On the Hill Relation and the Mean Reaction Time for Metastable Processes

Arnaud Guyader

Sorbonne Université / INRIA MATHERIALS

Joint work with Manon Baudel and Tony Lelièvre

CIRM, April 2023

Mean reaction time for metastable processes

- (X_t) a continuous-time ergodic Markov process
- A and B bounded disjoint open sets with smooth boundaries
- Goal: estimate T_{AB}, the mean reaction time at equilibrium.

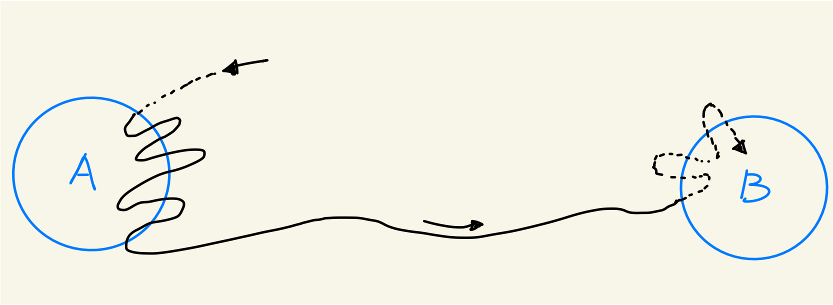

Reactive entrance distribution

- Set $\tau_0^A = \inf\{t > 0,\ X_t \in \bar{A}\}$, $\tau_0^B = \inf\{t > \tau_0^A,\ X_t \in \bar{B}\}$, ...
- Empirical entrance distribution in A: $\mu_N = \frac{1}{N}\sum_{n=1}^{N} \delta_{X_{\tau_n^A}}$
- Reactive entrance distribution ν_E = weak limit of μ_N
- Mean reaction time [E, Vanden Eijnden, 2006],[Lu, Nolen, 2015]:

$$T_{AB} := \mathbb{E}^{\nu_E}[\tau_B] = \lim_{N\to\infty} \frac{1}{N}\sum_{n=1}^{N}(\tau_n^B - \tau_n^A) \ a.s.$$

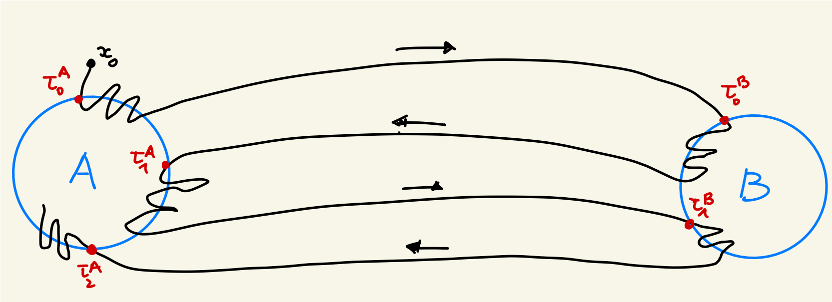

Metastability

- Example: overdamped Langevin dynamics

$$dX_t = -\nabla V(X_t)\, dt + \sqrt{2\beta^{-1}} dW_t$$

which is ergodic wrt

$$\mu(dx) = Z^{-1} \exp(-\beta V(x))\, dx$$

- Difficulty: A and B are typically metastable states, so that the mean reaction time T_{AB} is very large.

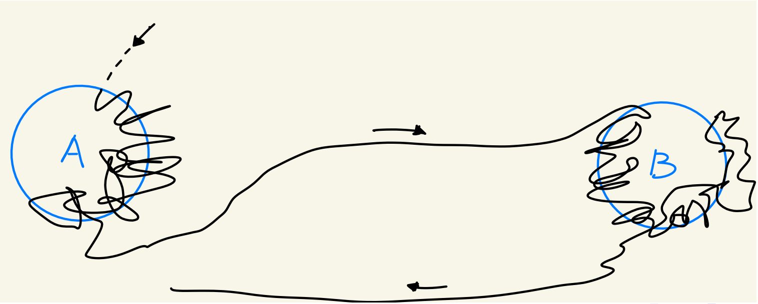

From (X_t) to (Y_n)

- Let Σ a co-dimension 1 submanifold in-between A and B
- Define $\mathcal{A} = \partial A$, $\mathcal{B} = \partial B$, and $Y_n = X_{\tau_n}$ where

$$\tau_n^\Sigma = \inf\{t > \tau_{n-1}, X_t \in \Sigma\}, \qquad \tau_n = \inf\{t > \tau_n^\Sigma, X_t \in \mathcal{A} \cup \mathcal{B}\}$$

- The Markov chain $(Y_n)_{n\geq 0}$ is with values in $\mathcal{A} \cup \mathcal{B}$.

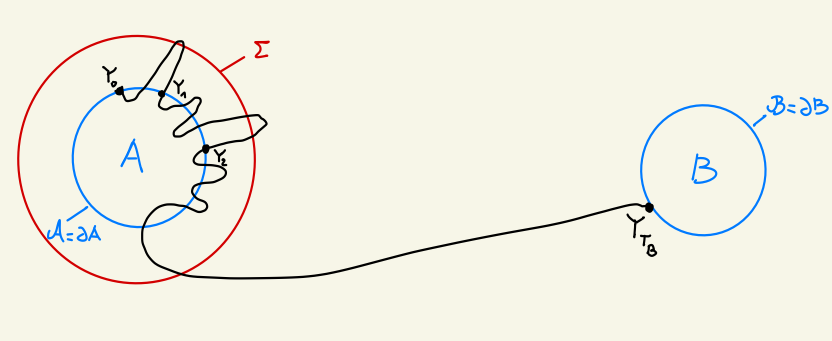

Reactive entrance distribution

- Successive entrance times in $\mathcal{A}$ and $\mathcal{B}$:

$$T_{k+1}^{\mathcal{A}} = \inf\{n > T_k^{\mathcal{B}},\ Y_n \in \mathcal{A}\}, \qquad T_{k+1}^{\mathcal{B}} = \inf\{n > T_{k+1}^{\mathcal{A}},\ Y_n \in \mathcal{B}\}$$

- Reactive entrance distribution in $\mathcal{A}$ at equilibrium

$$\frac{1}{N}\sum_{n=1}^{N} \delta_{Y_{T_n^{\mathcal{A}}}} \Rightarrow \nu_E$$

- Recall: ν_E does not depend on the choice of Σ.

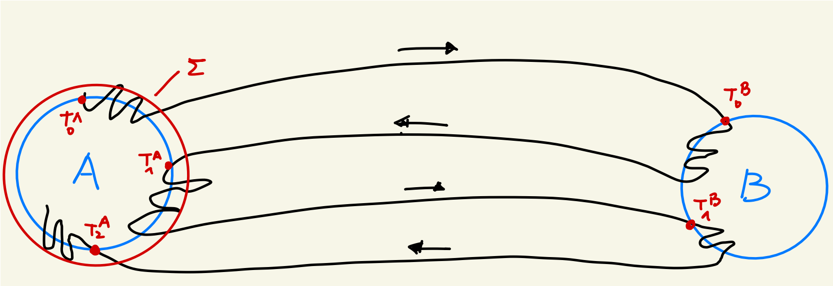

Back to the mean reaction time

- The mean reaction time rewrites as (strong Markov property)

$$T_{AB} := \mathbb{E}^{\nu_E}[\tau_B] = \mathbb{E}^{\nu_E}\left[\sum_{n=0}^{T_{\mathcal{B}}-1} \Delta(Y_n)\right]$$

where $T_{\mathcal{B}} = \inf\{n \geq 0,\ Y_n \in \mathcal{B}\}$ and, for all $x \in \mathcal{A}$,

$$\Delta(x) := \mathbb{E}^x[\tau_1] = \mathbb{E}^x[\tau_1 \mathbf{1}_{Y_1 \in \mathcal{A}}] + \mathbb{E}^x[\tau_1 \mathbf{1}_{Y_1 \in \mathcal{B}}].$$

- Two challenges: $\mathcal{A}$ and $\mathcal{B}$ are metastable, so that (i) $T_{\mathcal{B}}$ is very large, and (ii) ν_E is difficult to sample...
  - For (i), use the Hill relation
  - For (ii), replace ν_E with a quasi stationary distribution.

Assumptions and notation

- The Markov chain (Y_n) with kernel K satisfies
  - [A1] $\mathcal{A}$ and $\mathcal{B}$ are compact disjoint sets
  - [A2] (Y_n) is weak-Feller: $Kf \in \mathcal{C}(\mathcal{A} \cup \mathcal{B}, \mathbb{R})$ if $f \in \mathcal{C}(\mathcal{A} \cup \mathcal{B}, \mathbb{R})$
  - [A3] (Y_n) is positive Harris recurrent with unique stationary law π_0
  - [A4] $\pi_0(\mathcal{A}) > 0$ and $\pi_0(\mathcal{B}) > 0$.
- Example: these assumptions are satisfied for the Markov chain (Y_n) built from the overdamped Langevin dynamics.
- Notation: we write the block-decomposition of the kernel K of the chain (Y_n) over $\mathcal{A} \cup \mathcal{B}$ as follows:

$$K = \begin{bmatrix} K_{\mathcal{A}} & K_{\mathcal{A}\mathcal{B}} \\ K_{\mathcal{B}\mathcal{A}} & K_{\mathcal{B}} \end{bmatrix}.$$

The π-return process

- Let π be a probability measure on $\mathcal{A}$. The π-return process (Y_n^π) is the Markov chain with values in $\mathcal{A}$ and kernel K^π defined as: $\forall x \in \mathcal{A}$, $\forall C \subset \mathcal{A}$,

$$K^\pi(x, C) := \mathbb{P}^x(Y_1 \in C, T_\mathcal{B} > 1) + \mathbb{P}^x(Y_1 \in \mathcal{B})\pi(C).$$

i.e. (Y_n^π) is the chain (Y_n) "reset to π" each time Y_n enters $\mathcal{B}$.

- Lemma: $(Y_n^\pi)_{n\geq 0}$ admits the unique stationary distribution

$$R(\pi) := \frac{\pi(\mathrm{Id}_\mathcal{A} - K_\mathcal{A})^{-1}}{\mathbb{E}^\pi[T_\mathcal{B}]}.$$

- Remark: Such processes are classic in MD when people introduce a source in A and a sink in B to create a non-equilibrium flux from A to B [Farkas, 1927][Kramers, 1940]: Weighted Ensemble [Zuckerman, Aristoff], Milestoning [Elber, Vanden Eijnden], Transition Interface Sampling [Bolhuis, Van Erp], etc.

The Hill relation

- The Hill relation [Hill, 1977][Aristoff, 2018][Baudel, AG, Lelièvre, 2022]

$$\mathbb{E}^{\pi}\left[\sum_{n=0}^{T_{\mathcal{B}}-1} f(Y_n)\right] = \frac{\mathbb{E}^{R(\pi)}[f(Y_0)]}{\mathbb{P}^{R(\pi)}(Y_1 \in \mathcal{B})}.$$

- Application to $\pi = \nu_E$: The probability measure $R(\nu_E)$ is the stationary distribution π_0 restricted to $\mathcal{A}$:

$$R(\nu_E) = \frac{\pi_0 \mathbf{1}_{\mathcal{A}}}{\pi_0(\mathcal{A})} =: \pi_{0|\mathcal{A}}.$$

- Consequence: with $\Delta(x) = \mathbb{E}^x[\tau_1]$, we get

$$T_{AB} = \mathbb{E}^{\nu_E}\left[\sum_{n=0}^{T_{\mathcal{B}}-1} \Delta(Y_n)\right] = \frac{\mathbb{E}^{\pi_{0|\mathcal{A}}}[\tau_1]}{\mathbb{P}^{\pi_{0|\mathcal{A}}}(Y_1 \in \mathcal{B})}.$$

The Hill relation to compute T_{AB}

Another formulation:

$$T_{AB} = \Delta_{Loop}(\pi_{0|\mathcal{A}}) \left(\frac{1}{\mathbb{P}^{\pi_{0|\mathcal{A}}}(Y_1 \in \mathcal{B})} - 1 \right) + \Delta_{React}(\pi_{0|\mathcal{A}})$$

where

- $\Delta_{Loop}(\pi_{0|\mathcal{A}}) := \mathbb{E}^{\pi_{0|\mathcal{A}}}[\tau_1 | Y_1 \in \mathcal{A}]$ is the mean time for a loop from $\pi_{0|\mathcal{A}}$ back to $\mathcal{A}$,
- $\Delta_{React}(\pi_{0|\mathcal{A}}) := \mathbb{E}^{\pi_{0|\mathcal{A}}}[\tau_1 | Y_1 \in \mathcal{B}]$ is the mean time of a reactive trajectory from $\pi_{0|\mathcal{A}}$ to $\mathcal{B}$.

Problem: $\pi_{0|\mathcal{A}}$ is in general unknown and difficult to sample...
Hope: since $\mathcal{A}$ is metastable, maybe it is not needed to sample $\pi_{0|\mathcal{A}}$ since, typically, the process will reach a local equilibrium within $\mathcal{A}$ before going to $\mathcal{B}$.

The quasi-stationary distribution (QSD)

Lemma: Under our assumptions, (Y_n) admits a QSD ν_Q in $\mathcal{A}$, namely a probability measure on $\mathcal{A}$ such that: $\forall C \subset \mathcal{A}$,

$$\nu_Q(C) = \mathbb{P}^{\nu_Q}(Y_1 \in C | T_{\mathcal{B}} > 1).$$

Remarks:

- Starting from ν_Q, $T_{\mathcal{B}}$ is geometrically distributed, with parameter $\mathbb{P}^{\nu_Q}(Y_1 \in \mathcal{B})$.
- QSD and Yaglom limit: if $\mathcal{L}(Y_n | T_{\mathcal{B}} > n)$ admits a limit when $n \to \infty$, this limit is a QSD.
- The ν_Q-return process admits ν_Q as an invariant distribution:

$$R(\nu_Q) = \nu_Q$$

The Hill relation applied to $\pi = \nu_Q$

- Since $R(\nu_Q) = \nu_Q$, one has

$$\mathbb{E}^{\nu_Q}\left[\sum_{n=0}^{T_{\mathcal{B}}-1} \Delta(Y_n)\right] = \frac{\mathbb{E}^{\nu_Q}[\tau_1]}{\mathbb{P}^{\nu_Q}(Y_1 \in \mathcal{B})}.$$

- Back to the mean reaction time:

$$\mathbb{E}^{\nu_Q}\left[\sum_{n=0}^{T_{\mathcal{B}}-1} \Delta(Y_n)\right] = \Delta_{Loop}(\nu_Q)\left(\frac{1}{\mathbb{P}^{\nu_Q}(Y_1 \in \mathcal{B})} - 1\right) + \Delta_{React}(\nu_Q)$$

- What did we gain, compared to $\pi = \nu_E$? The probability distribution ν_Q can be sampled by brute force Monte Carlo.

The algorithm to estimate T_{AB}

To estimate

$$\mathbb{E}^{\nu_Q}\left[\sum_{n=0}^{T_{\mathcal{B}}-1} \Delta(Y_n)\right] = \Delta_{Loop}(\nu_Q)\left(\frac{1}{\mathbb{P}^{\nu_Q}(Y_1 \in \mathcal{B})} - 1\right) + \Delta_{React}(\nu_Q)$$

- Simulate (X_t) in a neighborhood of A, registering the successive loops $\mathcal{A} \to \Sigma \to \mathcal{A}$. This gives samples with "asymptotic" law ν_Q, and an estimate of $\Delta_{Loop}(\nu_Q)$.
- Use a rare event sampling algorithm (e.g., AMS) to simulate reactive trajectories, starting from ν_Q. This provides estimates of $\mathbb{P}^{\nu_Q}(Y_1 \in \mathcal{B})$ and $\Delta_{React}(\nu_Q)$.

Relative error

- Recall: we are interested in the mean reaction time

$$T_{AB} = \mathbb{E}^{\nu_E}\left[\sum_{n=0}^{T_{\mathcal{B}}-1} \Delta(Y_n)\right]$$

- Principle of the method: replace ν_E with ν_Q.
- Question: under which conditions do we have

$$ERR := \left| \frac{T_{AB} - \mathbb{E}^{\nu_Q}\left[\sum_{n=0}^{T_{\mathcal{B}}-1} \Delta(Y_n)\right]}{T_{AB}} \right| \ll 1?$$

Reaction time and convergence time to the QSD

- Define

$$p^+ := \sup_{x \in \mathcal{A}} \mathbb{P}^x(Y_1 \in \mathcal{B})$$

then, $\forall x \in \mathcal{A}$, one has $\frac{1}{p^+} \leq \mathbb{E}^x[T_{\mathcal{B}}]$ and therefore $\frac{1}{p^+} \leq T_{AB}$.

- The convergence time to the QSD is linked to

$$T_Q^E := \|\nu_E H_Q\|_{\mathrm{TV}}$$

where

$$H_Q f(x) := \mathbb{E}^x \left[\sum_{n=0}^{T_{\mathcal{B}}-1} (f(Y_n) - \nu_Q f) \right]$$

- Why can T_Q^E be seen as a convergence time to the QSD?

$$T_Q^E \leq \sum_{n=0}^{\infty} \|\mathcal{L}^{\nu_E}(Y_n | T_{\mathcal{B}} > n) - \nu_Q\|_{\mathrm{TV}}$$

Example: uniform geometric ergodicity

- In the context of overdamped Langevin dynamics, one can show that: $\exists \nu_Q, \exists \alpha > 0, \exists \rho \in (0,1), \forall x \in \mathcal{A}, \forall n \geq 0,$

$$\|\mathcal{L}^x(Y_n | T_{\mathcal{B}} > n) - \nu_Q\|_{TV} \leq \alpha \rho^n$$

- In this case,

$$T_Q^E \leq \min\left(\frac{\alpha}{1-\rho}, \inf_{c \in (0,\alpha)} \frac{2}{1-c}\left\lceil \frac{\ln(c/\alpha)}{\ln \rho} \right\rceil\right)$$

- Remark: if $\rho \to 0$, then the RHS goes to $\min(\alpha, 2)$.
- Examples of sufficient conditions to get geometric ergodicity: double-sided condition [Birkhoff 1957], Dobrushin condition [Dobrushin, 1970], Meyn-Tweedie like conditions [Champagnat, Villemonais, 2017], etc.

Relative error

Proposition: If $p^+ T_Q^E < 1$, then

$$ERR \leq \frac{p^+ T_Q^E}{1 - p^+ T_Q^E} \left(1 + \frac{\|\Delta\|_\infty}{\pi_{0|\mathcal{A}}(\Delta)}\right)$$

- Since $\frac{1}{p^+} \leq T_{AB}$, this shows that the relative error is small if the reaction time is large compared to the convergence time to the QSD, i.e.

$$T_Q^E \ll \frac{1}{p^+} \implies ERR \ll 1$$

- Remark: We have checked on examples that the upper bound is sharp in various ways. In particular, one can replace p^+ neither by $\mathbb{P}^{\nu_Q}(Y_1 \in \mathcal{B})$ nor by $\mathbb{P}^{\nu_E}(Y_1 \in \mathcal{B})$ in the RHS.

Where to place Σ?

- Approximation: if $p^+ T_Q^E \ll 1$, then

$$T_{AB} := \mathbb{E}^{\nu_E}\left[\sum_{n=0}^{T_{\mathcal{B}}-1} \Delta(Y_n)\right] \approx \mathbb{E}^{\nu_Q}\left[\sum_{n=0}^{T_{\mathcal{B}}-1} \Delta(Y_n)\right]$$

- Recall: quantities to estimate

$$\mathbb{E}^{\nu_Q}\left[\sum_{n=0}^{T_{\mathcal{B}}-1} \Delta(Y_n)\right] = \Delta_{Loop}(\nu_Q)\left(\frac{1}{\mathbb{P}^{\nu_Q}(Y_1 \in \mathcal{B})} - 1\right) + \Delta_{React}(\nu_Q)$$

- Trade-off:
  - if Σ is too far from A, the approximation is not valid, problems to sample ν_Q and to estimate $\Delta_{Loop}(\nu_Q)$,
  - if Σ is too close to A, the precision and complexity in the estimations of $\mathbb{P}^{\nu_Q}(Y_1 \in \mathcal{B})$ and $\Delta_{Loop}(\nu_Q)$ deteriorate...

References

- D. Aristoff, *Analysis and optimization of weighted ensemble sampling*, ESAIM: M2AN, 2018
- M. Baudel, AG, and T. Lelièvre, *On the Hill relation and the mean reaction time for metastable processes*, Stochastic Processes and their Applications, 2022
- F. Cérou, AG, T. Lelièvre, and D. Pommier, *A multiple replica approach to simulate reactive trajectories*, Journal of Chemical Physics, 2011
- T. Lelièvre, M. Ramil, and J. Reygner, *Estimation of statistics of transitions and Hill relation for Langevin dynamics*, Annales de l'IHP, 2023
- J. Lu and J. Nolen, *Reactive trajectories and the transition path process*, Probability Theory and Related Fields, 2015