{-12}	0.659

The results shown are based on a sample size of $n = 10^4$. The first column of the table indicates the methods employed. The second one shows the values of ϵ considered. The third one gives the estimates of the network unreliability computed by the different methods. The fourth column gives a 95% confidence interval, and the fifth and sixth ones give the standard deviation and the relative error for a single random variate. The results show a large degree of agreement in the estimations computed by all the methods. The CMC method results are consistent with a BRE situation, as the relative error first grows but then gets stable when ϵ approaches zero. Both the ZVA and the combined methods have higher relative error than CMC for the less reliable network ($\epsilon = 10^{-1}$), but improve substantially for smaller values of this parameter. Their behavior is also consistent with a BRE situation. The combined method has the best results, attaining a relative error almost half of the ZVA one.

Example 4. We now look at another classical topology, made of 20 nodes and 30 links, the dodecahedron topology as shown in Figure 4. Links are ordered arbitrarily, according to their numbering in the figure. The set \mathcal{P} is made of three paths, $\{3, 9, 18, 25, 26, 30\}$, $\{1, 5, 12, 19, 20, 28\}$ and $\{2, 7, 15, 23, 22, 29\}$. The empirical results appear in Table 2.

The results for ZVA algorithm are taken from L'Ecuyer et al. (2010b). The confidence intervals show again good agreement between the different estimations. As in the previous case, CMC provides a relative error which first grows and then seems to approach an upper bound, as could be expected in a BRE situation. On the other hand, the ZVA and the combined method yield similar results, both with much smaller relative errors than CMC. The relative errors moreover diminish substantially when ϵ gets smaller, as would be expected in a VRE situation.




Figure 4: Dodecahedron topology

Table 2: Empirical results for the dodecahedron topology, for $n = 10^4$ and four values of ϵ .

Method	ϵ	Estimate	95% Confidence interval	Std deviation	Relative error
CMC	10^{-1}	2.531×10^{-3}	$(2.218 \times 10^{-3}, 2.843 \times 10^{-3})$	1.594×10^{-2}	6.30
CMC	10^{-2}	1.864×10^{-6}	$(1.487 \times 10^{-6}, 2.241 \times 10^{-6})$	1.924×10^{-5}	10.3
CMC	10^{-3}	1.844×10^{-9}	$(1.456 \times 10^{-9}, 2.232 \times 10^{-9})$	1.980×10^{-8}	10.7
CMC	10^{-4}	1.856×10^{-12}	$(1.466 \times 10^{-12}, 2.247 \times 10^{-12})$	1.993×10^{-11}	10.7
ZVA	10^{-1}	2.896×10^{-3}	$(2.8276 \times 10^{-3}, 2.9645 \times 10^{-3})$	3.491×10^{-3}	1.2
ZVA	10^{-2}	2.0678×10^{-6}	$(2.0611 \times 10^{-6}, 2.0744 \times 10^{-6})$	3.425×10^{-7}	0.17
ZVA	10^{-3}	2.0076×10^{-9}	$(2.0053 \times 10^{-9}, 2.0099 \times 10^{-9})$	1.145×10^{-10}	0.057
ZVA	10^{-4}	2.0007×10^{-12}	$(2.0000 \times 10^{-12}, 2.0014 \times 10^{-12})$	3.464×10^{-14}	0.017
Combination	10^{-1}	2.901×10^{-3}	$(2.839 \times 10^{-3}, 2.963 \times 10^{-3})$	3.147×10^{-3}	1.08
Combination	10^{-2}	2.063×10^{-6}	$(2.054 \times 10^{-6}, 2.072 \times 10^{-6})$	4.607×10^{-7}	0.223
Combination	10^{-3}	2.008×10^{-9}	$(2.0053 \times 10^{-9}, 2.0111 \times 10^{-9})$	1.486×10^{-10}	0.0740
Combination	10^{-4}	2.0019×10^{-12}	$(1.9996 \times 10^{-12}, 2.00096 \times 10^{-12})$	3.466×10^{-14}	0.0173

Example 5. Consider finally the topology of Figure 2 where we add an arc from the third node on the highest path (between links 2 and 3) to the node just below it, and another one connecting the node between links 8 and 9 and the node just below it. We will refer to this topology as the “almost direct” one. Table 3 displays the results. BRE is actually verified in each case, but the smallest variance is obtained for the combined method. The improvement in the relative error is already significant for the less reliable case ($\epsilon = 10^{-2}$), and attains an order of magnitude in the most reliable case considered ($\epsilon = 10^{-4}$).

We have compared the accuracy of estimators for fixed values of sample sizes. On the other hand, computational times are also different due for instance to the time required to find mincuts with maximal probability at each step of the sampling process. Table 4 displays the relative time for the ZVA and combination techniques, compared with CMC, for all the above examples. We see that the additional CPU time is non-negligible when using importance sampling, but the BRE property is ensured, and even VRE can be obtained in some cases.

Table 3: Empirical results for the topology of Figure 2 with two additional links between direct paths, for $n = 10^4$ and two values of ϵ .

Method	ϵ	Estimate	95% Confidence interval	Std deviation	Relative error
CMC	10^{-2}	3.271×10^{-5}	$(3.168 \times 10^{-5}, 3.375 \times 10^{-5})$	5.273×10^{-5}	1.61
CMC	10^{-4}	3.273×10^{-11}	$(3.165 \times 10^{-11}, 3.381 \times 10^{-11})$	5.494×10^{-11}	1.69
ZVA	10^{-2}	3.300×10^{-5}	$(3.220 \times 10^{-5}, 3.380 \times 10^{-5})$	4.073×10^{-5}	0.81
ZVA	10^{-4}	3.255×10^{-11}	$(3.191 \times 10^{-11}, 3.319 \times 10^{-11})$	3.259×10^{-11}	1.02
Comb	10^{-2}	3.317×10^{-5}	$(3.289 \times 10^{-5}, 3.344 \times 10^{-5})$	1.411×10^{-5}	0.42
Comb	10^{-4}	3.306×10^{-11}	$(3.293 \times 10^{-11}, 3.317 \times 10^{-11})$	6.153×10^{-12}	0.185

Table 4: Empirical relative CPU times for ZVA and the combined methods with respect to CMC, when $n = 10^4$ and for various values of ϵ .

Topology	ϵ	ZVA	Combination
Arpanet	10^{-1}	57.5	100
Arpanet	10^{-3}	63.5	90.4
Arpanet	10^{-5}	63.9	91.1
Arpanet	10^{-6}	66.3	90.5
Dodecahedron	10^{-1}	71.8	110
Dodecahedron	10^{-2}	70.7	93.5
Dodecahedron	10^{-3}	70.6	91.2
Dodecahedron	10^{-4}	70.5	90.5
Almost direct	10^{-2}	35.6	60.2
Almost direct	10^{-4}	36.2	61.2

ACKNOWLEDGMENTS

This work has been supported by INRIA's associated team MOCQUASIN to all authors, and NSERC-Canada and a Canada Research Chair to the second author.

REFERENCES

Asmussen, S., and P. W. Glynn. 2007. *Stochastic simulation*. New York: Springer-Verlag.

Ball, M. O. 1986, Aug.. Computational complexity of network reliability analysis: An overview. *IEEE Transactions on Reliability* 35 (3): 230–239.

Cancela, H., M. El Khadiri, and G. Rubino. 2009. Rare event analysis by monte carlo techniques in static models. In *Rare Event Simulation Using Monte Carlo Methods*, ed. G. Rubino and B. Tuffin, 145–170. Wiley. Chapter 7.

Cancela, H., P. L'Ecuyer, M. Lee, G. Rubino, and B. Tuffin. 2009. 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*, ed. S. M. J. Faulin, A. A. Juan and E. Ramirez-Marquez. Springer Verlag.

Cancela, H., G. Rubino, and B. Tuffin. 2005. New measures of robustness in rare event simulation. In *Proceedings of the 2005 Winter Simulation Conference*, ed. M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines, 519–527: IEEE Press.

Colbourn, C. J. 1987. *The combinatorics of network reliability*. New York: Oxford University Press.

L'Ecuyer, P., J. H. Blanchet, B. Tuffin, and P. W. Glynn. 2010a. Asymptotic robustness of estimators in rare-event simulation. *ACM Transactions on Modeling and Computer Simulation* 20 (1): Article 6.

L'Ecuyer, P., G. Rubino, S. Saggadi, and B. Tuffin. 2010b. Approximate zero-variance importance sampling for static network reliability estimation. Submitted. Available at http://www.irisa.fr/dionysos/pages_perso/tuffin/Publis/static-IS.pdf.

- L'Ecuyer, P., and B. Tuffin. 2008. Approximate zero-variance simulation. In *Proceedings of the 2008 Winter Simulation Conference*, 170–181: IEEE Press.
- Nakayama, M. K. 1996. General conditions for bounded relative error in simulations of highly reliable Markovian systems. *Advances in Applied Probability* 28:687–727.
- Rubino, G., and B. Tuffin. (Eds.) 2009. *Rare event simulation using Monte Carlo methods*. John Wiley & Sons.
- Sedgewick, R., and M. Schidlowsky. 2003. *Algorithms in java, part 5: Graph algorithms*. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc.
- Shahabuddin, P. 1994. Importance sampling for the simulation of highly reliable Markovian systems. *Management Science* 40 (3): 333–352.

AUTHOR BIOGRAPHIES

HECTOR CANCELA holds a PhD. degree in Computer Science from the University of Rennes 1, INRIA Rennes, France (1996), and a Computer Systems Engineer degree from the Universidad de la República, Uruguay (1990). He is currently Full Professor and Director of the Computer Science Institute at the Engineering School of the Universidad de la República (Uruguay). He is also a Researcher at the National Program for the Development of Basic Sciences (PEDECIBA), Uruguay. His research interests are in Operations Research techniques, especially in stochastic process models and graph and network models, and in their application jointly with combinatorial optimization metaheuristics to solve different practical problems. He has published more than 40 papers in international journals and proceedings of refereed conferences. Prof. Cancela has participated as academic delegate at the National Committee for the Information Society hosted by the Presidency of the Republic of Uruguay. He has also been a national delegate at the Committee on Accreditation Standards for Engineering Studies at the MERCOSUR. He is currently President of ALIO, the Latin American Operations Research Association, and member of IFIP System Modeling and Optimization technical committee (TC7) . His email address is <cancela@fing.edu.uy>.

PIERRE L'ECUYER is Professor in the Département d'Informatique et de Recherche Opérationnelle, at the Université de Montréal, Canada. He holds the Canada Research Chair in Stochastic Simulation and Optimization. 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 Editor-in-Chief for *ACM Transactions on Modeling and Computer Simulation*, and Associate/Area Editor for *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. More information and 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) research 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, the main component of Quality of Experience. He recently co-edited with Bruno Tuffin 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. His email address is <rubino@inria.fr>.

BRUNO TUFFIN received his PhD degree in applied mathematics from the University of Rennes 1 (France) in 1997. Since then, he has been with INRIA in Rennes. He spent 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 email address is <bruno.tuffin@inria.fr>.