Importance sampling of Piecewise Deterministic Markov Processes for rare event simulation

Guillaume Chennetier [1,2] · Hassane Chraibi [1] · Anne Dutfoy [1] · Josselin Garnier [2]

[1] EDF Lab Paris-Saclay, Boulevard Gaspard Monge, 91120 Palaiseau, France.
[2] CMAP, Ecole Polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France.

Industrial motivation

Estimation of the probability of failure of industrial systems involved in the operation of nuclear power plants and dams.

- A computer code simulates the operation of the system.
  $\longrightarrow$ **P**iecewise **D**eterministic **M**arkov **P**rocesses.
- Typical probabilities of failure are very small (about 10^{-5}).
- Each simulation is numerically expensive.

$\hookrightarrow$ Crude Monte-Carlo methods are not feasible.

Piecewise Deterministic Markov Processes

Definition of a PDMP

Piecewise Deterministic Markov Process
(M.H.A Davis 1984)

Hybrid process: $Z_t = (X_t, M_t) \in E$

- position $X_t \in \mathcal{X}$ is continuous
- mode $M_t \in \mathcal{M}$ is discrete

1. **Flow** Φ $\rightarrow$ deterministic dynamics between two jumps
2. **Jump intensity** λ $\rightarrow$ law of the time of the random jumps
3. **Jump kernel** K $\rightarrow$ law of the state of the process after a jump

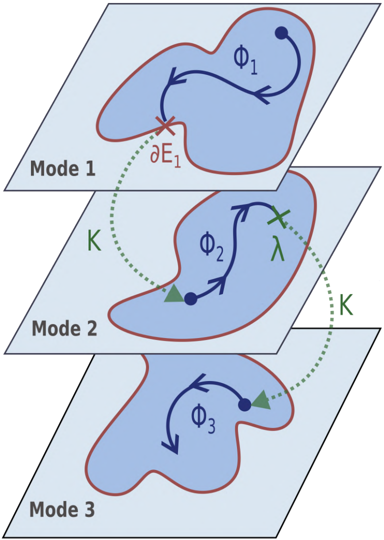

Likelihood of a PDMP trajectory

Let $\mathbf{Z} := (Z_t)_{t\in[0,t_{\max}]}$ be a PDMP trajectory of duration $t_{\max}$ on E.

Density function of a PDMP trajectory (*Thomas Galtier 2019*)

There is a dominant measure ζ for which a PDMP trajectory $\mathbf{Z}$ with $n_{\mathbf{Z}}$ jumps, inter-jump times $(t_k)_k$ and arrival states $(z_k)_k$ admits a probability density function f.

$$f(\mathbf{Z}) = \prod_{k=0}^{n_{\mathbf{Z}}} \left[\lambda\left(\Phi_{z_k}(t_k)\right)\right]^{\mathbb{1}_{t_k<\tau^{\partial}_{z_k}}} \exp\left[-\int_0^{t_k} \lambda\left(\Phi_{z_k}(u)\right) \mathrm{d}u\right] K\left(\Phi_{z_k}(t_k), z_{k+1}\right)^{\mathbb{1}_{k<n_{\mathbf{Z}}}} . \tag{1}$$

Take home message:

- explicit computation of the pdf of a PDMP trajectory,
- no need to recalculate the flow.

Rare event simulation

Main problem

Objective: estimate

$$\mathrm{P}_{\mathcal{F}} = \mathbb{P}_{f_0}\left(\mathbf{Z} \in \mathcal{T}_{\mathcal{F}}\right) = \mathbb{P}_{f_0}\left(\exists t \in [0, t_{\max}] : Z_t \in \mathcal{F}\right)$$

- $\mathbf{Z} = (Z_t)_{t\in[0,t_{\max}]}$ is a PDMP trajectory of fixed duration $t_{\max}$,
- $\mathbf{Z} \sim f_0$ the reference distribution of the PDMP trajectory,
- $\mathcal{T}_{\mathcal{F}}$ is the set of feasible PDMP trajectories that reach a critical region $\mathcal{F}$ of the state space before time $t_{\max}$.

Crude Monte-Carlo : $\widehat{\mathrm{P}}_{\mathcal{F}}^{\mathrm{CMC}} = \dfrac{1}{N}\sum_{k=1}^{N} \mathbb{1}_{\mathbf{Z}_k \in \mathcal{T}_{\mathcal{F}}}$ with $\mathbf{Z}_1, \ldots, \mathbf{Z}_N \overset{\text{i.i.d.}}{\sim} f_0$

$\hookrightarrow$ Requires on average $1/\mathrm{P}_{\mathcal{F}}$ simulations to obtain one realization of the event.

$\hookrightarrow$ High relative variance of $\widehat{\mathrm{P}}_{\mathcal{F}}^{\mathrm{CMC}}$ when $\mathrm{P}_{\mathcal{F}}$ is small.

Importance sampling (IS)

Idea: simulate trajectory $\mathbf{Z}$ according to an alternative distribution g which gives more weight on $\mathcal{T}_{\mathcal{F}}$ than f_0, then fix the bias with the likelihood ratio $w = f_0/g$.

Importance sampling trick with alternative distribution g :

$$\mathrm{P}_{\mathcal{F}} = \mathbb{P}_{f_0}\left(\mathbf{Z} \in \mathcal{T}_{\mathcal{F}}\right) = \mathbb{E}_{f_0}\left[\mathbb{1}_{\mathbf{Z}\in\mathcal{T}_{\mathcal{F}}}\right] = \int \mathbb{1}_{\mathbf{z}\in\mathcal{T}_{\mathcal{F}}} f_0(\mathbf{z})d\zeta(\mathbf{z}) \quad (2)$$

$$= \int \mathbb{1}_{\mathbf{z}\in\mathcal{T}_{\mathcal{F}}} \frac{f_0(\mathbf{z})}{g(\mathbf{z})} g(\mathbf{z})d\zeta(\mathbf{z}) = \mathbb{E}_g\left[\mathbb{1}_{\mathbf{Z}\in\mathcal{D}} \frac{f_0(\mathbf{Z})}{g(\mathbf{Z})}\right] \quad (3)$$

IS estimator : $$\widehat{\mathrm{P}}^{\mathrm{IS}}_{\mathcal{F}} = \frac{1}{N}\sum_{k=1}^{N} \mathbb{1}_{\mathbf{Z}_k\in\mathcal{T}_{\mathcal{F}}} \frac{f_0(\mathbf{Z}_k)}{g(\mathbf{Z}_k)} \quad \text{with } \mathbf{Z}_1,\ldots,\mathbf{Z}_N \overset{\text{i.i.d.}}{\sim} g \quad (4)$$

$\hookrightarrow$ Variance of $\widehat{\mathrm{P}}^{\mathrm{IS}}_{\mathcal{F}}$ relies on the choice of g

Optimal importance sampling

- **Optimal IS distribution:** $g_{\text{opt}} : \mathbf{z} \mapsto \frac{1}{\mathrm{P}_{\mathcal{F}}} \mathbb{1}_{\mathbf{z} \in \mathcal{T}_{\mathcal{F}}} f_0(\mathbf{z})$ produces a zero-variance IS estimator.

- **PDMP case**: the optimal IS distribution g_{opt} is fully determined by the so-called **committor function** U_{opt} of the process. Knowing U_{opt} is sufficient to generate PDMP trajectories under g_{opt}.

- **Committor function:** probability of realizing the rare event $\{\mathbf{Z} \in \mathcal{T}_{\mathcal{F}}\}$ knowing the state of the process at any given time $s \in [0, t_{\max}]$.

$$\mathrm{U}_{\text{opt}}\left(Z_s\right) = \mathbb{P}_{f_0}\left(\mathbf{Z} \in \mathcal{T}_{\mathcal{F}} \,|\, Z_s\right)^{\mathbb{1}_{\mathbf{Z}_s \notin \mathcal{T}_{\mathcal{F}}}} \quad \text{with} \quad \mathbf{Z}_s = \left(Z_t\right)_{t \in [0,s]}. \tag{5}$$

- **General committor function:** when estimating $\mathbb{E}_{f_0}\left[\varphi(\mathbf{Z})\right]$ we have

$$\mathrm{U}_{\text{opt}}\left(\mathbf{Z}_s\right) = \mathbb{E}_{f_0}\left[\varphi(\mathbf{Z}) \mid \mathbf{Z}_s\right] \quad \text{with} \quad \mathbf{Z}_s = \left(Z_t\right)_{t \in [0,s]}.$$

Optimal biasing with committor and edge committor function

Edge committor function U^-_{opt} : mean value of the committor function knowing the process is about to jump with reference jump kernel K_0.

$$\text{"}\ U^-_{\text{opt}}(Z_s^-) = \mathbb{E}_{K_0(Z_s^-,\cdot)}\left[U_{\text{opt}}\left(Z_s\right)\right]\ \text{"}. \tag{6}$$

Optimal jump intensity and jump kernel: *(Thomas Galtier 2019)*

$$\text{"}\ \lambda_{\text{opt}} = \lambda_0 \times \frac{U^-_{\text{opt}}}{U_{\text{opt}}}\ \text{"} \quad \text{and} \quad \text{"}\ K_{\text{opt}} = K_0 \times \frac{U_{\text{opt}}}{U^-_{\text{opt}}}\ \text{"}. \tag{7}$$

If the process is c times more likely to realize the event:

1. by jumping now from state z, then $\lambda_{\text{opt}}(z)$ should be c times $\lambda_0(z)$,
2. by jumping to state z from state z^- rather than jumping randomly from z^-, then $K_{\text{opt}}(z^-, z)$ should be c times $K_0(z^-, z)$.

Our method in a nutshell

Chennetier, Chraibi, Dutfoy, Garnier (2022), Adaptive importance sampling based on fault tree analysis for piecewise deterministic Markov process. *arXiv preprint arXiv:2210.16185*.

1. Building a family of approximations of the committor function U_{opt}.
   - *First contribution:* Fault tree analysis (minimal path sets and cut sets),
   - *Current work:* Mean hitting times of a random walk on a graph.

2. The best representative of this family is sequentially determined using a cross-entropy procedure coupled with a recycling scheme for past samples.

3. A consistent and asymptotically normal post-processing estimator of the final probability $P_{\mathcal{F}}$ is returned.

Approximating U_{opt} with graph-based mean hitting times

PDMP approximated by a random walk on a graph

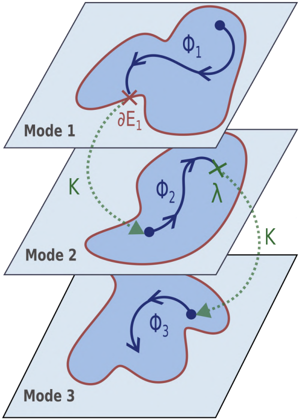

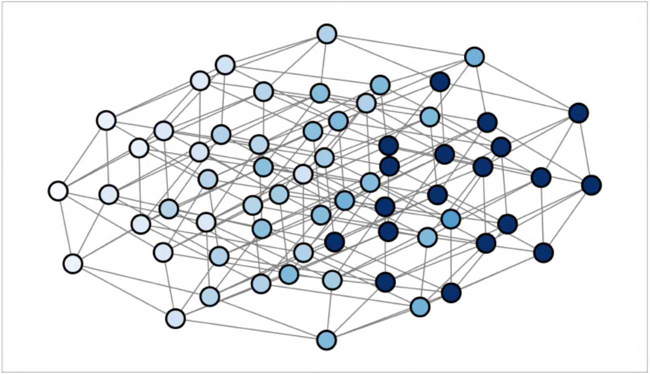

Figure 1: PDMP with 64 modes, $\mathcal{M}_{\mathcal{F}}$ in dark blue.

$\mathbf{Z} \in \mathcal{T}_{\mathcal{F}}$ only if the trajectory stays long enough in a mode of $\mathcal{M}_{\mathcal{F}}$.

Mean hitting times for a Markovian homogeneous random walk

- Let $(Y_t)_t$ be a time-continuous random walk on the mode set $\mathcal{M}$ with an infinitesimal generator matrix Q.
- We note $h_m = \mathbb{E}\left[\tau_m\left(\mathcal{M}_{\mathcal{F}}\right)\right]$ with $\tau_m\left(\mathcal{M}_{\mathcal{F}}\right) = \inf_{t\geq 0}\left\{Y_t \in \mathcal{M}_{\mathcal{F}} \mid Y_0 = m\right\}$.
- If the random walk is time-homogeneous then $(h_m)_{m\in\mathcal{M}}$ the vector of mean hitting times of $\mathcal{M}_{\mathcal{F}}$ is explicit and solution of the linear system:

$$h_{m_1} = 0 \;\forall m_1 \in \mathcal{M}_{\mathcal{F}} \text{ and } \sum_{m_2 \notin \mathcal{M}_{\mathcal{F}}} Q[m_1, m_2]\, h_{m_2} = -1 \;\forall m_1 \notin \mathcal{M}_{\mathcal{F}}. \tag{8}$$

Idea: compute $(h_m)_{m\in\mathcal{M}}$ for a matrix Q chosen such that $(Y_t)_t$ "behaves like" $(M_t)_t$ the mode part of the PDMP trajectory $(Z_t)_t$.

$\hookrightarrow$ *In practice even using the simple random walk gives good results.*

Minimal support condition: for any $m_1, m_2 \in \mathcal{M}$, $Q[m_1, m_2] > 0$ only if there are $x_1, x_2 \in \mathcal{X}$ such that $K\left((x_1, m_1), (x_2, m_2)\right) > 0$.

Proximity score and approximation of $\mathrm{U}_{\mathrm{opt}}$

1. For each mode $m \in \mathcal{M}$, we set ρ_m the proximity score to the set $\mathcal{M}_{\mathcal{F}}$:

$$\rho_m = 1 - \frac{h_m}{\max\limits_{m' \in \mathbb{M}} \{h_{m'}\}} \in [0, 1].$$

2. We define a family $(U_{\boldsymbol{\theta}})_{\boldsymbol{\theta} \in \Theta}$ of approximations of $\mathrm{U}_{\mathrm{opt}}$ parameterized by a vector $\boldsymbol{\theta} \in \Theta \subset \mathbb{R}^{d_\Theta}$ of arbitrary size d_Θ.

$$U_{\boldsymbol{\theta}}((x, m)) = \exp\left(\sum_{k=1}^{d_\Theta} \theta_k \times \psi_{k, d_\Theta}(\rho_m)\right) \tag{9}$$

The sequence $(\psi_{k,\infty})_{k \in \mathbb{N}^*}$ is typically a basis of $L^2([0,1])$. For example:

- Polynomial: $\psi_{k,d_\Theta}(\rho) = \rho^k$.
- Piecewise linear: $\psi_{k,d_\Theta}(\rho) = \rho \, \mathbb{1}_{\rho > \frac{k-1}{d_\Theta}}$.

Example with a simple random walk

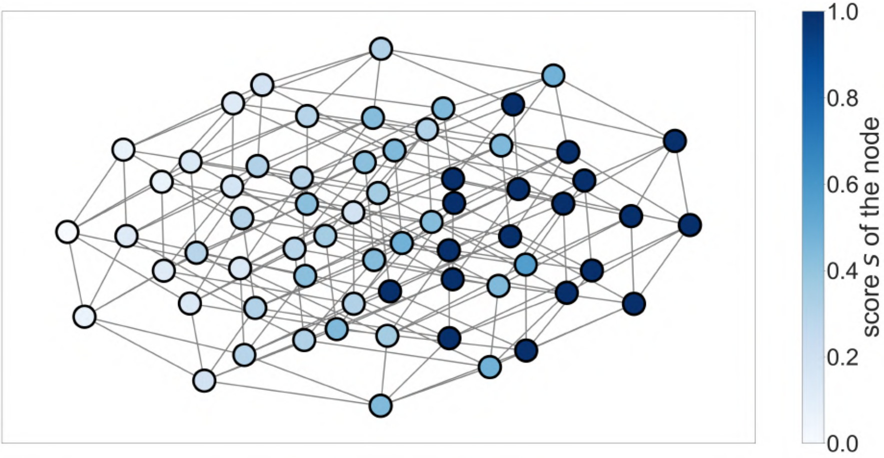

Figure 2: **Scores on a graph with 64 vertices.** $\mathcal{M}_{\mathcal{F}}$ is given by the vertices with score 1.

Recycling adaptive IS

The cross entropy procedure

How to find the best candidate within the family $(U_{\boldsymbol{\theta}})_{\boldsymbol{\theta}\in\Theta}$?

To each candidate $U_{\boldsymbol{\theta}} \in (U_{\boldsymbol{\theta}})_{\theta\in\Theta}$ corresponds an importance distribution $g_{\boldsymbol{\theta}} \in (g_{\boldsymbol{\theta}})_{\theta\in\Theta}$. We look for the closest distribution $g_{\boldsymbol{\theta}}$ to g_{opt} in the sense of the Kullback-Leibler divergence.

$$
\begin{aligned}
\underset{\boldsymbol{\theta}\in\Theta}{\arg\min}\, \mathcal{D}_{\text{KL}}\left(g_{\text{opt}} \| g_{\boldsymbol{\theta}}\right) &= \underset{\boldsymbol{\theta}\in\Theta}{\arg\min}\, \mathbb{E}_{g_{\text{opt}}}\left[\log\left(\frac{g_{\text{opt}}(\mathbf{Z})}{g_{\boldsymbol{\theta}}(\mathbf{Z})}\right)\right] \\
&= \underset{\boldsymbol{\theta}\in\Theta}{\arg\min} \int -\log\left(g_{\boldsymbol{\theta}}(\mathbf{Z})\right) \frac{\mathbb{1}_{\mathbf{Z}\in\mathcal{T}_{\mathcal{F}}}\, f_0(\mathbf{Z})}{\mathrm{P}_{\mathcal{F}}} d\zeta\left(\mathbf{Z}\right) \\
&= \underset{\boldsymbol{\theta}\in\Theta}{\arg\max}\, \mathbb{E}_{f_0}\left[\mathbb{1}_{\mathbf{Z}\in\mathcal{T}_{\mathcal{F}}} \log\left(g_{\boldsymbol{\theta}}\left(\mathbf{Z}\right)\right)\right].
\end{aligned}
$$

This last quantity can be minimized iteratively by successive Monte-Carlo approximations with importance sampling.

Adaptive algorithm with recycling of past samples

Start with an initial parameter $\boldsymbol{\theta}^{(1)}$. At iteration $j = 1, \dots, J$:

1. **Simulation step**: generate a new sample of n_j trajectories

$$\mathbf{Z}_{j,1}, \dots, \mathbf{Z}_{j,n_j} \overset{\text{i.i.d.}}{\sim} g_{\boldsymbol{\theta}^{(j)}}$$

2. **Optimization step**: compute the next iterate $\boldsymbol{\theta}^{(j+1)}$ by solving:

$$\boldsymbol{\theta}^{(j+1)} \in \underset{\boldsymbol{\theta} \in \Theta}{\arg\max} \sum_{i=1}^{j} \sum_{k=1}^{n_j} \mathbb{1}_{\mathbf{Z}_{i,k} \in \mathcal{T}_{\mathcal{F}}} \frac{f_0\left(\mathbf{Z}_{i,k}\right)}{g_{\boldsymbol{\theta}^{(i)}}\left(\mathbf{Z}_{i,k}\right)} \log\left[g_{\boldsymbol{\theta}}\left(\mathbf{Z}_{i,k}\right)\right] \tag{10}$$

Estimation step: at iteration J, the final estimator of the probability $\mathrm{P}_{\mathcal{F}}$ is:

$$\widehat{\mathrm{P}}_{\mathcal{F}} = \frac{1}{\sum_{j=1}^{J} n_j} \sum_{j=1}^{J} \sum_{k=1}^{n_j} \mathbb{1}_{\mathbf{Z}_{j,k} \in \mathcal{T}_{\mathcal{F}}} \frac{f_0\left(\mathbf{Z}_{j,k}\right)}{g_{\boldsymbol{\theta}^{(j)}}\left(\mathbf{Z}_{j,k}\right)} \tag{11}$$

Recycling scheme: past samples are reused during optimization and estimation.

We proved **consistency** and **asymptotic normality** of $\widehat{\mathrm{P}}_{\mathcal{F}}$ for the PDMP case.

Numerical results

Performances on the spent fuel pool

Test case: Spent fuel pool from nuclear industry. The corresponding graph has 32, 768 vertices.

Method	N	$\widehat{\mathrm{P}}_{\mathcal{F}}$	$\widehat{\sigma}/\widehat{\mathrm{P}}_{\mathcal{F}}$	95% confidence interval
CMC	10^5	2×10^{-5}	223.60	$[0\,;\,4.77 \times 10^{-5}]$
	10^6	1.3×10^{-5}	277.35	$[5.93 \times 10^{-6}\,;\,2.01 \times 10^{-5}]$
	10^7	1.77×10^{-5}	237.68	$[1.51 \times 10^{-5}\,;\,2.03 \times 10^{-5}]$
AIS–MHT	10^3	1.86×10^{-5}	1.62	$[1.67 \times 10^{-5}\,;\,2.04 \times 10^{-5}]$
	10^4	2.01×10^{-5}	0.86	$[1.98 \times 10^{-5}\,;\,2.05 \times 10^{-5}]$

Table 1: **Comparison between crude Monte-Carlo (CMC) and our adaptive importance sampling method with mean hitting times (AIS-MHT).**

$\hookrightarrow$ Variance reduction by a factor of 10,000.

Robustness in practice

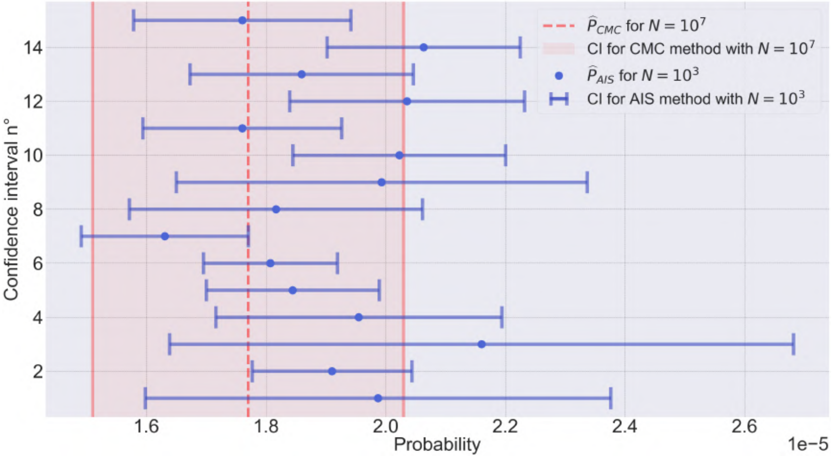

Figure 3: **15 confidence intervals with AIS–MHT method and sample size of 1000 vs 1 confidence interval with CMC method and sample size of** 10^7.

References

- M. H. A. Davis (1984), Piecewise-Deterministic Markov Processes: A general class of non-diffusion stochastic models. *J. R. Stat. Soc. B 46:353-388*.

- B. de Saporta, F. Dufour and H. Zhang (2015), Numerical methods for simulation and optimization of piecewise deterministic Markov processes : application to reliability. *Mathematics and statistics series Wiley-ISTE*.

- R. Rubinstein (1999), The Cross-Entropy method for combinatorial and continuous optimization. *Methodology and Computing in Applied Probability 1:127–190*.

- B. Delyon and F. Portier (2018), Asymptotic optimality of adaptive importance sampling. *NIPS 2018*.

- H. Chraibi, A. Dutfoy, T. Galtier et J. Garnier (2019), On the optimal importance process for piecewise deterministic Markov process. *ESAIM: PS. 23:893–921*.

- G. Chennetier, H. Chraibi, A. Dutfoy, J. Garnier (2022), Adaptive importance sampling based on fault tree analysis for piecewise deterministic Markov process. *arXiv preprint arXiv:2210.16185*.

The end

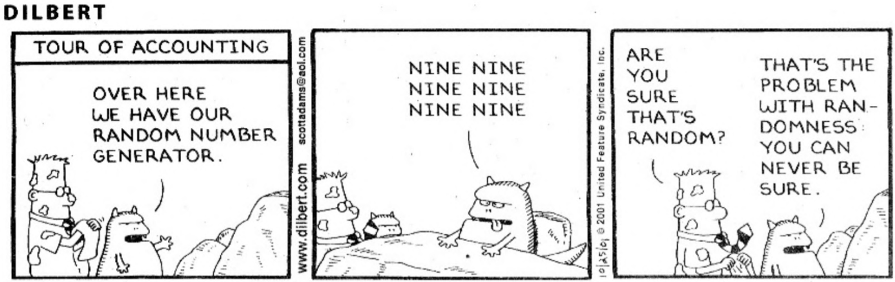

Thank you for your attention.

Supplementary material

Example: the spent fuel pool

If the system does not cool the pool, the nuclear fuel evaporates the water then damages the structure and contaminates the outside.

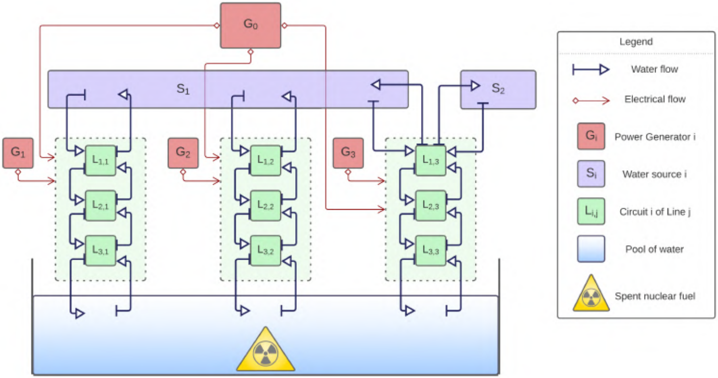

Aim: estimating the probability of the water level falling below a set threshold.

Mathematical details

- **Flow** Φ : solution of differential equations. Can be costly to solve. When no jump between time s and $s+t$:

$$Z_{s+t} = \Phi_{Z_s}(t).$$

- **Deterministic jumps** : when the position reaches ∂E the boundaries of E.

$$\tau_z^{\partial} = \inf\{t > 0 : \Phi_z(t) \in \partial E\}.$$

- **Jump intensity** λ : parameter of the distribution of the time T_z of the next random jump knowing current state z.

$$\mathbb{P}(\tau_z > t \mid Z_s = z) = \mathbb{1}_{t<t_z^{\partial}} \exp\left(- \int_0^t \lambda\left(\Phi_z(u)\right) du \right). \tag{12}$$

- **Jump kernel** K : for any departure state z^-, density $z \mapsto K(z^-, z)$ of a Markovian kernel $\mathcal{K}_{z^-}$ with respect to some measure ν_{z^-} on E.

Likelihood of a PDMP trajectory

Let $\mathbf{Z} := (Z_t)_{t\in[0,t_{\max}]}$ be a PDMP trajectory of duration $t_{\max}$ on E.

Density function of a PDMP trajectory (*Thomas Galtier 2019*)

There is a dominant measure ζ for which a PDMP trajectory $\mathbf{Z}$ with $n_{\mathbf{Z}}$ jumps, inter-jump times $t_1, \ldots, t_{n_{\mathbf{Z}}}$ and arrival states $z_1, \ldots, z_{n_{\mathbf{Z}}}$ admits a probability density function f.

$$f(\mathbf{Z}) = \prod_{k=0}^{n_{\mathbf{Z}}} \left[\lambda\left(\Phi_{z_k}(t_k)\right)\right]^{\mathbb{1}_{t_k<\tau^{\partial}_{z_k}}} \exp\left[-\int_0^{t_k} \lambda\left(\Phi_{z_k}(u)\right) \mathrm{d}u\right] K\left(\Phi_{z_k}(t_k), z_{k+1}\right)^{\mathbb{1}_{k<n_{\mathbf{Z}}}} . \tag{13}$$

Take home message:

- explicit computation of the pdf of a PDMP trajectory,
- no need to recalculate the flow.

Committor function for importance sampling

Committor function

Committor function: probability of realizing the rare event $\{\mathbf{Z} \in \mathcal{T}_{\mathcal{F}}\}$ knowing that at a fixed time $s > 0$ the process is in a given state z.

$$U_{\text{opt}}(z, s) = \mathbb{P}_{f_0}\left(\mathbf{Z} \in \mathcal{T}_{\mathcal{F}} \mid Z_s = z\right). \tag{14}$$

(in general $U_{opt}\left(\mathbf{Z}\right) = \mathbb{E}_{f_0}\left[\varphi(\mathbf{Z}) \mid \mathbf{Z}_s\right]$ *with* $\mathbf{Z}_s = (Z_t)_{t\in[0,s]}$ *when estimating* $\mathbb{E}_{\pi_0}\left[\varphi(\mathbf{Z})\right]$*)*

Knowing U_{opt} is sufficient to build the optimal IS estimator.

To lighten the future equations we also note the variant committor function U^{-}_{opt} :

$$U^{-}_{\text{opt}}(z^{-}, s) = \int_{z\in E} U_{\text{opt}}(z, s)\, K\left(z^{-}, z\right) d\nu_{z^{-}}. \tag{15}$$

U^{-}_{opt} is the probability of realizing the rare event $\{\mathbf{Z} \in \mathcal{T}_{\mathcal{F}}\}$ knowing that at a fixed time $s > 0$ the process jumps from a given state z^{-}.

Optimal IS for PDMP

Optimal jump intensity and jump kernel: *(Thomas Galtier 2019)*

$$\lambda_{\text{opt}}(\Phi_z(t)\,;\,s) = \lambda_0(\Phi_z(t)) \times \frac{U^-_{\text{opt}}(\Phi_z(t), s+t)}{U_{\text{opt}}(\Phi_z(t), s+t)}, \tag{16}$$

$$K_{\text{opt}}(z^-, z\,;\,s) = K_0(z^-, z) \times \frac{U_{\text{opt}}(z, s)}{U^-_{\text{opt}}(z^-, s)}. \tag{17}$$

If the process is k times more likely to realize the event:

1. by jumping now from state z, then $\lambda_{\text{opt}}(z)$ should be k times λ_0,
2. by going to state z after a jump from state z^-, then $K_{\text{opt}}(z^-, z)$ should be k times $K_0(z^-, z)$.

Approximation with MPS

Approximation of the committor function with minimal path sets

The path sets of a system are the sets of components such that:

1. keeping all components of any path set intact prevents system failure.
2. keeping one component broken in each path set ensures system failure.

A **M**inimal **P**ath **S**et is a path set that does not contain any other path set.

We note:

- d_{MPS} the number of MPS (they are unique if the system is coherent),
- $\beta^{(\mathrm{MPS})}(z)$ the number of MPS with at least one broken component.

A good U_θ should therefore be increasing in $\beta^{(\mathrm{MPS})}(z)$.

Minimal path sets: the spent fuel pool case

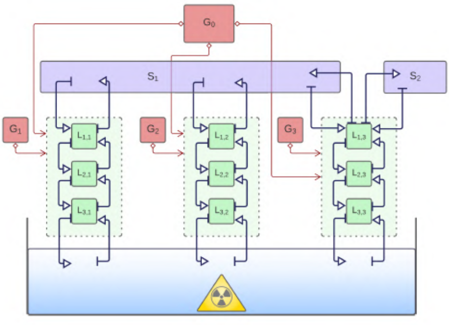

Figure 4: Physical representation of the SFP

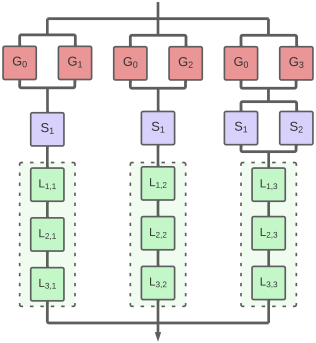

Figure 5: Functionnal diagram of the SFP

8 MPS in the spent fuel pool system: (with $\boldsymbol{L_j} = (L_{i,j})_{i=1}^{3}$ for $j = 1, 2, 3$)

$$(G_0, S_1, \boldsymbol{L_1}), (G_1, S_1, \boldsymbol{L_1}), (G_0, S_1, \boldsymbol{L_2}), (G_2, S_1, \boldsymbol{L_2}),$$
$$(G_0, S_1, \boldsymbol{L_3}), (G_3, S_1, \boldsymbol{L_3}), (G_0, S_2, \boldsymbol{L_3}), (G_3, S_2, \boldsymbol{L_3}).$$

Our MPS-based proposition

For $\boldsymbol{\theta} \in \mathbb{R}_+^{d_{\mathbf{MPS}}}$ we propose:

$$U_{\boldsymbol{\theta}}^{(\mathrm{MPS})}(z) = \exp\left[\left(\sum_{i=1}^{\beta^{(\mathbf{MPS})}(z)} \theta_i\right)^2\right]. \tag{18}$$

Flexible dimension of $\boldsymbol{\theta}$: imposing equality on some coordinates of $\boldsymbol{\theta}$ reduce its effective dimension and simplify the search for a good $\boldsymbol{\theta}$ when d_{MPS} is large.

→ Example for dimension 1 with $\theta_1 = \cdots = \theta_{d_{\mathbf{MPS}}}$:

$$U_{\boldsymbol{\theta}}^{(\mathrm{MPS})}(z) = \exp\left[\left(\theta_1\, \beta^{(\mathrm{MPS})}(z)\right)^2\right]. \tag{19}$$

The form $x \mapsto \exp(x^2)$ garantees that the ratios $U_{\boldsymbol{\theta}}^{-}/U_{\boldsymbol{\theta}}$ are strictly increasing in $\beta^{(\mathrm{MPS})}$. Without this condition, it is increasingly difficult to break new components and they are repaired faster and faster as they are lost.

Minimal cut sets

Minimal cut sets: smallest sets of components that if left broken ensure system failure. (permanent repair of one component in each group prevents the failure)

In this system: there is 69 minimal cut sets for 15 components.

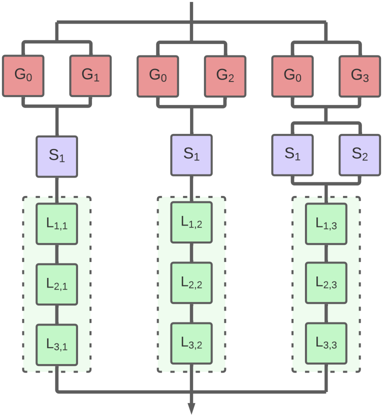

Figure 6: Functionnal diagram of the SFP

Adaptive algorithm

Asymptotic confidence interval

Assumptions

1. The functions λ, K, and $(U_\theta)_{\theta \in \Theta}$ are bounded on their support below and above by strictly positive constants,
2. $\theta_{\text{opt}} \in \Theta$ is the unique maximizer of $\theta \mapsto \mathbb{E}_{f_0}\left[\mathbb{1}_{\mathbf{Z} \in \mathcal{T}_{\mathcal{F}}} \log g_\theta(\mathbf{Z})\right]$,
3. there is $t_\varepsilon > 0$ such that $t_z^\partial \geq t_\varepsilon$ for any $z^- \in \partial E$ and any $z \in \operatorname{supp} K(z^-, \cdot)$.

Under these assumptions, with $\widehat{\sigma}^2 = \frac{1}{\sum_{j=1}^J n_j} \sum_{j=1}^J \sum_{k=1}^{n_j} \mathbb{1}_{\mathbf{Z}_{j,k} \in \mathcal{T}_{\mathcal{F}}} \frac{f_0\left(\mathbf{Z}_{j,k}\right)^2}{g_{\theta(j)}\left(\mathbf{Z}_{j,k}\right)^2} - \widehat{\mathrm{P}}_{\mathcal{F}}^2$ the estimator of the asymptotic variance $\mathbb{E}_{f_0}\left[\mathbb{1}_{\mathbf{Z} \in \mathcal{T}_{\mathcal{F}}} \frac{f_0(\mathbf{Z})}{g_{\theta_{\text{opt}}}}(\mathbf{Z})\right] - \mathrm{P}_{\mathcal{F}}^2$, and with $v_{1-\alpha/2}$ the $(1-\alpha/2)$-quantile of the $\mathcal{N}(0,1)$ distribution, we have :

$$\mathbb{P}\left(\mathrm{P}_{\mathcal{F}} \in \left[\widehat{\mathrm{P}}_{\mathcal{F}} - v_{1-\alpha/2}\sqrt{\widehat{\sigma}^2/N_J}\,;\, \widehat{\mathrm{P}}_{\mathcal{F}} + v_{1-\alpha/2}\sqrt{\widehat{\sigma}^2/N_J}\right]\right) \xrightarrow[N_J \to \infty]{} 1 - \alpha.$$

Bandit problem

Off-policy best arm identification in multi-armed bandit

- Several nominal distributions $\pi_1, \ldots, \pi_d$,
- $\mu_i := \mathbb{E}_{\pi_i}[\varphi(\mathbf{Z})]$ for $i = 1, \ldots, d$ and a function φ (example: $\varphi = \mathbb{1}_{\mathcal{D}}$).

Objective: find the best distribution $\underset{i \in \{1,\ldots,d\}}{\arg\min}\, \mathbb{E}_{\pi_i}[\varphi(\mathbf{Z})]$

Reverse importance sampling: if we draw $(\mathbf{Z}_1, \ldots, \mathbf{Z}_N) \sim \left(\otimes_{k=1}^N q_k\right)$ then:

$$\widehat{\mu}_i = \frac{1}{N} \sum_{k=1}^N \varphi(\mathbf{Z}_k) \frac{\pi_i(\mathbf{Z}_k)}{q_k(\mathbf{Z}_k)} \quad \text{for any } i = 1, \ldots, d$$

Best sequential sampling policy?

Stability of the method

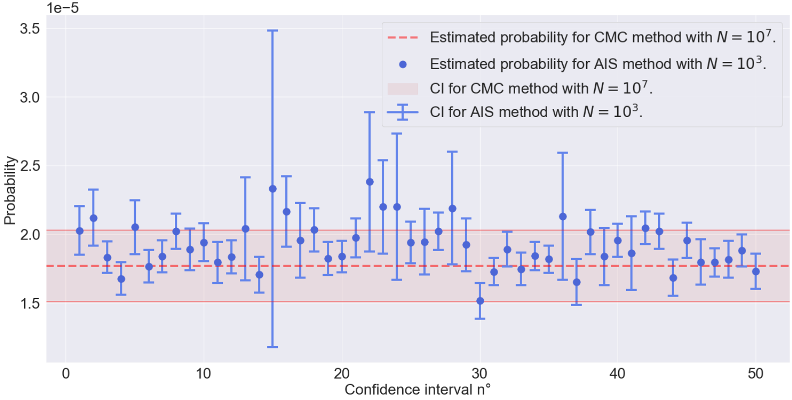

Figure 7: **50 confidence intervals with AIS–MPS method and sample size of 1000 vs 1 confidence interval with CMC method and sample size of 10^7.**