Path-ZVA Simulation Method for Time-Bounded Rare Events in a Semi-Markov Chain

Daniël Reijsbergen, Pieter-Tjerk de Boer,
and Werner Scheinhardt

SINGAPORE UNIVERSITY OF TECHNOLOGY AND DESIGN

UNIVERSITY OF TWENTE.

Table of Contents

Reliability Modelling

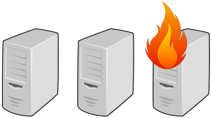

Reliability models: how reliable is my system?

- *Unavailability*: long-term average fraction of time that system is failed.
- *Unreliability*: probability that system fails within τ time units, $\tau \in (0, \infty)$.

The *failures* and *repairs* of system components can be modelled using a *semi-Markov Chain*.

Given description of system as semi-Markov chain, unavailability and unreliability can be estimated using *simulation*.

Prior work:[1,2] unavailability, so we focus on ***unreliability***.

[1] R., de Boer, Scheinhardt, and Juneja. "*Path-ZVA: General, efficient, and automated importance sampling for highly reliable Markovian systems.*" ACM TOMACS, 2018.

[2] Ruijters, R., de Boer, and Stoelinga. "*Rare event simulation for dynamic fault trees.*" Reliability Engineering & System Safety, 2019.

Unreliability

Let $\mathcal{X}$ be the countable state space of the system model, i.e., which, or how many, components are working or not.

Some states represent system failure.

- w.l.o.g.: single "goal" state g.

Unreliability: probability π that, starting from $s \in \mathcal{X}$, we reach state $g \in \mathcal{X}$ before time $\tau > 0$.

Toy example: single component with 4 spares. System fails when component and all spares have failed.

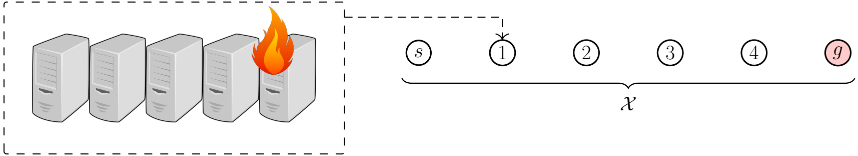

Reliability Modelling: Toy Example

Expected time before failure of active component: λ^{-1}. Here, λ need not be rate of exponential distribution!

Group repair: repair process is reset whenever a new component fails, but when it is completed all components are repaired. ($\Rightarrow$ semi-Markov!)

As with failures, the probability distribution of repair times need not be exponential – e.g., Weibull or approximately normal.

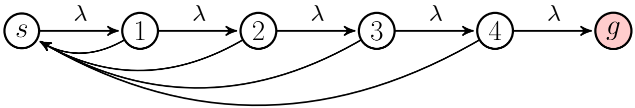

Semi-Markov Chains

Semi-Markov chain: model the transitions between states in $\mathcal{X}$.[3]

Two stochastic processes:

1. $X_n \in \mathcal{X}$, the system state after the nth state transition,
2. $T_n \in \mathbb{R}^+$, the time spent in state X_{n-1}.

The process X_n is a *discrete-time Markov chain (DTMC)*, i.e., only depends on the current state and not on previous states.

If T_n are drawn from the exponential distribution, then (X_n, T_n) is a continuous-time Markov chain.

[3] Limnios and Oprisan. *Semi-Markov processes and reliability*. 2012.

Semi-Markov Chains: Definition

Two common definitions of semi-Markov process.

Our focus: let T'_{nz} be the *potential* transition times for all steps $n \in \mathbb{N}$ and $x \in \mathcal{X}$, defined as follows:

$$\mathbb{P}(T'_{nx} \leq t) = F_{X_{n-1}x}(t).$$

Then $X_n = x$ and $T_n = t$ if and only if:

$$x = \operatorname*{argmin}_{y \in \mathcal{X}} T'_{ny} \quad \text{and} \quad t = T'_{nx}.$$

Simulation: in each time step n, draw all potential transition times t'_{ny}, determine which time $t = \min_y t'_{ny}$ is smallest, then move to corresponding state at time t.

Rarity in Semi-Markov Chains

Unreliability is probability π of reaching g before time bound τ.

If π is small, e.g., because component repairs occur more rapidly than component failures, then π is hard to estimate.

Formal notion of *rarity*: rarity parameters ϵ and r_{xy}.
If $r_{xy} > 0$, then if $\epsilon \downarrow 0$ it becomes unlikely for the potential transition time from x to y to be small.

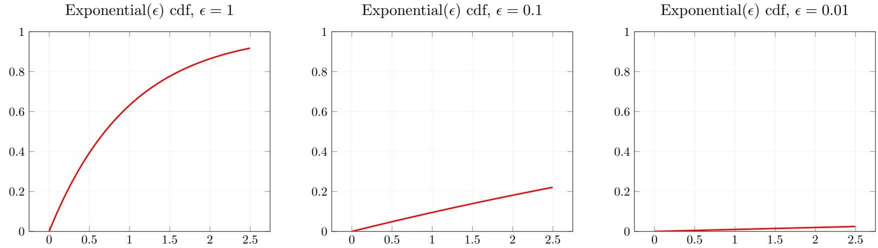

Rarity in Semi-Markov Chains

For each pair of states $x, y \in \mathcal{X}$ we define $r_{xy} \geq 0$ such that

$$F_{xy}(t) = \Theta(\epsilon^{r_{xy}}),$$

that is,

$$\lim_{\epsilon \downarrow 0} \frac{F_{xy}(t)}{\epsilon^{r_{xy}}} = c_t, \text{ such that } 0 < c_t < \infty \text{ for all } t \in \mathbb{R}^+$$

Typically, $r_{xy} > 0$ for failure transitions, and $r_{xy} = 0$ for repair transitions.

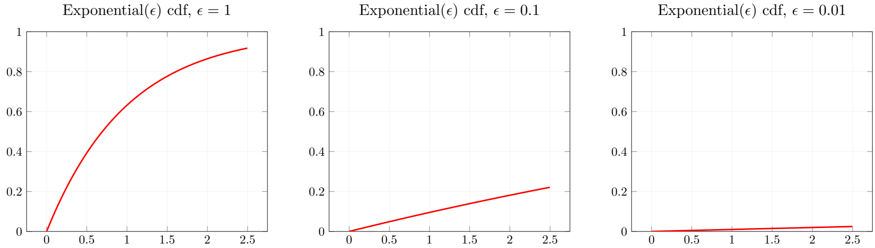

Importance Sampling

Importance sampling: in each step, instead of original combined distribution F_x for potential transition times from x, use distribution G_x that makes the rare event more likely.

To produce an unbiased estimate, we weight each *sample path* ω of size n with its *likelihood ratio* L:

$$L(\omega) = \prod_{i=1}^{n} \frac{dF_{x_{i-1}}((t'_{iy})_{y \in \mathcal{X}})}{dG_{x_{i-1}}((t'_{iy})_{y \in \mathcal{X}})}$$

Question: how to choose the new distribution G_x?

Non-trivial choice: if made poorly, variance of estimator will be higher (or even infinite).

Zero-Variance Approximation

How to choose G_x?

- *Zero-Variance Approximation (ZVA)*:[4] base the new probability of jumping to a state on the approximated likelihood of observing the rare event from that state.

Let $v(x) \in [0, 1]$ be a guess for the probability of observing the rare event, given current state $x \in \mathcal{X}$.

- *Path-ZVA*: base $v(x)$ on the shortest paths (in terms of ϵ) from x to the system failure state g.

  We determine the shortest paths numerically, i.e., using Dijkstra's method – no need to consider states that are harder to reach from s than g.

  Numerical method also removes *high-probability cycles*.

[4] L'Ecuyer and Tuffin. "*Approximate zero-variance simulation.*" Winter Simulation Conference, 2008.

Proposed Approach

Proposed approach:

$$G_x((t_y)_{y\in\mathcal{X}}) = \sum_{z\in\mathcal{X}} p_{xz} H_{xz}(t_{xz}) \prod_{\substack{y\in\mathcal{X}\\ y\neq z}} F_{xy}(t_{xy})$$

Intuition: in each step, pick a "target" state Z using v:

$$p_{xz} = \mathbb{P}(Z = z | X = x) \sim \frac{\epsilon^{r_{xz}} v(z)}{v(x)}.$$

Speed up the transition to z if this step is ϵ-hard. Baseline:

$$H_{xz}(t) = F_{xz}(\epsilon^{-r_{xz}} t).$$

However, we will also explore alternatives.

Experiments

We use simulation experiments to evaluate the performance of our estimator $\hat{\pi}$. We use 10^7 runs in all cases.

Goal: Bounded Relative Error (BRE):

$$\lim_{\epsilon \to 0} \frac{\sqrt{\mathrm{Var}(\hat{\pi})}}{\pi} < \infty$$

BRE $\Leftrightarrow$ relative c.i. width bounded.

General conditions known for BRE are known for DTMCs.[5]

[5] Nakayama. "*General conditions for bounded relative error in simulations of highly reliable Markovian systems.*" Advances in Applied Probability, 1996.

Experiments

We consider two distributions for failures/repairs:

1. Exponential(λ)
2. Weibull($2, \lambda$)

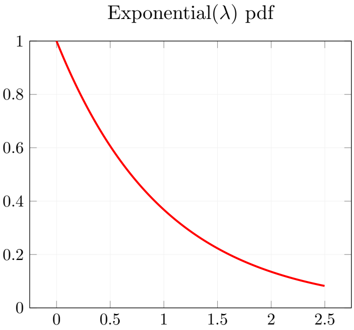

$$f(x) = \lambda e^{-\lambda x}$$

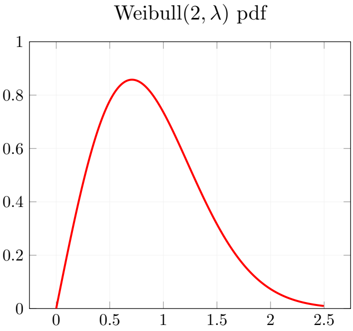

$$f(x) = 2\lambda^2 x e^{-(\lambda x)^2}$$

Experiment 1: Toy Example

We consider four settings for the toy example:

setting	failures	failures sim.	repairs	rel. c.i. half-width ($N = 10^7$)
1	Exp(ϵ)	Exp(1)	Exp(1)	≈ 0.05
2	Exp(ϵ)	Exp(1)	Weib(2,1)	≈ 0.05
3	Weib(2,ϵ)	Weib(2,1)	Weib(2,1)	≈ 0.75
4	Weib(2,ϵ)	Exp(1)	Weib(2,1)	≈ 0.25

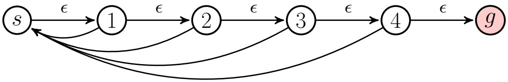

Experiment 1: Toy Example

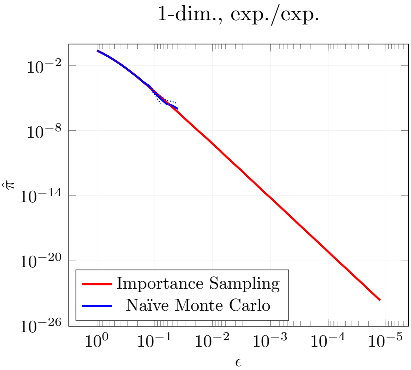

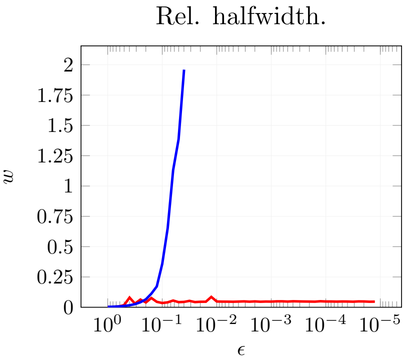

Experiment 1: Toy Example

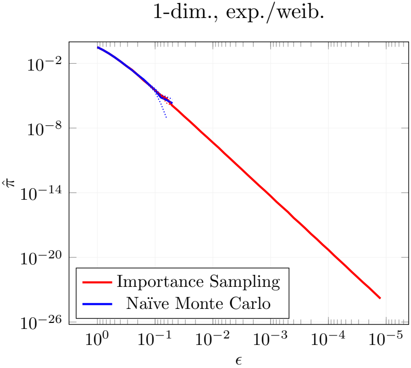

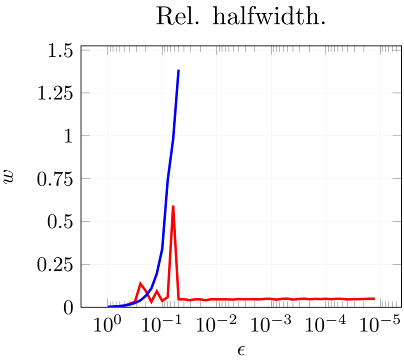

Experiment 1: Toy Example

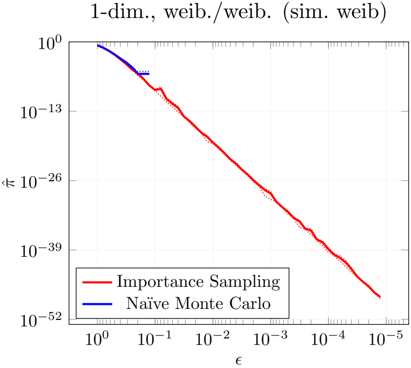

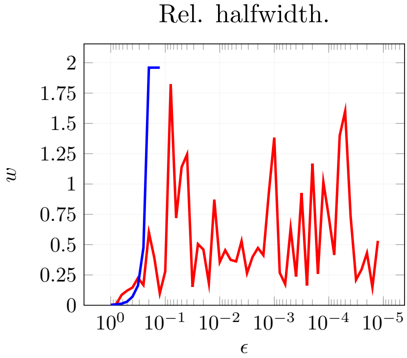

Experiment 1: Toy Example

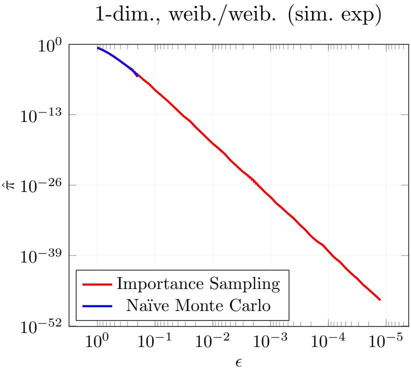

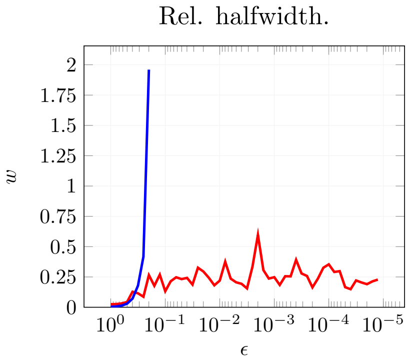

Experiment 2: 2-dimensional system

Generalization: two component types, system fails if there are 5 or more failed components in total.

We consider the three cases as before (exp./exp., exp./weib., weib./weib.). Also, a *bad* IS scheme, in which we bias *both* failure transitions in each round instead of choosing a target.

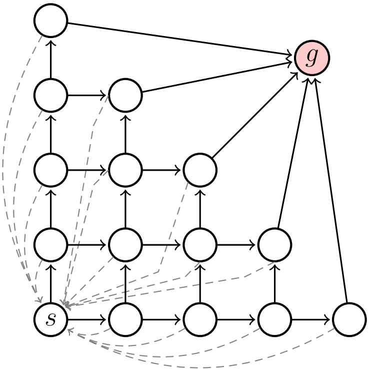

Experiment 2: 2-dimensional system

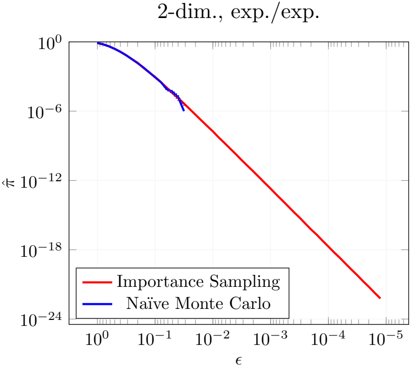

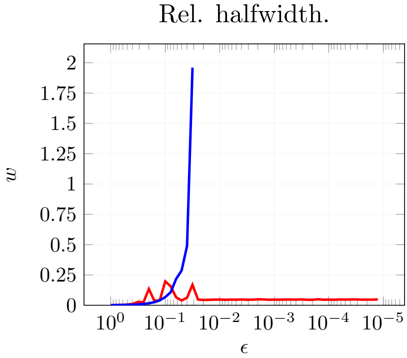

Experiment 2: 2-dimensional system

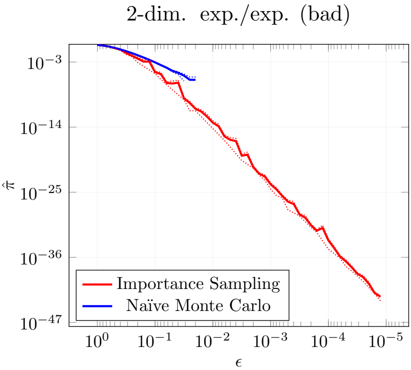

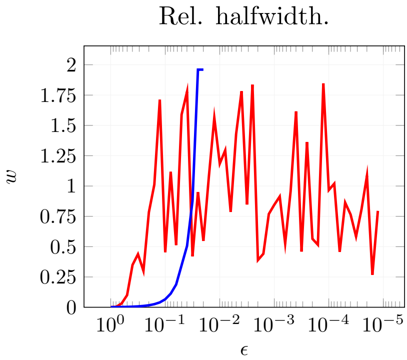

Experiment 2: 2-dimensional system

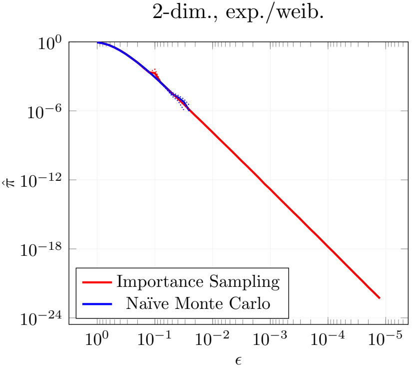

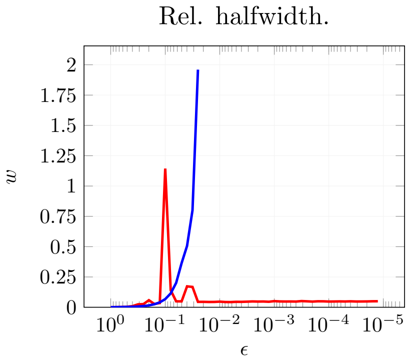

Experiment 2: 2-dimensional system

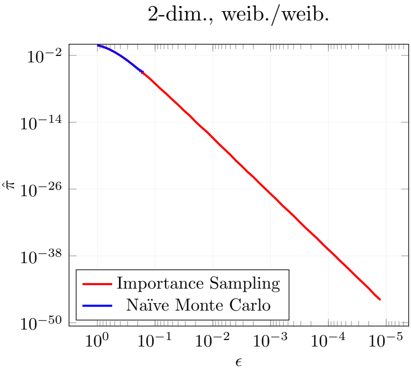

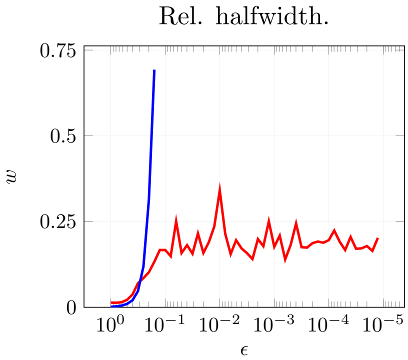

Experiment 3: Realistic System

So far, toy examples. Realistic system model:
Distributed Database System (DDS).

Nine component types, system fails if two of any type have failed. Dedicated repair per type.

State space size: thousands. Numerical solutions doable for Markov chains, not for general distributions (such as Weibull).

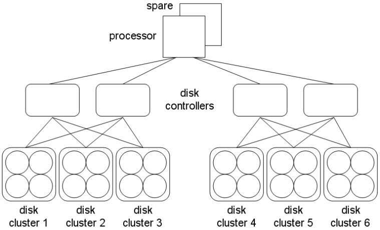

Experiment 3: Realistic System

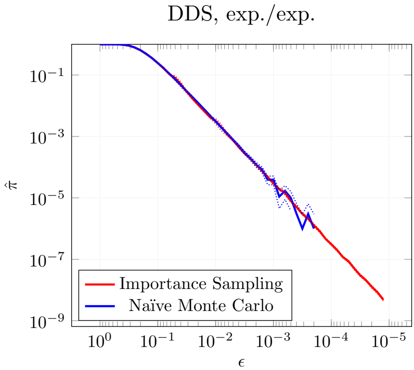

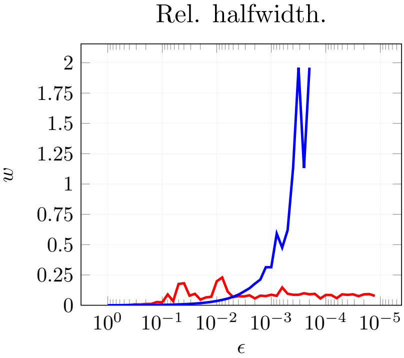

Experiment 3: Realistic System

DDS, exp./weib.

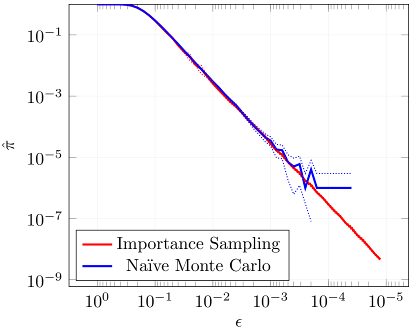

Rel. halfwidth.

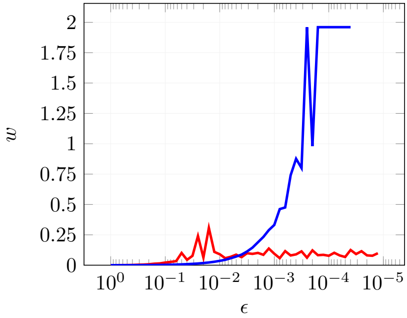

Experiment 3: Realistic System

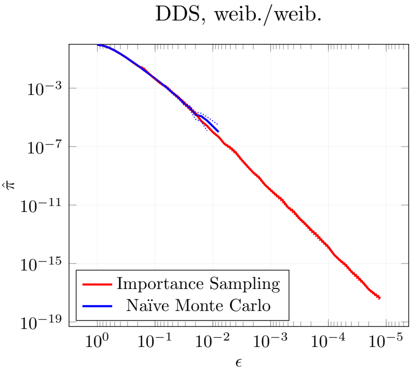

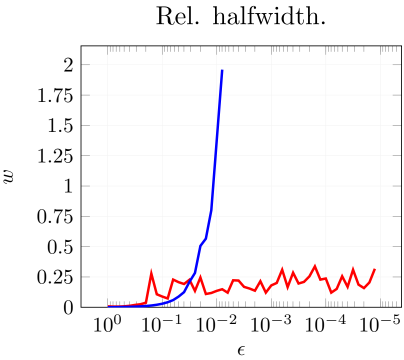

Conclusions

Experiments show that our method works well for some classes of semi-Markov processes.

Future work:

- Good performance for distributions with an exponential tail, but what about other (e.g., power tail) distributions?
- Prove bounded relative error. Can we also show vanishing relative error?
- Can our approach be extended to generalized semi-Markov processes (GSMPs)?

Conclusions

Thank you!