

On the estimation of the mean time to failure by simulation

Peter Glynn, Marvin Nakayama, Bruno Tuffin

► **To cite this version:**

Peter Glynn, Marvin Nakayama, Bruno Tuffin. On the estimation of the mean time to failure by simulation. 2017 Winter Simulation Conference, Dec 2017, Las Vegas, United States. hal-01551078

HAL Id: hal-01551078

<https://hal.inria.fr/hal-01551078>

Submitted on 29 Jun 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On the estimation of the mean time to failure by simulation

Peter W. Glynn

Department of Management Science and Engineering
Stanford University
475 Via Ortega
Stanford, CA 94305, USA

Marvin K. Nakayama

Department of Computer Science
New Jersey Institute of Technology
Newark, NJ 07102, USA

Bruno Tuffin

Inria
Campus de Beaulieu, 263 Avenue Général Leclerc
35042 Rennes, FRANCE

ABSTRACT

The mean time to failure (MTTF) of a stochastic system is often estimated by simulation. One natural estimator, which we call the direct estimator, simply averages independent and identically distributed copies of simulated times to failure. When the system is regenerative, an alternative approach is based on a ratio representation of the MTTF. The purpose of this paper is to compare the two estimators. We first analyze them in the setting of crude simulation (i.e., no importance sampling), showing that they are actually asymptotically identical in a rare-event context. The two crude estimators are inefficient in different but closely related ways: the direct estimator requires a large computational time because times to failure often include many transitions, whereas the ratio estimator entails estimating a rare-event probability. We then discuss the two approaches when employing importance sampling; for highly reliable Markovian systems, we show that using a ratio estimator is advised.

1 INTRODUCTION

Dependability analysis is of primary importance in many areas, such as nuclear power plants, telecommunications, manufacturing, transport systems, and computer science; for examples, see [8] and [14]. Even if system failures are rare, their occurrence may have dramatic consequences and therefore need to be analyzed with care. We focus here on one common dependability metric, the *mean time to failure* (MTTF), which is the expected value of the random time to reach failure.

An example of the type of system we are considering is one with components subject to failures and repairs exponentially distributed over time. Such a system is then represented by a Markov chain which can in principle be solved analytically, but for practical problems the state space is usually so large that it would require an enormous computation time. Monte Carlo simulation then becomes a relevant option.

A crude simulation of the model entails simulating failures and repairs of components up to the failure of the whole system. We obtain the direct estimator of the MTTF by repeating the experiment many independent times and averaging the obtained times to system failure. But in the case when individual components are highly reliable (in the sense that failure rates are much smaller than the repair rates), this often requires a very long computation time because it involves, with high probability, a large number of transitions before a failure of the system since when components are failed, it is more likely to have repairs than other failures. In the literature, another estimator is often instead used. It exploits the regenerative structure of the model and expresses the MTTF as a ratio of quantities over regenerative cycles. Estimating by crude simulation the numerator in this expression is efficient, but it is not the case for the denominator, which is the probability of a rare event. Many rare-event simulation techniques have been developed to obtain efficient estimators using the ratio expression [8, 12, 13].

The purpose of this paper is to review and discuss the relative merits of the two estimators: the direct and a ratio-based one. We highlight the following results we obtain:

- We first show that crude estimators based on direct simulation of times to failure and on the ratio expression are asymptotically similar in performance, in rare-event settings. Both estimators are inefficient, suffering from different but closely related issues: the direct estimator requires large computation times because replications are often very long, whereas the ratio estimator encounters difficulties from estimating a rare-event probability.
 - To analyze the asymptotics as the event of interest becomes rarer, we consider a sequence $(A_b : b \geq 1)$ of (failure) sets, with $\mathbb{P}(A_b) \rightarrow 0$ as

$b \rightarrow \infty$. We prove that the two estimators are asymptotically equivalent when estimating the expected hitting time to A_b as $b \rightarrow \infty$.

- Moreover, we provide numerical results that the same is true for highly reliable Markovian systems, in which the asymptotic regime differs in that the failure set is fixed but the component failure rates shrink with the repair rates fixed.
- Given that crude estimators are equivalent in performance, we next compare the *importance sampling* (IS) versions of the estimators. We show that in the setting of highly reliable Markovian systems, it is not possible in full generality to design efficient direct IS estimators, so the ratio-based estimators are then rather advised.

The rest of the paper develops as follows. Section 2 presents the two versions of the crude estimators of the MTTF, with their associated central limit theorems (CLTs), from which one can construct asymptotically valid confidence intervals. Section 3 compares the two estimators in rare-event settings and shows that they are equivalent in terms of accuracy as the computational budget grows large. Instead of the well-known IS ratio estimation, Section 4 discusses direct estimators via IS. Based on simple examples, we show that, and explain why, designing efficient direct estimators is difficult. Section 5 concludes the paper.

2 CRUDE MTTF ESTIMATORS

2.1 The estimators

Let $X = (X(t) : t \geq 0)$ be an S -valued non-delayed (classically) regenerative process [18] with regeneration times $0 = \Gamma(0) < \Gamma(1) < \dots$. For $k \geq 1$, let $\tau(k) = \Gamma(k) - \Gamma(k-1)$ be the length of the k th regenerative cycle. Given a set $A \subset S$, the goal is to compute $\alpha = \mathbb{E}[T]$, where $T = \inf\{t \geq 0 : X(t) \in A\}$ is the hitting time of A and $\mathbb{E}[\cdot]$ is the expectation operator. When A corresponds to the set of states for which the simulated system is failed, α represents the MTTF. We assume throughout that $\mathbb{E}[\tau^2(1)] < \infty$.

Because X is classically regenerative, we have that

$$((\tau(k), (X(\Gamma(k-1) + s) : 0 \leq s < \tau(k)) : k \geq 1)$$

is a sequence of independent and identically distributed (IID) cycles. For real-valued x and y , define $x \wedge y = \min(x, y)$. For $k \geq 1$, let $W(k) = \inf\{t \geq 0 : X(\Gamma(k-1) + t) \in A\}$ be the first hitting to A after regeneration time $\Gamma(k-1)$. The classical regenerative property implies that $((\tau(k), W(k) \wedge \tau(k), I(k)) : k \geq 1)$ is an IID

sequence of triplets, where $I(k) = \mathcal{I}(W(k) < \tau(k))$ and $\mathcal{I}(\cdot)$ is the indicator function. Define $\tau = \tau(1)$, $W = W(1) = T$, and $p = \mathbb{P}(T < \tau)$. A proof of the following ratio representation for α appears, e.g., in [7].

Proposition 1 If $p > 0$, then

$$\alpha = \frac{\mathbb{E}[T \wedge \tau]}{\mathbb{P}(T < \tau)}. \quad (1)$$

Set $N(0) = 0$, and for $j \geq 1$, let $N(j) = \inf\{k > N(j-1) : I(k) = 1\}$ be the index k of the cycle corresponding to the j th cycle in which A is hit. Let $T(1) = T$, and $T(j) = \inf\{t \geq 0 : X(\Gamma(N(j-1)) + t) \in A\}$ for $j \geq 2$. Then α can be estimated either by the *direct estimator*

$$\alpha_1(m) = \frac{1}{m} \sum_{j=1}^m T(j)$$

or via the *ratio estimator*

$$\alpha_2(n) = \frac{(1/n) \sum_{k=1}^n [W(k) \wedge \tau(k)]}{(1/n) \sum_{k=1}^n I(k)}.$$

2.2 Central Limit Theorems

Let \Rightarrow denote weak convergence (e.g., [2]), and let $\mathcal{N}(a, s^2)$ be a normal random variable with mean a and variance s^2 . Then the direct estimator $\alpha_1(m)$ satisfies the following CLT.

Proposition 2 If $p > 0$, then

$$m^{1/2}[\alpha_1(m) - \alpha] \Rightarrow \sigma_1 \mathcal{N}(0, 1)$$

as $m \rightarrow \infty$, where

$$\sigma_1^2 = \alpha^2 + \frac{\mathbb{E}[(T \wedge \tau)^2]}{p} - 2\alpha \frac{\mathbb{E}[T \mathcal{I}(T < \tau)]}{p}. \quad (2)$$

Proof. Because $p > 0$ and $\mathbb{E}[\tau^2] < \infty$, all the expectations below are finite. Note that $\sigma_1^2 = \mathbb{E}[T^2] - \alpha^2$ because $\alpha_1(m)$ averages IID copies of T . To show that σ_1^2 satisfies (2), observe that $T \stackrel{\mathcal{D}}{=} (T \wedge \tau) + \mathcal{I}(T \geq \tau)T'$, where $\stackrel{\mathcal{D}}{=}$ denotes ‘‘equality in distribution’’ and $T' \stackrel{\mathcal{D}}{=} T$ is independent of $(T \wedge \tau, \mathcal{I}(T \geq \tau))$. Hence, $\mathbb{E}[T^2] = \mathbb{E}[(T \wedge \tau)^2] + 2\mathbb{E}[(T \wedge \tau)\mathcal{I}(T \geq \tau)]\mathbb{E}[T] + \mathbb{E}[\mathcal{I}(T \geq \tau)]\mathbb{E}[T^2]$. Therefore, because

$p = 1 - \mathbb{E}[\mathcal{I}(T \geq \tau)]$ and $\alpha = \mathbb{E}[T] = \mathbb{E}[T \wedge \tau]/p$ by (1), we get

$$\begin{aligned}\mathbb{E}[T^2] &= \frac{\mathbb{E}[(T \wedge \tau)^2]}{p} + 2\alpha \frac{\mathbb{E}[(T \wedge \tau)\mathcal{I}(T \geq \tau)]}{p} \\ &= \frac{\mathbb{E}[(T \wedge \tau)^2]}{p} + 2\alpha \left(\frac{\mathbb{E}[T \wedge \tau]}{p} - \frac{\mathbb{E}[T\mathcal{I}(T < \tau)]}{p} \right) \\ &= \frac{\mathbb{E}[(T \wedge \tau)^2]}{p} + 2\alpha^2 - 2\alpha \frac{\mathbb{E}[T\mathcal{I}(T < \tau)]}{p},\end{aligned}$$

establishing (2). \square

On the other hand, the second estimator $\alpha_2(n)$ satisfies the following CLT, which also appears in [7] but we include its proof here for completeness.

Proposition 3 If $p > 0$, then

$$n^{1/2}[\alpha_2(n) - \alpha] \Rightarrow \sigma_2 \mathcal{N}(0, 1)$$

as $n \rightarrow \infty$, where

$$\sigma_2^2 = \frac{\mathbb{E}[(T \wedge \tau)^2]}{p^2} - 2\alpha \frac{\mathbb{E}[T\mathcal{I}(T < \tau)]}{p^2} + \frac{\alpha^2}{p}. \quad (3)$$

Proof. Note that

$$n^{1/2}[\alpha_2(n) - \alpha] = n^{-1/2} \frac{\sum_{k=1}^n [(W(k) \wedge \tau(k)) - \alpha I(k)]}{\sum_{k=1}^n I(k)/n}.$$

But $n^{-1} \sum_{k=1}^n I(k) \rightarrow p$ almost surely as $n \rightarrow \infty$, and $n^{-1/2} \sum_{k=1}^n [(W(k) \wedge \tau(k)) - \alpha I(k)] \Rightarrow \tilde{\sigma}_2 \mathcal{N}(0, 1)$ as $n \rightarrow \infty$, where

$$\tilde{\sigma}_2^2 = \mathbb{E}[(T \wedge \tau - \alpha \mathcal{I}(T < \tau))^2] = \mathbb{E}[(T \wedge \tau)^2] - 2\alpha \mathbb{E}[T\mathcal{I}(T < \tau)] + \alpha^2 \mathbb{P}(T < \tau),$$

so (3) holds by Slutsky's theorem. \square

3 COMPARISON OF THE CRUDE ESTIMATORS

Although the CLTs in Propositions 2 and 3 for $\alpha_1(m)$ and $\alpha_2(n)$, respectively, may appear to be different, they actually are very similar. In fact, they agree at the instants at which the $T(j)$'s occur (i.e., hitting times of A), which [17] also note.

Proposition 4 For $m \geq 1$, we have

$$\alpha_2(N(m)) = \alpha_1(m).$$

Proof. For $j \geq 1$, note that $T(j) = \tau(N(j-1) + 1) + \tau(N(j-1) + 2) + \dots + \tau(N(j) - 1) + W(N(j))$. We have that $\tau(k) < W(k)$ when $N(j-1) < k < N(j)$, whereas $W(k) < \tau(k)$ for $k = N(j)$. Thus, $T(j) = \sum_{k=N(j-1)+1}^{N(j)} \tau(k) \wedge W(k)$. Also, observe that $m = \sum_{k=1}^{N(m)} I(k)$, so that

$$\alpha_1(m) = \frac{\sum_{l=1}^m \sum_{k=N(j-1)+1}^{N(j)} \tau(k) \wedge W(k)}{\sum_{k=1}^{N(m)} I(k)} = \alpha_2(N(m)).$$

□

3.1 Equivalence with a decreasing sequence of reachable sets

We will now show that these estimators are asymptotically identical in the rare-event setting in which A is a “rare” set. Consider a sequence $(A_b : b \geq 1)$ of subsets of S for which $p_b \equiv \mathbb{P}(T_b < \tau) \rightarrow 0$ as $b \rightarrow \infty$, where $T_b = \inf\{t \geq 0 : X(t) \in A_b\}$. For each fixed b , we can then define $W_b(k)$, $I_b(k)$, $N_b(j)$, $T_b(j)$, etc., analogously to $W(k)$, $I(k)$, $N(j)$, $T(j)$, etc., but with A_b instead of A ; e.g., $W_b(k) = \inf\{t \geq 0 : X(\Gamma(k-1) + t) \in A_b\}$ and $I_b(k) = \mathcal{I}(W_b(k) < \tau(i))$. Suppose that it takes $c > 0$ units of computer time to generate c simulated time units of the process X , where the computer time units may differ from the time units of X , e.g., milliseconds vs. days. The number of the $T_b(j)$ generated in c units of computer time is then given by $\beta_b(c) = \sup\{m \geq 0 : \sum_{j=1}^m T_b(j) \leq c\}$, so that the estimator (that is analogous to $\alpha_1(m)$) available after expending c units of computer time is

$$\hat{\alpha}_{1,b}(c) = \frac{1}{\beta_b(c)} \sum_{j=1}^{\beta_b(c)} T_b(j)$$

when $\beta_b(c) \geq 1$, and $\hat{\alpha}_{1,b}(c) = 0$ when $\beta_b(c) = 0$. Similarly, let $\Lambda_b(c)$ be the number of $(W_b(k) \wedge \tau(k), I_b(k))$ generated in c units of computer time, so that $\Lambda_b(c) = \sup\{l \geq 0 : \sum_{k=1}^l [W_b(k) \wedge \tau(k)] \leq c\}$. The estimator (that is analogous to $\alpha_2(n)$) available after c units of computer time is

$$\hat{\alpha}_{2,b}(c) = \frac{\sum_{k=1}^{\Lambda_b(c)} W_b(k) \wedge \tau(k)}{\sum_{k=1}^{\Lambda_b(c)} I_b(k)}$$

if $\Lambda_b(c) \geq 1$, and $\hat{\alpha}_{2,b}(c) = 0$ if $\Lambda_b(c) = 0$. Note that it takes on average roughly $\mathbb{E}[\tau]/p_b$ units of computer time to observe one visit to A_b . So, to hope for consistency and CLTs, we need a computational budget t_b for which $t_b p_b \rightarrow \infty$ as $b \rightarrow \infty$.

We next consider the relative accuracy of the two estimators $\hat{\alpha}_{1,b}(\cdot)$ and $\hat{\alpha}_{2,b}(\cdot)$, showing that $\hat{\alpha}_{1,b}(c)$ and $\hat{\alpha}_{2,b}(c)$ are asymptotically identical in the regime in which $t_b \gg 1/p_b$.

Proposition 5 Assume $\mathbb{E}[\tau^3] < \infty$. If $t_b p_b \rightarrow \infty$ as $b \rightarrow \infty$, then we have that as $b \rightarrow \infty$,

$$\sqrt{t_b p_b} \left(\frac{\hat{\alpha}_{i,b}(t_b)}{\mathbb{E}[T_b]} - 1 \right) \Rightarrow \sqrt{\mathbb{E}[\tau]} \mathcal{N}(0, 1), \quad i = 1, 2, \quad \text{and} \quad (4)$$

$$\sqrt{t_b p_b} \left(\frac{\hat{\alpha}_{1,b}(t_b)}{\mathbb{E}[T_b]} - \frac{\hat{\alpha}_{2,b}(t_b)}{\mathbb{E}[T_b]} \right) \Rightarrow 0. \quad (5)$$

Proof. To prove (4) for $i = 1$, we will first establish that

$$\left(\sqrt{t_b p_b} \frac{1}{[t_b p_b s]} \sum_{j=1}^{[t_b p_b s]} \left(\frac{T_b(j)}{\mathbb{E}[T_b]} - 1 \right) : 0 < s < \infty \right) \Rightarrow \left(\frac{B(s)}{s} : 0 < s < \infty \right) \quad (6)$$

as $b \rightarrow \infty$ in $D(0, \infty)$, where $B = (B(s) : 0 < s < \infty)$ is a standard Brownian motion and $D(0, \infty)$ is the space of right-continuous functions with left limits on $(0, \infty)$, and then employ a random-time-change argument. We now show (6) by applying Theorem 1.4(b), p. 339, of [4] after verifying the two sufficient conditions in their (1.19) and (1.17). Note that the process $\mathbf{A}_b(\cdot)$ of their theorem is just $\mathbf{A}_b(s) = ([t_b p_b s]/(t_b p_b)) \text{Var}[T_b/\mathbb{E}[T_b]]$. But our Proposition 2 implies that $\text{Var}[T_b] = (\mathbb{E}[T_b])^2(1 + o(1))$ as $b \rightarrow \infty$ because $\mathbb{E}[T_b \mathcal{I}(T_b < \tau)] = o(1)$ as $b \rightarrow \infty$, where we use the notation that a function $f(b) = o(g(b))$ for another function g if $f(b)/g(b) \rightarrow 0$ as $b \rightarrow \infty$. Consequently, for each $s \geq 0$, we have that $\mathbf{A}_b(s) \rightarrow s$ as $b \rightarrow \infty$, establishing the condition (1.19) of Theorem 1.4(b) of [4].

The other condition, (1.17), in Theorem 1.4(b) of [4] requires verifying that

$$\frac{1}{t_b p_b} E \left[\max_{1 \leq j \leq [t_b p_b s]} V_b(j) \right] \rightarrow 0 \quad (7)$$

as $b \rightarrow \infty$, for each fixed $s \geq 0$, where the $V_b(j)$, $j \geq 1$, are IID copies of $((T_b/\mathbb{E}[T_b]) - 1)^2$. But

$$\frac{1}{t_b p_b} E \left[\max_{1 \leq j \leq [t_b p_b s]} V_b(j) \right] = \frac{1}{t_b p_b} \int_0^\infty P \left(\max_{1 \leq j \leq [t_b p_b s]} V_b(j) > x \right) dx \quad (8)$$

and $\mathbb{P}(\max_{1 \leq j \leq [t_b p_b s]} V_b(j) > x)/(t_b p_b) \rightarrow 0$ as $b \rightarrow \infty$ for each $x \geq 0$. Hence, (7) follows from (8) if we can apply the dominated convergence theorem to (8). Note that

$$\begin{aligned} \frac{1}{t_b p_b} P \left(\max_{1 \leq j \leq [t_b p_b s]} V_b(j) > x \right) &\leq \frac{1}{t_b p_b} \sum_{j=1}^{[t_b p_b s]} \mathbb{P}(V_b(j) > x) = \frac{[t_b p_b s]}{t_b p_b} P \left(\left(\frac{T_b}{\mathbb{E}[T_b]} - 1 \right)^2 > x \right) \\ &\leq \frac{s}{x^{3/2}} E \left[\left(\frac{T_b}{\mathbb{E}[T_b]} - 1 \right)^3 \right], \end{aligned} \quad (9)$$

where the last step follows from Markov's inequality. Because $\mathbb{E}[\tau^3] < \infty$ by assumption, we can apply an argument similar to that used to prove Proposition 2 to compute $\mathbb{E}[T_b^3]$. This computation shows that $\mathbb{E}[(T_b/\mathbb{E}[T_b] - 1)^3]$ is bounded as a function of b . The inequality (9) then proves that the integrand of (8) is uniformly dominated by an integrable function of x , so that (7) holds. This proves (6).

To enable applying a random-time-change theorem to (6), we next prove that

$$\frac{\beta_b(t_b s)}{t_b p_b} \Rightarrow \frac{s}{\mathbb{E}[\tau]} \quad (10)$$

as $b \rightarrow \infty$. We note that for $\varepsilon > 0$,

$$P\left(\beta_b(t_b s) \geq \frac{t_b p_b s(1 + \varepsilon)}{\mathbb{E}[\tau]}\right) = \mathbb{P}(T_b(1) + \dots + T_b(n_b) \leq t_b s), \quad (11)$$

where $n_b = \lceil t_b p_b s(1 + \varepsilon)/\mathbb{E}[\tau] \rceil$, and observe that $n_b \mathbb{E}[T_b] = t_b s(1 + \varepsilon)(1 + o(1))$ as $b \rightarrow \infty$. Chebyshev's inequality then implies that (11) is bounded above by

$$P(|T_b(1) + \dots + T_b(n_b) - n_b \mathbb{E}[T_b]| > t_b s \varepsilon (1 + o(1))) \leq \frac{n_b \text{Var}[p_b T_b]}{(t_b p_b s \varepsilon)^2 (1 + o(1))} \rightarrow 0$$

as $b \rightarrow \infty$, because $\text{Var}[p_b T_b]$ is bounded as a function of b (since $\mathbb{E}[(T_b/\mathbb{E}[T_b])^2] \rightarrow 1$ as $b \rightarrow \infty$). A similar argument shows that $\mathbb{P}(\beta_b(t_b s) \leq t_b p_b s(1 - \varepsilon)/\mathbb{E}[\tau]) \rightarrow 0$ as $b \rightarrow \infty$, for each fixed $\varepsilon > 0$, thus proving (10).

In view of (6) and (10), the random-time-change theorem (e.g., Theorem 14.4 of [2]) then yields the conclusion that

$$\sqrt{t_b p_b} \frac{1}{\beta_b(t_b s)} \sum_{j=1}^{\beta_b(t_b s)} \left(\frac{T_b(j)}{\mathbb{E}[T_b]} - 1 \right) \Rightarrow \frac{B(s/\mathbb{E}[\tau])}{s/\mathbb{E}[\tau]}$$

as $b \rightarrow \infty$. Consequently, (4) for $i = 1$ follows by setting $s = 1$.

We next show that (5) holds, which will imply (4) holds for $i = 2$ by (4) for $i = 1$ and the converging-together lemma. First note that $\hat{\alpha}_{1,b}(c) \leq \hat{\alpha}_{2,b}(c) \leq \sum_{j=1}^{\beta_b(c)+1} T_b(j)/\beta_b(c)$ and $\sum_{j=1}^{\beta_b(c)+1} T_b(j)/\beta_b(c) - \hat{\alpha}_{1,b}(c) = T_b(\beta_b(c) + 1)/\beta_b(c)$, which implies that (5) holds if

$$\frac{1}{\sqrt{t_b p_b}} \max_{1 \leq j \leq \lceil t_b p_b s \rceil} \frac{T_b(j)}{\mathbb{E}[T_b]} \Rightarrow 0 \quad (12)$$

as $b \rightarrow \infty$ for $s \geq 0$. But this follows immediately from (7) and Markov's inequality. \square

3.2 Equivalence in a Highly Reliable Markovian Systems Setting

We next consider a model of highly reliable Markovian systems (HRMS) commonly studied in the literature, e.g., see [3, 15, 16, 17], among others. Basically (the reader is advised to read the above references for more details) the state space S is decomposed into the set A of failed states and the set of operational states. Transitions of the Markov chains are repairs and failures of components. Failures are assumed to be rare events with respect to repairs, so that a rarity parameter $0 < \varepsilon \ll 1$ is introduced. Failure transitions are assumed to have a rate $O(\varepsilon)$, while repair transitions have a rate $\Theta(1)$, where we use the notation that a function $f(\varepsilon)$ is $O(g(\varepsilon))$ if $|f(\varepsilon)/g(\varepsilon)|$ remains bounded when $\varepsilon \rightarrow 0$ and it is $\Theta(g(\varepsilon))$ if $|f(\varepsilon)/g(\varepsilon)|$ is bounded and also bounded away from 0, when $\varepsilon \rightarrow 0$. The smaller ε is, the smaller the probability to reach A from an initial operational state. In contrast to Section 3.1, where b was the rarity parameter and we considered a sequence of sets $(A_b : b \geq 1)$ as $b \rightarrow \infty$, we now change the rarity parameter to ε and examine the asymptotics as $\varepsilon \rightarrow 0$ for a fixed set A of failed states.

Index by ε the probability measure driving the system and denote it \mathbb{P}_ε . Indexing by ε will allow us to highlight the properties of estimators when ε gets close to 0. The time to failure (that is, to reach a failed state) is denoted T_ε , and the direct and ratio estimators can be used for both crude simulation and importance sampling. Here we consider the embedded discrete-time Markov chain (DTMC), where the time spent in each state is taken as the expected value of the exponential holding time in the state of the continuous-time chain. This discrete-time conversion reduces the variance of the estimator and simplifies the analysis in next section. But looking at either the continuous or discrete version does not change the conclusions of the comparison between the direct and the ratio-based estimators.

As was done in Section 3.1, we can also examine the asymptotic equivalence of the two estimators as $\varepsilon \rightarrow 0$. Instead, we simply present numerical results to illustrate this property.

A numerical comparison. Consider a system with 3 component types, with $n_1 = n_2 = n_3 = 3$, where n_i is the redundancy of component type i . Each component has an exponentially distributed time to failure with rate λ_i for components of type i , where $\lambda_i = \varepsilon$, for some parameter ε . Any failed component has an exponentially distributed repair time with rate 1. Times to failure and repair times are all independent. The system is down whenever fewer than two components of any one type are operational.

The results with various values of ε and various sample sizes m are provided in Table 1 for the direct estimator and Table 2 for the regenerative (ratio-based) estimator (n is then the number of independent cycles). The last column displays the work-normalized variance, defined as the variance of the estimator multiplied

by the CPU time. It balances the computational effort and variance and basically represents the expected variance for a unit of computational budget. It can be seen that the work-normalized variances are basically the same for each value of ϵ .

Table 1: Results for the direct DTMC crude estimator

m	ϵ	Est.	Confidence Interval	Variance	CPU	Work Norm. Var.
10^5	0.1	8.755	(8.708e+00, 8.802e+00)	5.840e+01	0.17	9.733e-05
10^7	0.1	8.769	(8.764e+00, 8.774e+00)	5.879e+01	17.7	1.041e-04
10^5	0.01	5.818e+02	(5.782e+02, 5.854e+02)	3.366e+05	1.33	4.488e+00
10^7	0.01	5.841+02	(5.838e+02, 5.845e+02)	3.343e+05	134	4.482e+00
10^5	0.001	5.5925825e+04	(5.558e+04, 5.627e+04)	3.126e+09	12.86	4.022e+05
10^7	0.001	5.5844640e+04	(5.581e+04, 5.588e+04)	3.117e+09	1316.5	4.104e+05

Table 2: Results for the regenerative crude DTMC estimator

n	ϵ	Est.	Confidence Interval	Variance	CPU	Work Norm. Var.
10^5	0.1	8.692	(8.595e+00, 8.788e+00)	2.412e+02	0.05	1.206e-04
10^7	0.1	8.772	(8.762e+00, 8.782e+00)	2.484e+02	4.283	1.064e-04
10^5	0.01	5.805e+02	(5.558e+02, 6.051e+02)	1.580e+07	0.0166	2.633e+00
10^7	0.01	5.812e+02	(5.788e+02, 5.837e+02)	1.586e+07	2.917	4.627e+00
10^5	0.001	5.496e+04	(4.742e+04, 6.249e+04)	1.478e+12	0.0166	2.463e+05
10^7	0.001	5.535e+04	(5.459e+04, 5.611e+04)	1.510e+12	2.800	4.227e+05

It is interesting to note that the direct estimator does not encounter rare-event problems: as $\epsilon \rightarrow 0$, the relative variance (i.e., variance divided by the square of the expected value) is kept bounded in Table 1. But the computation (CPU) time increases because the number of steps before reaching system failure increases. Table 2 instead shows the opposite for the ratio-based estimator: the computational time is bounded but we have a rare-event estimation issue. In other words, the computational time issue is replaced by a rare-event estimation problem. Compared to computational time issues, rare-event problems are more extensively studied and variance-reduction techniques can be applied.

4 IMPORTANCE SAMPLING ESTIMATORS

Because the crude estimators are asymptotically equivalent, one may wonder why in the rare-event setting IS techniques have only been developed for the ratio estimator [8, 14] but not for its direct counterpart. To illustrate our arguments and give counter-examples, we introduce the following very simple example which will be used throughout the section.

Example 1 Consider a system made of a single type of components with two components failing with rate ε and a single repairman with repair rate 1. The state space and transition-rate diagram of the continuous-time Markov chain are represented in Figure 1, where state x means x failed components. The system is failed when both components are failed (state 2). For this example, the embedded




Figure 1: HRMS with a single type of components and 2 components.

DTMC is described in Figure 2, with mean sojourn times $1/(2\varepsilon)$ and $1/(1 + \varepsilon)$ in states 0 and 1, respectively. A path to failure is simply described as a transition




Figure 2: HRMS with a single type of components and 2 components: embedded DTMC.

from 0 to 1, followed by $n \geq 0$ cycles $1 \rightarrow 0 \rightarrow 1$, and finally a transition from 1 to 2. Hence, the MTTF can be directly computed as

$$\mathbb{E}_\varepsilon(T_\varepsilon) = \sum_{n=0}^{\infty} (n+1) \left(\frac{1}{2\varepsilon} + \frac{1}{1+\varepsilon} \right) \left(\frac{1}{1+\varepsilon} \right)^n \frac{\varepsilon}{1+\varepsilon} = \frac{1+3\varepsilon}{2\varepsilon^2}.$$

Note that this expression could also have been found from the system of equations obtained by conditioning on the first move: $\mathbb{E}_\varepsilon(T_\varepsilon) = 1/(2\varepsilon) + \mathbb{E}_{\varepsilon,1}(T_\varepsilon)$ and $\mathbb{E}_{\varepsilon,1}(T_\varepsilon) = 1/(1 + \varepsilon) + 1/(1 + \varepsilon)\mathbb{E}_\varepsilon(T_\varepsilon)$ (with $\mathbb{E}_{\varepsilon,1}(\cdot)$ the expected value when starting from state 1 instead of state 0), but it is insightful to decompose into the sum over paths with n cycles, whose expected length is $(n+1) \left(\frac{1}{2\varepsilon} + \frac{1}{1+\varepsilon} \right)$ and whose probability is $\left(\frac{1}{1+\varepsilon} \right)^n \frac{\varepsilon}{1+\varepsilon}$.

For the crude Monte Carlo estimator using the embedded DTMC, the second

moment is

$$\mathbb{E}_\varepsilon[(T_\varepsilon)^2] = \sum_{n=0}^{\infty} (n+1)^2 \left(\frac{1}{2\varepsilon} + \frac{1}{1+\varepsilon} \right)^2 \left(\frac{1}{1+\varepsilon} \right)^n \frac{\varepsilon}{1+\varepsilon} = \frac{(2+\varepsilon)(1+3\varepsilon)^2}{4(1+\varepsilon)\varepsilon^4}$$

leading to a variance of $\frac{1}{4} \frac{(1+3\varepsilon)^2}{(1+\varepsilon)\varepsilon^4}$. The relative variance is bounded, illustrating again that this estimator does not encounter a rare-event problem.

Let N be the (random) number of transitions in a run (to which the computation time is proportional), and its expectation is

$$\mathbb{E}_\varepsilon(N) = \sum_{n=0}^{\infty} (2+2n) \left(\frac{1}{1+\varepsilon} \right)^n \frac{\varepsilon}{1+\varepsilon} = \frac{2(1+\varepsilon)}{\varepsilon}.$$

The work-normalized relative variance $WNRV$, defined as the variance multiplied by the computation time and divided by the square of the expected value, is then

$$WNRV = \frac{2}{\varepsilon}.$$

The work-normalized relative variance is a good measure of efficiency of an estimator, since if bounded as $\varepsilon \rightarrow 0$, it means that the computational budget to ensure a predefined accuracy level is independent of the rarity. Here the computational time to achieve a given accuracy level increases as $\varepsilon \rightarrow 0$, as we had previously noted.

4.1 Failure Biasing

If we wish to apply IS to the direct estimator, a natural idea is to apply failure biasing similarly to what has been done for the regenerative estimator: from any state except the initial one with all components up, change the probability of making a failure transition to be ρ , independent of ε . This will make reaching a failed state more likely because under the original system dynamics, the probability of taking a failure transition is $O(\varepsilon)$ as failure rates are $O(\varepsilon)$ but repair rates are $\Theta(1)$. Several implementations exist, including simple failure biasing (SFB) and balanced failure biasing (BFB) [16]. Under SFB, the probability of any failure given that a failure occur is proportional to its original probability, but with BFB, it is uniform; i.e., if f failure transitions are possible out of a state, then each failure transition is given probability $1/f$. For repairs, the conditional probabilities are always taken proportional to the original ones.

Note that when simulating a Markov chain, the number of steps can be large and a change of probability matrix can lead to very poor results if the probabilities are changed too much because the likelihood ratio is subject to large variations. But

usually for HRMS the number of steps in direct paths to failure is small (otherwise the failure biasing methods would not be efficient), so we could (wrongly as we will see) expect failure biasing to be efficient.

Example 2 When applying IS to Example 1 by changing the probabilities of the Markov chain, the only latitude we have is to change to probabilities from state 1 as described in Figure 3 by using ρ as the probability to reach the failed state from 1. Failure biasing is then the only possibility.




Figure 3: HRMS with a single type of components and 2 components: Failure biasing on the embedded DTMC.

The probability under IS of the path to failure with $n \geq 0$ cycles is $(1 - \rho)^n \rho$. Letting $\tilde{\mathbb{E}}_\varepsilon$ denote the expectation operator under the IS probability distribution, the second moment of the IS estimator is

$$\tilde{\mathbb{E}}_\varepsilon[(T_\varepsilon L)^2] = \mathbb{E}_\varepsilon[(T_\varepsilon)^2 L] = \sum_{n=0}^{\infty} (n+1)^2 \left(\frac{1}{2\varepsilon} + \frac{1}{1+\varepsilon} \right)^2 \frac{\left(\left(\frac{1}{1+\varepsilon} \right)^n \frac{\varepsilon}{1+\varepsilon} \right)^2}{(1-\rho)^n \rho}.$$

For the sum to converge, we need $1/((1+\varepsilon)^2(1-\rho)) < 1$, or equivalently,

$$\rho < 1 - \frac{1}{(1+\varepsilon)^2} = 2\varepsilon - 3\varepsilon^2 + o(\varepsilon^2). \quad (13)$$

In other words, the failure probability from state 1 cannot be increased too much; otherwise, the likelihood ratio will build up too much and lead to infinite variance.

This is not a new story as this type of issue has already been encountered in the simulation of HRMS by [9] and [10], where deferred repairs are considered so that at some states, failures are not rare (something we also have here from state 0). The point is that failures from intermediate states can have a probability $\Theta(1)$, introducing cycles with large probability, which can cause the variance of the likelihood to potentially explode. As a remedy, they propose to apply a so-called small failure biasing, assigning a failure probability $\delta \ll \rho$ to the whole set of failures from some states. While the issue is the same, here the assumptions are different:

1. We do not have deferred repairs and the problem is only at the initial state 0, which is not a problem with regenerative simulation because the simulation of a cycle stops when we have a return to this state;
2. In [9] and [10], the small probability δ is independent of ε , which is not possible here from what we have seen to avoid an infinite variance.

Thus the implementation of IS to the direct estimator is problematic since, already on our simple example, it is not possible to make unrare the failure transition from state 1.

The average number of transitions under IS is

$$\tilde{\mathbb{E}}_{\varepsilon}(N) = \sum_{n=0}^{\infty} (2+2n)(1-\rho)^n \rho = \frac{2}{\rho}. \quad (14)$$

So the average simulation time for a single run will increase to infinity as $\varepsilon \rightarrow 0$ by (13).

Remark: This analysis can be related to [6], where it is shown that we can typically expect the IS variance to grow exponentially in the (deterministic) number of simulated steps. With HRMS, the asymptotic analysis is instead with respect to $\varepsilon \rightarrow 0$ rather than an increasing (fixed) number of transitions. Although failure biasing with $\rho = \Theta(1)$ leads to the system failing after usually only a few transitions, (13) stipulates that $\rho = O(\varepsilon)$ to ensure finite variance, resulting in the mean of the (random) number of transitions until failure to diverge asymptotically to infinity, as seen by (14).

But even if the computation time increases as $\varepsilon \rightarrow 0$, the direct estimator could still be efficient if the relative variance vanishes with ε . To see if that can happen, let us focus on the zero-variance IS scheme.

4.2 Zero-Variance Approximation

In the literature, there indeed exists a zero-variance IS scheme for Markov chains and HRMS, as derived in [1] and [12]. In order to implement it, we cannot use the estimator $T_{\varepsilon}L$ but rather the still-unbiased $T_{\varepsilon}^{IS} = \sum_{j=0}^{\tau_{\varepsilon}-1} (1/\lambda(Y_k))L_k$, where L_k is the likelihood ratio for step 0 to step k . This type of estimator is often called a *filtered importance sampling estimator* [1, 5]. For any states x and y , let $\mathbb{E}_{\varepsilon,y}(T_{\varepsilon})$ be the MTTF starting from y and $(P_{x,y})_{x,y \in \mathcal{S}}$ the original transition matrix of the DTMC. Then using the IS matrix

$$\tilde{P}_{x,y} = P_{x,y} \frac{1/\lambda(x) + \mathbb{E}_{\varepsilon,y}(T_{\varepsilon})}{\mathbb{E}_{\varepsilon,x}(T_{\varepsilon})}$$

yields an estimator with variance zero, which follows from a direct application of the framework described in [1, 11, 12].

Example 3 On our example in Figure 3, we can modify the probabilities from state 1 only, with

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

We can see that the probability to reach State 2 directly from 1 is $\Theta(\varepsilon^3)$, so even if the variance is zero, the estimation takes on average longer time, $\frac{2}{\rho} = \Theta(\varepsilon^{-3})$, as ε gets closer to zero. If the optimal ρ is not known and an approximation of order $\rho = \theta(\varepsilon^3)$ is used, we need a relative variance $O(\varepsilon^{-3})$, that is a variance $O(\varepsilon)$, to ensure a bounded work-normalized variance.

Let us investigate if it is easily attainable. The second moment of the estimator T_ε^{IS} with ρ as the probability to go directly from 1 to 2 is

$$\tilde{\mathbb{E}}_\varepsilon[(T_\varepsilon^{IS})^2] = \sum_{n=0}^{\infty} \left(\frac{1}{2\varepsilon} + \sum_{k=1}^n \left(\frac{1}{1+\varepsilon} + \frac{1}{2\varepsilon} \right) \left(\frac{1}{1+\varepsilon} \right)^k \frac{1}{(1-\rho)^k} + \frac{1}{1+\varepsilon} \frac{(\varepsilon/(1+\varepsilon)(1/(1+\varepsilon))^n)}{\rho(1-\rho)^n} \right)^2 (1-\rho)^n \rho.$$

A closed-form expression is easily obtainable, but the derivation is very long and not insightful. However, we make the following observations:

- For $\rho = \frac{2\varepsilon^3}{(1+\varepsilon)^2(1+2\varepsilon)}$, we retrieve $\tilde{\mathbb{E}}_\varepsilon[(T_\varepsilon^{IS})^2] = \frac{1}{4} \frac{(1+3\varepsilon)^2}{\varepsilon^4} = (\tilde{\mathbb{E}}_\varepsilon[T_\varepsilon^{IS}])^2$, that is, a variance zero.
- For $\rho = \varepsilon^3$ (i.e., an approximation of the probability of good asymptotic order), the variance is $\Theta(\varepsilon^{-2})$, two orders of magnitude better than the variance of the crude estimator, but this gain is lost on the computational time.
- For $\rho = 2\varepsilon^3$ (i.e., the exact first-order term of the zero-variance change of measure), the variance is $\Theta(1)$, which is better but still not sufficient to yield a bounded work-normalized variance.

Thus, we see that much better than an exact first-order approximation of transition probabilities is required. This seems hard to obtain in practice. With the ratio estimator, a first-order approximation is used in [12], which yields bounded normalized variance (even a vanishing one), and the estimator does not suffer from an increasing computational time.

5 CONCLUSIONS

In conclusion, we have reviewed and compared two standard estimators of the MTTF for (classically) regenerative processes: a direct one expressed as the average of simulated times to failure, and a second one making use of the regenerative

structure and expressing the MTTF as a ratio of expected values. We have highlighted that

1. Crude direct and ratio-based estimators are asymptotically equivalent as the probability to reach the specified failure set decreases, with the computational issue for the direct estimator just being replaced by a rare-event problem for the ratio estimator.
2. When failures are rare, we may want to apply IS to obtain more efficient estimators, but we have illustrated that for the direct estimator, designing an efficient IS procedure might be difficult, while many efficient ones exist for the ratio expression. The latter is then advised.

ACKNOWLEDGMENTS

This work has been supported in part by the National Science Foundation under Grant No. CMMI-1537322. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

- [1] H. P. Awad, P. W. Glynn, and R. Y. Rubinstein. Zero-variance importance sampling estimators for markov process expectations. *Mathematics of Operations Research*, 38(2):358–388, 2013.
- [2] P. Billingsley. *Convergence of Probability Measures*, volume Second. John Wiley and Sons, New York, 1999.
- [3] H. Cancela, G. Rubino, and B. Tuffin. MTTF estimation by Monte Carlo methods using Markov models. *Monte Carlo Methods and Applications*, 8(4):312–341, 2002.
- [4] S.N. Ethier and T.G. Kurtz. *Markov Processes: Characterization and Convergence*. Springer-Verlag, New York, 1986.
- [5] P. Glasserman. Filtered Monte Carlo. *Mathematics of Operations Research*, 18(3):610–634, 1993.
- [6] P. W. Glynn. Importance sampling for Markov chains: Asymptotics for the variance. *Stochastic Models*, 10(4):701–717, 1994.

- [7] A. Goyal, P. Shahabuddin, P. Heidelberger, V. F. Nicola, and P. W. Glynn. A unified framework for simulating Markovian models of highly reliable systems. *IEEE Transactions on Computers*, C-41:36–51, 1992.
- [8] P. Heidelberger. Fast simulation of rare events in queueing and reliability models. *ACM Transactions on Modeling and Computer Simulation*, 5(1):43–85, 1995.
- [9] S. Juneja and P. Shahabuddin. Fast simulation of Markovian reliability/availability models with general repair policies. In *Proceedings of the Twenty-Second International Symposium on Fault-Tolerant Computing*, pages 150–159. IEEE Computer Society Press, 1992.
- [10] S. Juneja and P. Shahabuddin. Fast simulation of Markov chains with small transition probabilities. *Management Science*, 47(4):547–562, 2001.
- [11] P. L’Ecuyer and B. Tuffin. Effective approximation of zero-variance simulation in a reliability setting. In *Proceedings of the 2007 European Simulation and Modeling Conference*, pages 48–54, Ghent, Belgium, 2007. EUROSIS.
- [12] P. L’Ecuyer and B. Tuffin. Approximating zero-variance importance sampling in a reliability setting. *Annals of Operations Research*, 189(1):277–297, 2012.
- [13] G. Rubino and B. Tuffin. Markovian models for dependability analysis. In G. Rubino and B. Tuffin, editors, *Rare Event Simulation using Monte Carlo Methods*, pages 125–144. John Wiley & Sons, 2009. Chapter 6.
- [14] G. Rubino and B. Tuffin. *Rare Event Simulation using Monte Carlo Methods*. Wiley, 2009.
- [15] P. Shahabuddin. Fast transient simulation of Markovian models of highly dependable systems. *Performance Evaluation*, 20:267–286, 1994.
- [16] P. Shahabuddin. Importance sampling for the simulation of highly reliable Markovian systems. *Management Science*, 40(3):333–352, 1994.
- [17] P. Shahabuddin, V. F. Nicola, P. Heidelberger, A. Goyal, and P. W. Glynn. Variance reduction in mean time to failure simulations. In *Proceedings of the 1988 Winter Simulation Conference*, pages 491–499. IEEE Press, 1988.
- [18] W. L. Smith. Regenerative stochastic processes. *Proceedings of the Royal Society: Series A. Mathematical and Physical Sciences*, 232:6–31, 1955.