

Non-Stationary Forward Flux Sampling

Nils B. Becker,¹ Rosalind J. Allen,² and Pieter Rein ten Wolde¹

¹*FOM Institute for Atomic and Molecular Physics (AMOLF), Science Park 104, 1098 XG Amsterdam, The Netherlands*

²*SUPA, School of Physics and Astronomy, The University of Edinburgh, James Clerk Maxwell Building, The King's Buildings, Mayfield Road, Edinburgh EH9 3JZ, UK*

We present a new method, Non-Stationary Forward Flux Sampling, that allows efficient simulation of rare events in both stationary and non-stationary stochastic systems. The method uses stochastic branching and pruning to achieve uniform sampling of trajectories in phase space and time, leading to accurate estimates for time-dependent switching propensities and time-dependent phase space probability densities. The method is suitable for equilibrium or non-equilibrium systems, in or out of stationary state, including non-Markovian or externally driven systems. We demonstrate the validity of the technique by applying it to a one-dimensional barrier crossing problem that can be solved exactly, and show its usefulness by applying it to the time-dependent switching of a genetic toggle switch.

I. INTRODUCTION

Rare events, which occur infrequently but have important consequences, control the dynamical behavior of many physical systems, both in and out of equilibrium – classic examples include crystal nucleation, protein folding, earthquakes and traffic jams. When simulating such systems on a computer, some form of enhanced sampling is usually needed in order to generate any significant number of rare event samples on the time scale of the simulation. While a number of enhanced sampling methods are available for systems in steady state, many important rare event processes happen in non-stationary systems, for which most existing methods are unsuitable. In this article, we introduce a new enhanced sampling scheme, Non-Stationary Forward Flux Sampling, which allows efficient simulation of rare events in both stationary and non-stationary stochastic systems.

For systems in thermodynamic equilibrium, a large variety of rare-event techniques have been developed. One is the Bennett-Chandler method^{1,2}, which involves a calculation of the free energy along a predetermined reaction coordinate, followed by a computation of the kinetic pre-factor by firing off trajectories from the top of the free-energy barrier. Other techniques are transition path sampling³ and transition interface sampling (TIS)⁴, which employ Monte Carlo sampling of the ensemble of transition paths, approximate schemes such as partial-path TIS⁵ and milestoning⁶, which use a series of interfaces in phase space between the initial and final state, and string methods⁷. While these schemes have been successfully applied to a large class of problems, they do require knowledge of the phase-space density, which limits their use to systems in thermodynamic equilibrium.

For non-equilibrium systems, the phase-space density is generally not known. This severely limits the possibilities for devising enhanced sampling schemes to calculate transition rates. Yet, for non-equilibrium systems that are in stationary state, recently a number of rare-event techniques have been developed. One is the weighted-ensemble (WE) method^{8,9}, where phase space is divided into bins, and trajectories are selected and re-weighted

bin-wise to achieve uniform coverage of the phase space. Another technique is the non-equilibrium minimum action method¹⁰, which allows the characterization of transition paths but not rate constants. Non-equilibrium umbrella sampling¹¹ coarse-grains systems with Markovian dynamics on overlapping grids in state space and biases inter- vs. intra-bin transitions. Forward-Flux Sampling (FFS)¹² uses a series of interfaces in phase space between the initial and final state to drive the system over the barrier in a ratchet-like manner, by capitalizing on those fluctuations that move the system from one interface to the next. While these methods do not require thermal equilibrium, they rely on the system being in stationary state.

In reality, however, many important rare event processes happen in systems which are not in stationary state. For these processes, the propensity (probability per unit time) for the rare event to occur is time-dependent; this time dependence may be caused by external driving, by transient relaxation of the system from an out-of-equilibrium initial state, or by the presence of memory in the dynamics on a relevant time scale. In fact many real-life instances of the rare event processes mentioned above are time-dependent, such as: crystal nucleation during flash-freezing (e.g. when preparing cryo-electron microscopy samples); protein folding during transient association with a chaperone protein; and triggering of traffic jams by brief disturbances on the road. Other interesting cases include transitions between multiple limit cycles in neural networks under time-dependent stimuli (as suggested for epileptic seizures, e.g.¹³), and the response of metastable biochemical networks to transient signals, e.g. in cell differentiation¹⁴ or in viral life cycle progression, see sec. IV. These important types of rare events are not accessible to any existing enhanced-sampling techniques, with the exception of the FFS-inspired method of Berryman and Schilling¹⁵ which relies on mapping the systems dynamics onto a time-inhomogeneous Markov process. The noise-sampling method of Crooks and Chandler¹⁶ allows sampling of transition paths in non-stationary systems but cannot be used to compute time-dependent transition rates, since




Figure 1. Barrier crossing in stationary state. Equilibration within the metastable region A occurs within τ_A . After a mean waiting time τ_{AB} transitions from A to B occur, with a typical duration τ_C .

that would require knowledge of the initial phase-space density, which is unavailable in general.

In this article, we describe a new method, Non-Stationary Forward Flux Sampling (NS-FFS), that allows efficient simulation of rare events in time-dependent stochastic dynamical systems. NS-FFS constitutes a time-dependent generalization of FFS, and is conceptually straightforward and easy-to-implement. NS-FFS achieves uniform sampling of trajectories crossing a pre-defined region in time and phase space, by combining interfaces in phase space as used in FFS^{12,17} with a flat-histogram pruned-enriched Rosenbluth method originally developed for polymer simulations^{18,19}. In NS-FFS, trajectories are branched (proliferated) or pruned (terminated) based on their progression towards the final state, using interfaces in phase space *and* time. The scheme can be employed to sample the time-dependent phase-space density and time-dependent crossing fluxes, with uniform relative error. It thereby gives access to time-dependent transition rate functions²⁰, including their low-propensity tails.

The article is structured as follows. In section II we provide a theoretical background, contrasting the well-studied setting of stationary, Markovian barrier-crossing with more general time-dependent rare event problems. Section III presents the NS-FFS algorithm, together with corresponding pseudo-code. The correctness and efficiency of the algorithm are demonstrated in Section IV using two simple examples: diffusive escape in a one-dimensional W-shaped potential, and time-dependent switching in a genetic toggle switch. We conclude by discussing the main features of the method, and possible extensions, in Section V.

II. TIME-DEPENDENT RARE EVENTS

A. Transition rate constants for Markovian systems in stationary state

We first discuss rare transitions occurring in a system in stationary state. They may be visualized in terms of a

static free-energy landscape²¹, as shown in Fig. 1. A typical trajectory starts inside the metastable state A and fully explores the basin within a time τ_A ; after a waiting time τ_{AB} it makes a rapid transition (on a timescale τ_C) over a high free-energy barrier C into state B . The essential observation is that if the equilibration time τ_A is much shorter than the mean waiting time in the A basin τ_{AB} , then in the regime $\tau_A, \tau_C \ll t \ll \tau_{AB}$, transitions from A to B occur in a Markovian, memoryless fashion, effectively starting from a stationary state within A . Since the system is still in A with probability $\simeq 1$, switches also happen with a constant propensity, whose value equals the rate constant $k_{AB} = \tau_{AB}^{-1}$ (see also²⁰). Numerical techniques for simulating rare events in stationary systems exploit this fact by generating biased ensembles of *short* transition paths of duration τ_A, τ_C , which nevertheless allow reliable estimates of the much *longer* waiting time $\tau_{AB} \gg \tau_A, \tau_C$.

We briefly review how this works in FFS¹². Given a progress coordinate λ which increases from A to B , one defines a set of interfaces at successive levels $\{\lambda_l\}_{1 \leq l \leq L}$. A sample of points at the initial interface λ_1 is generated by a quasi-stationary simulation within state A . These are then used to initialize a set of trajectories which are propagated to interface λ_2 , or terminated if they re-enter A , whichever happens first. This procedure is repeated, starting from λ_2 and stopping at λ_3 or A , and so on. By propagating trajectories in segments from one λ -interface to the next, FFS capitalizes on the rare fluctuations towards the transition. The resulting transition paths are of length $\tau_C \lesssim t \ll \tau_{AB}$. This leads to efficiency gains which grow exponentially with the barrier height.

B. Time-dependent rare transition events

In this paper we are interested in situations where the metastable states A and B can still be identified, but transitions between them happen with time-dependent propensity. For example, let us suppose that the generic system illustrated in Fig. 1 and discussed above is exposed to weak external forcing with protocol $\phi(t)$, $0 < t < T$. For a macroscopic description with a time resolution coarser than τ_A , the system is macroscopically Markovian and one can still define a transition rate from A to B , but this now depends on time: $k_{AB} = k_{AB}(t)$ ²⁰.

Transition events may then be ‘uniformly rare’ so that $k_{AB}^{-1}(t) \gg T$ for all t and the survival probability $S_A(t) \simeq 1$ up to time T . However, if the transition rate is high in a particular time window, then the survival probability $S_A(t)$ may drop significantly below 1 during the time interval $(0, T)$ of interest (as in figs. 9,10 discussed below), and has to be taken into account in computing $k_{AB}(t)$ ²⁰. One then needs to measure both the first-passage time distribution or flux $q_{AB}(t)$ from A to B and the survival probability $S_A(t)$, to extract the time-dependent rate $k_{AB}(t) = q_{AB}(t)/S_A(t)$; see the accompanying paper²⁰ for a detailed discussion.

Alternatively, transitions from A to B may be time-dependent even in systems without external driving, due to “macroscopic memory”, in which the system’s dynamics evolves on a time scale τ_{slow} such that $\tau_A < \tau_{\text{slow}} < \tau_{AB}$. In this case, relaxation within the A basin will no longer be effectively instantaneous, and for $t < \tau_{\text{slow}}$ the exit propensity from A will depend on the history of the trajectory. While transitions from A to B cannot be described by a rate constant, one may be able to characterize such systems in terms of a rate kernel $k_{AB}(t|t')$, which quantifies the propensity to switch from A to B for the first time at time t , given that the previous switch happened at $t' < t$ (see²⁰ for a detailed discussion). To extract the rate kernel from a simulation, one needs to measure the probability to stay in A without interruption from t' to t ; this requires a simulation over a time interval $(0, T)$ where $T \gtrsim \tau_{\text{slow}}$. In some systems memory effects may be combined with external driving – for example barrier escape in an underdamped, driven system.

In all of these scenarios, the system dynamics is non-trivial and interesting over a time window $(0, T)$ (determined by the external driving or internal memory) which is longer than the typical transition time τ_{AB} . This fact makes it impossible to speed up the simulation by generating only an ensemble of short trajectories of length $\gtrsim \tau_A, \tau_C$, as is done in the rare event techniques for stationary systems discussed above. To capture the physical behavior of non-stationary systems, trajectories must extend over the *entire* time window $(0, T)$ of interest.

Nevertheless, enhanced sampling is both useful and feasible for non-stationary systems. Clearly, if in the time window of interest $0 < t < T$ an event occurs with low probability, then brute-force simulations will fail to generate more than a few, if any, events on this time scale. The goal of an enhanced sampling method is therefore to generate an ensemble of trajectories of full length T which is biased towards the transition. By reweighting this ensemble one can compute properties such as the time-dependent probability density, the transition flux $q_{AB}(t)$, the transition rate function $k_{AB}(t)$ or the transition rate kernel $k_{AB}(t|t')$ for the original system, over the time window $0 < t < T$ of interest. NS-FFS, presented below, is precisely such a technique: by proliferating trajectories that evolve towards the final state and terminating those that do not, it generates transition events in the time window of interest; moreover, the branching/pruning strategy is such that the relative sampling error is uniform over the space-time region of interest.

III. NON-STATIONARY FORWARD FLUX SAMPLING

The aim of the NS-FFS method is to generate a biased set of trajectories which sample transitions from state A to state B , defined by a progress coordinate λ , as a function of time, for non-stationary stochastic systems. To bias the set of trajectories towards the transition, one would like the flux of trajectories in the biased ensemble

to be uniform in λ (as in methods like FFS); to sample accurately the time-dependent behavior of the system (i.e. to obtain good sampling of early as well as late transition events), one would also like the flux of trajectories to be uniform in time. NS-FFS achieves both of these objectives, by generating a set of trajectories that uniformly cover a specified region of interest \mathcal{R} in time and in the progress coordinate. Rare excursions into low-probability regions within \mathcal{R} are sampled with the same accuracy as common events, and early excursions are sampled with the same accuracy as late ones.

The method is conceptually simple. First, one defines the region of λ -time space \mathcal{R} that is of interest. One then partitions the region \mathcal{R} using a series of interfaces; the interfaces may be defined as a level set either in the progress coordinate or in the time coordinate. Each interface is then partitioned into a set of bins; the bins are defined as intervals in the respective other coordinate (time for λ interfaces and λ for time interfaces).

The simulations start by generating an ensemble of initial conditions, for example by performing a brute-force simulation in the initial state A . One then proceeds by firing off a trajectory from a randomly chosen initial condition, and propagating it according to the given dynamical equations of the system. The trajectory is assigned a statistical weight, which is initially unity. Upon crossing one of the pre-defined interfaces, the trajectory may be branched (split into several ‘child’ trajectories, with new statistical weights), or pruned (terminated). Repeating this procedure recursively for all child trajectories, one generates a ‘trajectory tree’ which extends from 0 to T . The whole procedure is then repeated by firing off a new trajectory from a randomly chosen initial condition.

The probability that a trajectory is branched or pruned when crossing a given interface bin depends on a running histogram which monitors the weighted flux of trajectories crossing that bin. The branching/pruning rule is set up such that trajectories which arrive at a bin which has previously been under-sampled are likely to branch; those which arrive at a previously over-sampled bin are likely to be pruned. The algorithm successively approaches a steady state in which the numbers of trajectories passing through all bins on all interfaces are equal – i.e. uniform sampling is achieved both in time and in the progress coordinate.

In this section, we discuss each of the three ingredients—stochastic branching and pruning; interfaces in phase-time space; the flat-histogram rule—in more detail, relegating mathematical details to Appendixes. We then present pseudo-code of the resulting NS-FFS algorithm. Finally, we comment on the main features of the method. Again, implementation details are given in the Appendix.




Figure 2. Stochastic branching/pruning. A branching/pruning move with average child number \bar{n} proliferates or terminates branches. If $\bar{n} \geq 1$, surviving branches are decreased or increased in weight, respectively (in the diagram, line thickness represents statistical weight). The target child number \bar{n} may be an arbitrary function of the chosen coordinates q, q' , as long as the reweighting factor r satisfies Eq. 1.

A. Stochastic branching and pruning

A key element of the NS-FFS algorithm is the branching and pruning of trajectories, which allows control of the trajectory density. In a branching move, independent copies of the trajectory are created with a common history up to time t , while in a pruning move, the trajectory is terminated. The essential observation here is that one is free to branch or prune trajectories at will, as long as the statistical weights of the child trajectories are adjusted (reweighted) appropriately.

Suppose that at time t , a trajectory with statistical weight w is randomly branched into $n = 1 \dots n_{\max}$ children with probability $b(n)$, or pruned with probability $b(n=0)$ (Fig. 2). Each child branch is assigned a new weight $w' = r(n) \times w$. Clearly this branching/pruning move will be statistically unbiased only if the weight factor $r(n)$ is chosen correctly.

A necessary and sufficient condition for the combination of b and r to be correct is that weight be conserved on average over branching/pruning outcomes. That is, we may choose the branch number distribution and reweighting factor at will as long as they satisfy

$$\langle nr \rangle = \sum_{n=0}^{n_{\max}} b(n)nr(n) = 1, \text{ where } \sum_{n=0}^{n_{\max}} b(n) = 1. \quad (1)$$

This holds under very general conditions, including non-stationary system dynamics with memory and dependence of b on arbitrary parameters, as shown in app. A.

Thus we are free to adjust b to yield a desired mean branch number $\langle n \rangle \in (0, n_{\max})$, and thereby enrich or dilute the density of sample paths based on any chosen criterion, as long as we also adapt r to satisfy Eq. 1.

B. Interfaces in phase-time space

In order to extend the FFS algorithm to non-stationary systems, we trigger branching/pruning moves whenever

a trajectory crosses an interface. Since NS-FFS deals with systems whose dynamics are intrinsically time-dependent, the region of interest $\mathcal{R} = [\lambda_1, \lambda_L] \times [t_1, t_I]$ is two dimensional. To achieve uniform sampling of trajectories in both λ and time, the branching/pruning rule needs to depend on both these coordinates. To accommodate this, we define a set of interfaces as level sets in one of the coordinates (λ or t), and partition each interface into a set of bins, which are intervals along the other coordinate. The region of interest \mathcal{R} is thus covered by a two-dimensional grid of subdivisions (see Fig. 3). The interfaces are used to trigger the branching/pruning moves; the bins are used to determine the target child number \bar{n} for these moves.

The most direct generalization of FFS arises when interfaces are placed at a set of λ -levels, and subdivided into time-bins (Fig. 3a). The interface-bin grid is then given by the sets $B_{li} = \{(x, t) | \lambda(x, t) = \lambda_l, t_i \leq t < t_{i+1}\}$ where bin B_{li} refers to the i -th time bin on the l -th λ -interface. This interface arrangement will be referred to as ‘ λ -if’ (λ -interfaces). One can then measure the total probability weight that has passed through bin B_{li} , $H_{li} = \sum_{a=1}^{N_{li}} w_a$ where N_{li} is the running number of trajectories with weights $\{w_a\}$ that have reached B_{li} . The *crossing flux* per tree is

$$j_{li} = H_{li}/S, \quad (2)$$

where S is the running number of sampled trees. In computing H_{li} , one can choose either to count crossings only in the forward direction (increasing λ), or in both directions, depending on the system property of interest. When counting only forward crossings, the quantity $\langle j_{li} \rangle$ measures the forward probability flux across B_{li} ; formally

$$\langle j_{li} \rangle = \int_{t_i}^{t_{i+1}} \langle \delta(\lambda(t) - \lambda_l) \dot{\lambda} \theta(\dot{\lambda}) \rangle dt, \quad (3)$$

where θ is the Heaviside step function, and $\dot{\lambda} \theta(\dot{\lambda})$ is the forward probability flux. Although j_{li} is strictly speaking a (unitless) crossing probability, it is proportional to the forward probability flux averaged over the bin B_{li} , which justifies the name ‘crossing flux’. Importantly, when an absorbing boundary condition is imposed at the last interface λ_L which is located beyond the top of the barrier, then

$$\langle j_{Li} \rangle = \int_{t_i}^{t_{i+1}} q_{AB}(t) dt \quad (4)$$

where q_{AB} is the first-passage time probability density, or exit flux, into B . This makes the λ -if setup naturally suited for estimating exit fluxes.

Alternatively, one may interchange the roles of time and the progress coordinate, by placing interfaces at a set of time points t_i , and subdividing them into bins along the λ -direction (Fig. 3b). We call this interface arrangement ‘ t -if’ (time-interfaces). The interface-bin grid is now described by $B_{ti} = \{(x, t) | \lambda_l \leq \lambda(x, t) < \lambda_{l+1}, t = t_i\}$;




Figure 3. In the λ -if setup (a), branching/pruning events are triggered by the crossing of levels $\{\lambda_i\}$ of the progress coordinate. Since in this case, the interface crossing flux j_{li} corresponds to the flux of probability along λ , the λ -if setup is naturally suited for measuring fluxes and time-dependent transition propensities $q_{A\lambda}(t)$. In the t -if setup (b), branching/pruning is triggered at a set of times $\{t_i\}$. Here, the interface crossing flux j_{li} provides a measure of the local probability density, making the t -if setup naturally suited for sampling $p(\lambda, t)$.

bin B_{li} refers to the l -th λ bin on the i -th time interface. Crossing fluxes are still defined according to Eq. 2, but these have a different meaning in the t -if setup: j_{li} simply estimates the probability to find the system in the interval $(\lambda_l, \lambda_{l+1})$ at time t_i :

$$\langle j_{li} \rangle = \int_{\lambda_l}^{\lambda_{l+1}} \langle \delta(\lambda(t_i) - \lambda) \rangle d\lambda = \int_{\lambda_l}^{\lambda_{l+1}} p(\lambda, t_i) d\lambda. \quad (5)$$

Thus the t -if setup lends itself naturally to estimating densities or potentials of mean force.

In both the λ -if and t -if setups, the branching/pruning rules are set up to ensure uniform sampling of j_{li} . The λ -if setup leads to uniform sampling of (forward) fluxes, while the t -if setup implements uniform sampling of phase space densities. We stress however that either setup could be used to measure either quantity. The relative efficiencies of the two methods depend on the quantity used for biasing but also on more technical aspects, as discussed below.

C. Sampling with uniform error

The final ingredient in the algorithm is the rule for setting the child number probability $b(n)$ for branching/pruning moves. In NS-FFS, a target child number is set for each bin B_{li} , depending on the statistics of previous crossings of that bin; the goal of the branching rule is to sample the crossing flux j_{li} through each of the bins B_{li} with uniform relative error.

The relative error in j_{li} in an NS-FFS run may be approximated as (see app. B)

$$\frac{\langle \delta j_{li}^2 \rangle}{\langle j_{li} \rangle^2} \simeq \frac{\alpha_N}{\langle N_{li} \rangle} \left[1 + \alpha_w \frac{\langle \delta w_a^2 \rangle}{\langle w_a \rangle^2} \right]. \quad (6)$$

Here w_a denotes the (stochastic) weight of a trajectory reaching B_{li} , and averages and variances refer to an ensemble of NS-FFS runs with S trees. The constants α_N, α_w approach unity in the ‘ideal’ case where trajectories which reach B_{li} are uncorrelated. This expression shows that the error is controlled by the number N_{li} of

trajectories that cross the bin, with an extra contribution arising from the spread in their weights $\langle \delta w_a^2 \rangle / \langle w_a \rangle^2$. Thus, to obtain a uniform relative error in j_{li} (for a given total computational cost $\propto \sum N_{li}$), the branching rule needs to equalize the number of trajectories reaching each bin, while keeping the distribution of trajectory weights within each bin sharply peaked. In NS-FFS, these requirements are met by using a somewhat simplified version of the flatPERM rule¹⁹:

Algorithm 1 Branching rule for bin B_{li}

1. Calculate the target child number as $\bar{n} = w_a / j_{li}$ where w_a is the weight of the incoming trajectory, and j_{li} is the current flux estimate.
2. Set the child number probabilities

$$b(n) = \delta_{n, \lceil \bar{n} \rceil} (\bar{n} - \lfloor \bar{n} \rfloor) + \delta_{n, \lfloor \bar{n} \rfloor} (\lceil \bar{n} \rceil - \bar{n}),$$

where $\lceil \cdot \rceil$ and $\lfloor \cdot \rfloor$ denote the ceiling and floor functions, respectively.

3. Draw a child number $n \in \{\lfloor \bar{n} \rfloor, \lceil \bar{n} \rceil\}$ from b . If $n > 0$, set all child branch weights to $w' \leftarrow j_{li}$.
-

It is easily verified that the number of children is on average $\langle n \rangle = \bar{n} = w_a / j_{li}$. If $\bar{n} < 1$, pruning occurs with probability $1 - \bar{n}$. The weight w' of the children (if any) is given by the parent weight w_a multiplied by the reweighting factor $r = w' / w_a = 1 / \bar{n}$, which is independent of n . It is easy to see that the condition for unbiased statistics, Eq. 1, is satisfied by this branching/pruning rule.

The branching rule, algorithm 1, produces a uniform error in the crossing flux estimate j_{li} , since it tends to equalize counts between bins and minimize the weight variance within a bin. To see this, consider an NS-FFS simulation which is in steady state, i.e. after the crossing flux estimates j_{li} have converged to their average values $\langle j_{li} \rangle$. In this situation the branching rule assigns each trajectory leaving bin B_{li} the ‘perfect’ weight $w \rightarrow w_{li}^\infty = \langle j_{li} \rangle$. Since all child trajectories are assigned this same weight, indeed the variance of trajectory weights leaving B_{li} is minimized (in the ideal case it is zero). We also note that the weight w_{li}^∞ is equal to the system’s intrinsic

probability to cross bin B_{li} . It follows that on average exactly one trajectory per tree (with weight w_{li}^∞) will emanate from bin B_{li} – or equivalently $N_{li}/S \times j_{li} \rightarrow 1 \times \langle j_{li} \rangle$ as the simulation converges. Thus NS-FFS achieves uniform sampling of the region \mathcal{R} , with on average an equal number of trajectories crossing each bin, on each interface.

D. The NS-FFS algorithm

Combining the above ingredients, we now describe the full NS-FFS algorithm. To set up the simulation, one needs to generate initial configurations at time $t = 0$, according to an initial distribution $\rho(x_0)$ (which need not be known explicitly as a function of x_0); these could simply be generated via a brute-force simulation in the A state. Next, one identifies a suitable progress coordinate λ , and places a set of interfaces, subdivided into bins $\{B_{li}\}$, over the region \mathcal{R} of phase-time space of interest, according to either the λ -if or the t -if setup. One also specifies a dynamic range for reweighting, by setting minimum and maximum trajectory weights (w_{\min}, w_{\max}) (these enhance convergence, see below). The simulation then proceeds according to algorithm 2:

Algorithm 2 NS-FFS

Init. Define a histogram of total crossing weights H_{li} and set $H_{li} \leftarrow 0$ for all l, i . Set a tree counter $S \leftarrow 0$. Define a queue of pending trajectories G and set $G \leftarrow \{\}$.

Run Iterate the following steps, until the desired accuracy is reached:

1. Start a new trajectory at $t = 0$, from a new initial state x_0 . Insert the trajectory into G , and assign an initial weight $w \leftarrow 1$. Increment the tree counter, $S \leftarrow S + 1$.
 2. While G is non-empty, iterate:
 - (a) Pick and remove a trajectory from G .
 - (b) Propagate the trajectory forward in time while recording any observables of interest, weighted with w , until either the final time T is reached, or an interface is crossed. In the latter case:
 - i. Determine the bin B_{li} which was crossed and increment H_{li} by the current trajectory weight w .
 - ii. If $w \in (w_{\min}, w_{\max})$, carry out a branching/pruning move: draw a child number n and set the child weight w' according to algorithm 1. Otherwise, set $n \leftarrow 1$ and $w' \leftarrow w$.
 - iii. Generate n child trajectories and insert them into the queue G for further propagation.
-

When setting up the region of interest \mathcal{R} , clearly the transition region should be covered to enhance transition

paths. It is equally important to let \mathcal{R} extend well into the metastable basins; this allows for pruning of trajectories which would otherwise accumulate in the basins, degrading performance. Note also that by default, trajectories are not terminated when they leave \mathcal{R} before they reach the final time T , ensuring that they may re-enter \mathcal{R} and contribute at a later time. It is of course possible to explicitly add absorbing boundaries, which may increase performance, in cases where later reentry is not required.

The weight limits (w_{\min}, w_{\max}) which appear in Alg. 2 are not strictly necessary for the existence of a steady state with uniform sampling of the crossing flux j_{li} . They do, however, greatly enhance convergence towards it. In the initial phase of an NS-FFS simulation, the weight histogram H_{li} is sparsely populated such that the flux estimates j_{li} are subject to large fluctuations. This can result in avalanches of correlated low-weight paths which degrade performance. For the method to be useful, it is necessary to have an effective way of controlling these bursts. Among a number of possible remedies including negative feedback control of tree size, branching thresholds, or explicit flat-histogram branching based on number densities¹⁹, we found weight limits to be particularly simple and very effective. In practice, the reasonable rule-of-thumb to choose $w_{\min} \lesssim \min_{l,i} \{\langle j_{li} \rangle\}$ and $w_{\max} \gtrsim 1$ was found to work well.

The output of algorithm 2 consists of weighted trees of trajectories in which all trajectory segments end either at an interface (branching/pruning point) or at the final time T (completion) (Note that a simulation may be stopped only after a full tree is finished). Trees may be generated depth-first (children before sisters), or breadth-first (sisters before children), depending on whether the queue G in Alg. 2 is of the last-in-first-out or first-in-first-out type. We found no significant difference in performance between depth-first and breadth-first traversal, in contrast to other reports for the case of PERM²². Trajectory trees may also be written to disk in a recursive data structure for offline analysis.

IV. APPLICATIONS

A. Crossing of a linear barrier

As a simple model for rare barrier crossing events, we consider overdamped Brownian motion in a linear double-ramp potential $U = -a|x|$ with boundaries at $x = \pm 1$ and barrier height $a > 0$, see Fig. 4. We consider two systems, which differ in their boundary conditions: one with reflecting boundaries at both $x = -1$ and $x = 1$, meaning that probability is conserved; another one with a reflecting boundary condition at $x = -1$ and an absorbing one at $x = 1$. The two systems will be referred to as the reflecting/reflecting or reflecting/absorbing systems, respectively. This choice of model potential allows for an exact calculation of the Green's function, so that the correctness of our simulations can be assessed directly.




Figure 4. Linear ramp potential U (left) and the first 500 branches of path trees generated during NS-FFS simulations using the t -if setup (top, reflecting boundaries) and the λ -if setup (bottom, reflecting/absorbing boundaries). Although the total simulated time that is shown, is $\tau_{sim} \lesssim 6 \ll \tau_{rxn}$, crossing events are successfully generated. Bins B_{li} are depicted in blue, branch weight is indicated by line shading and branch points are shown as circles.

Particles are injected at $x = -1$, $t = 0$ and diffuse according to the overdamped Langevin equation

$$\dot{x} = -D\partial_x U + \xi. \quad (7)$$

We use thermal units where $k_B T = 1$, so that ξ represents Gaussian white noise with $\langle \xi(t) \rangle = 0$, $\langle \xi(t)\xi(t') \rangle = 2D\delta(t - t')$. In these units, the diffusion constant D is equal to the mobility coefficient, and a has units of inverse length.

In this barrier escape problem both the time scale τ_A for equilibration in region A for $x \gtrsim -1$ and the crossing time τ_C , are much faster than the waiting time between crossings $\tau_{AB} = \tau_{BA} = 2\tau_{rxn}$ where τ_{rxn} is the global relaxation time of the system. The Laplace-transformed probability density $p(x, s) = \int_0^\infty e^{-st} p(x, t) dt$ can be calculated exactly using standard methods, see app. D, leading to explicit expressions for the time scales

$$\tau_A = \frac{4}{a^2 D}, \tau_C = \frac{2}{aD}, \text{ and } \tau_{rxn} = \frac{2e^a}{a^2 D}. \quad (8)$$

In the present examples, the barrier height and diffusion constant were set to $a = 15$ and $D = 1$, respectively, such that $(\tau_A, \tau_C, \tau_{rxn}) = (0.02, 0.13, 2.9 \times 10^4)$, respectively. Moreover, the full probability density $p(x, t)$, and for the reflecting/absorbing case, the exit propensity $j(x, t)$, could be obtained as functions of time using an efficient numerical contour integration method²³.

Fig. 4 depicts the potential, the interface bins, and partial trajectory trees taken from NS-FFS simulations

of this system from $t = 0$ to $t = T = 1$. The region of interest was taken to be $\mathcal{R} = [-1, 1] \times [0, 1]$ in (x, t) ; the progress coordinate λ was defined trivially as $\lambda \equiv x$.

We first study the probability density $p(\lambda, t)$ in the reflecting/reflecting system, which relaxes towards the Boltzmann distribution $\propto e^{-U}$ over times longer than τ_{rxn} . The system was simulated using the t -if setup which is most natural for measuring $p(\lambda, t)$, since the branching factors are controlled by the local density. $I = 199$ interfaces were placed across \mathcal{R} , at regular time intervals (from $t = 0$ to 1), and partitioned into $L = 40$ equal-sized bins each (from $\lambda = -1$ to 1). In Fig. 5 we compare $p(\lambda, t)$ as obtained as via exact analytical calculation, brute-force and NS-FFS simulation. While the probability density p varies over 8 decades within the region \mathcal{R} , the sampling density in NS-FFS is constant within a factor of 5 (Fig. 5c). After reweighting, p is correctly reproduced throughout \mathcal{R} (even where it is very small; Fig. 5d). This was achieved within a simulation time that would generate only a handful of transition paths in a brute force simulation (Fig. 5b). The region around the cusp of the potential is somewhat under-sampled. This is due to the fact that the force in our model is discontinuous at $x = 0$; to achieve complete sampling uniformity in this region, one would require a bin width on the order of the length scale of variation of the potential.

We next consider the time-dependent exit flux through the absorbing boundary in the reflecting/absorbing system. For this NS-FFS calculation, we used the λ -if setup; this setup is natural since here the sampling bias is based on crossing fluxes over the λ -interfaces, and the last crossing flux coincides with the observable of interest. $L = 19$ λ -interfaces were placed at regular intervals and partitioned into $I = 50$ time-bins each; only crossings in the positive λ -direction counted towards H_{li} . Fig. 6 shows the probability density $p(\lambda, t)$ computed using the exact solution (a) and NS-FFS (c) as well as the unweighted crossing fluxes and the reweighted exit flux (d) for this system.

As expected, the number fluxes of trajectories emanating from each bin in the positive lambda direction (d, right axis) superimpose in a narrow band, confirming that NS-FFS indeed produces uniform sampling across \mathcal{R} . Only the earliest bins at the farthest interfaces are visited less often, since the system dynamics does not allow them to be reached in the required time with sufficient probability.

In contrast, the number density of trajectories over λ, t (b) is roughly uniform over \mathcal{R} , but does exhibit a systematic bias favoring the A state. This can be understood as follows. The branching rule in the chosen λ -if setup biases towards uniform number flux in the forward λ -direction, not uniform density. While in the region $x > 0$ trajectories spontaneously move in positive flux direction, in the region $x < 0$, trajectories tend to drift downhill, against the positive flux direction. In this region, the λ -if NS-FFS simulation maintains a uniform population of uphill trajectories by proliferating. The to-




Figure 5. Time-dependent density $p(\lambda, t)$ for the linear ramp potential with initial condition $p(\lambda, 0) = \delta(\lambda + 1)$, with reflecting boundary conditions. (a) numerically exact solution; (b) brute force sampling. In the t -if NS-FFS run, raw counts (c) cover \mathcal{R} almost uniformly while weighted counts (d) reproduce the exact density. Taking the B state boundary to be $q_B = 0.5$, the occupation of the B state (d) fits a linear growth model with slope $k_{AB} = (3.41 \pm .03) \times 10^{-5}$ and delay $0.099 \pm .003$ (blue, as extracted from NS-FFS; black, exact solution). Total simulated time was $10^5 \simeq 3.4\tau_{\text{rxn}}$. Counts refer to the histogram bins of size $(\Delta x, \Delta t) = (.05, .01)$ used in this figure.

tal (uphill+downhill) trajectory density is thus increased in the region $x < 0$.

This observation clearly demonstrates the difference in sampling biases between the t -if and λ -if setups: the former generates a uniform density while the latter generates uniform fluxes in the λ -direction. Nevertheless, as Fig. 6b shows, in practice the methods provide almost uniform sampling of both forward flux and total number density, so that it is certainly possible to sample fluxes using the t -if setup and densities with the λ -if setup, without dramatic loss in performance.




Figure 6. Time-dependent probability density (as in Fig. 5) and exit flux, but for reflecting/absorbing boundary conditions. The exact density p is shown in (a). Counts generated in a λ -if NS-FFS run show roughly uniform sampling (b, see text). Weighted counts (c) reproduce the exact density over the full dynamic range. Unweighted number fluxes of trajectories emanating from each bin (d, right axis) for all 19 λ -interfaces (gray value increasing with λ) collapse, indicating uniform number flux. As a consequence the exit probability flux j over $x = 1$ (d, left axis) is sampled uniformly including its low-probability onset (black, exact solution; blue, NS-FFS). An estimate of the stationary exit flux from these data for $t > .24$ gives $j = 3.424(\pm 0.02) \times 10^{-5}$; the exact value is $j_\infty = \tau_{AB}^{-1} = 3.44 \times 10^{-5}$.

In this context it is worth noting that depending on the observable to be estimated, there is the additional freedom of setting up the simulation to measure either density or exit flux. For instance, the slope in Fig. 5(e) and the plateau value of the exit flux in Fig. 6(d) coincide, even though trajectories may re-cross the barrier from B to A in the reflecting/reflecting system, but not in the reflecting/absorbing system. This is in agreement with the exact expressions in the limit $t \ll \tau_{\text{rxn}}$; indeed, as long as $t \ll \tau_{\text{rxn}}$, the occupation of the B state is negligibly small so that back-crossings occur with a probability even much smaller than forward crossings. Effectively, the B state initially appears as absorbing even in the purely reflecting system. Thus, either simulation may be used to measure the rate constant for this system.

The relative errors in the estimated probability den-




Figure 7. Relative errors corresponding to the probability densities shown in figs. 5 and 6 for the t -if and λ -if (top and bottom, respectively). Relative errors are computed as absolute differences between sampled and exact densities, normalized by the exact density.


sity $|\Delta p|/p = |p_{\text{sim}}(\lambda, t) - p_{\text{exact}}(\lambda, t)|/p_{\text{exact}}(\lambda, t)$ shown in Fig. 7, further illustrate the uniform sampling over \mathcal{R} generated by NS-FFS. The relative error scatters uniformly over \mathcal{R} and in particular, does not scale with $p^{-1/2}$ as would be the case for brute-force sampling. Larger errors remain only in the fringes of the accessible region, where fewer trajectories are sampled. The residual stripe pattern in the t -if case carries the signature of correlated trajectories originating at the cusp of the potential; the undersampling right at the cusp which causes this can be considered a pathological feature of the force discontinuity in our model. Notice that although the λ -if setup equalizes positive fluxes rather than number density, the probability density of trajectories nevertheless exhibits uniform error in this example, despite the weak asymmetry across \mathcal{R} in the number of trajectories sampled, visible in Fig. 6b.

B. Genetic toggle switch

As an intrinsically non-equilibrium example, we next consider a bistable gene regulatory network which can be seen as a simplified version of the λ -phage genetic switch. This ‘toggle switch’ consists of two genes that mutually repress each other. In the ‘exclusive’ variant considered here, the two genes, which produce proteins A and B respectively, share a common DNA operator region O, such that when the dimer A_2 is bound to the operator, protein B cannot be produced, and vice versa. This model is discussed in detail in refs^{24,25}.

In this simplified model, production of proteins is represented by a Poisson process. The model consists of the

symmetric reaction set:



The fact that only protein dimers may bind to the operator site, and the fact that the state OA_2B_2 is disallowed, together make this system a robust bistable switch²⁴. Each metastable state is characterized by an abundance of only one species, and transitions between these states occur on a much longer timescale than relaxation within them. We use the same rate constants as in¹²: $\mu = k/4$, $k_f = k_b = 5k$, $k_{\text{on}} = 5k$ and $k_{\text{off}} = k$, and we measure time in units of k^{-1} .

A natural progress coordinate for this non-equilibrium system is given by the difference in total monomer numbers,

$$\lambda = n_B + 2n_{\text{OB}_2} + 2n_{\text{B}_2} - (n_A + 2n_{\text{OA}_2} + 2n_{\text{A}_2}).$$

We are interested in a region $\mathcal{R} = \{(\lambda, t) \in (-40, 40) \times (0, 10^3)\}$ which spans both metastable states and the transition region. Using the t -if setup, we define $I = 500$ equidistant time interfaces. $L = 16$ λ -bins were defined with boundaries at $\lambda = \pm\{40, 24, 22, 18, 15, 12, 9, 4\}$ (These are the interface locations used in¹², augmented by bins in the basins A, B .) We take the metastable states as $A = \{(x, t) | \lambda < \lambda_A\}$ and $B = \{(x, t) | \lambda > \lambda_B\}$ where $-\lambda_A = \lambda_B = 24$.

1. Unbiased relaxation

Fig. 8 shows the result of an NS-FFS simulation of this model genetic switch from $t = 0$ to $T = 10^3$, with initial molecule numbers fixed to 0 except $n_A = n_{\text{A}_2} = 10$, so that $\lambda(0) = -30 \in A$. As in the previous one-dimensional example, the region of interest is sampled approximately uniformly in NS-FFS. Measuring the occupancy of the B state $\langle h_B(t) \rangle = \langle \theta(\lambda(t) - \lambda_B) \rangle$, and fitting to a delayed linear rise $(t - \tau_{\text{lag}})/\tau_{\text{AB}}$, we obtain a lag time $\tau_{\text{lag}} = (129 \pm 5)$, and recover a waiting time for barrier crossing $\tau_{\text{AB}} = (1.07 \pm 0.01) \times 10^6$ in accordance with the previously measured value $1/k_{\text{AB}} = (1.06 \pm 0.02) \times 10^6$ ¹².

2. Response to time-dependent forcing

We now consider the reaction of the toggle switch to a time-dependent external bias. This case is inspired by the phage- λ switch in the bacterium *Escherichia coli*, where an increase in intracellular RecA concentration triggers the transition from the lysogenic to the lytic phase of




Figure 8. Bin counts (a), probability density $p(\lambda, t)$ (b), and cumulative crossing probability $\langle h_B(t) \rangle$ (c) for the toggle switch, simulated using NS-FFS with the t -if setup. A linear fit for $t > 250$ is shown in orange. Error bars were generated by bootstrap resampling from 25 independent simulation runs with total simulated time 10^6 each.

the virus life cycle²⁶. As a simplified model for the action of RecA we introduce a species R which degrades A monomers:



The degradation of A by R forces the switch towards the B state; thus R can be regarded as an ‘external force’ acting on the switch, whose strength can be measured by the steady-state bias γn_R^∞ , where $n_R^\infty = k_R/\mu_R$ is the number of molecules of R in steady state. Relaxation of n_R towards n_R^∞ is exponential with a relaxation time μ_R^{-1} .

First, we initialize the switch in state A with $n_R = 0$ copies of R and switch on the production of R. The steady state bias γn_R^∞ is chosen such that in steady state the switch is fully driven to the B state. The switch then flips from A to B with a distribution of switching times. Switching events result from favorable fluctuations in the copy numbers of the molecules that constitute the switch (eqs. 9). These are partly due to the intrinsic stochastic nature of the switch, and partly induced by (extrinsic) fluctuations in the number of biasing molecules R (eqs. 10). The distribution of switching times thus reflects both intrinsic noise of the switch and extrinsic noise originating from fluctuations in the number of R molecules. We investigated these effects by using NS-FFS to obtain the switching dynamics as a function of the level of noise in the bias. To modulate

the latter, we varied the equilibrium copy number of R between $n_R^\infty = 1$ and $n_R^\infty = 100$, while keeping the average bias γn_R^∞ and bias time constant $\mu_R^{-1} \equiv 500$ fixed. Therefore, on a mean-field level, all bias protocols were kept the same. However, the individual bias trajectories $n_R(t)$ are markedly different: At $n_R^\infty = 100$, each trajectory $n_R(t)$ exhibits a nearly deterministic and exponential rise in time, while at $n_R^\infty = 1$, each individual trajectory $n_R(t)$ is a single off-on event, which is exponentially distributed in time. Fig. 9a shows the mean exponential rise of the bias protocol and the level of fluctuations around it. Note that in a linear system where the state occupancies are linear functions of the external forcing history, this variation of parameters would lead to a probability of having switched after the pulse which is independent of the pulse duration μ_R^{-1} .

Fig. 9(b-e) illustrates the switch response to a bias at various noise levels. For a nearly deterministic bias, with $n_R^\infty = 100$, the switch response is characterized by a gradual increase of λ towards a threshold, followed by a transition over the threshold (Fig. 9b). The transition times have a relatively narrow distribution around $t = 500$. Since further increasing the expression level n_R^∞ does not sharpen the switching time distribution (not shown), the result shown in panel b corresponds to the intrinsic limit in precision of the toggle switch at the given rate of biasing μ_R .

As the driving becomes noisier (Fig. 9c,d), the switch response exhibits a broader distribution of switching times, with both early and late crossings, and in addition, a weak new metastable state develops at the transition state (d). Fig. 9e summarizes these results, showing the probability that the switch has flipped, as a function of time, for several different values of n_R^∞ . It is clear that noise in the driving force has an important effect on the switching trajectories.

The qualitative change in the switch flipping trajectories for low values of n_R^∞ can be understood by a picture in which the switch dynamics are enslaved to fluctuations in R. In the extreme case $n_R^\infty = 1$ (Fig. 9d), individual bias trajectories switch to full bias strength suddenly, at random times. The switch then responds to the strong bias by rapid flipping to the B state in a stereotyped way; the time course of the switching event (not shown) includes a pause at the threshold which is responsible for the zone of higher density around $\lambda = 0$ (This pause originates from the fact that for $n_R^\infty = 1$, the A molecules are very rapidly degraded when the single R molecule becomes present, while it still takes time to produce the B molecules, which ultimately flip the switch).

We now consider the switch response to transient pulses of biasing molecules R. In these simulations, after the system reaches a quasi-stationary state in region A, $n_R^{\text{pulse}} = 10^2$ biasing molecules are flushed in instantaneously. We set $k_R = 0$, so n_R then decays to 0 over a pulse duration μ_R^{-1} . During the pulse, the switch is biased towards the B state. To isolate the effect of different pulse durations on the switch re-




Figure 9. Response of the toggle switch to a bias of R molecules, whose production begins at time $t = 0$ (Eq. 10). The mean bias strength approaches $\gamma n_R^\infty = k$ (a, lines) but with different amounts of noise, corresponding to different choices for n_R^∞ (a, shaded areas lie between the 5th and 95th percentiles). The switch response becomes more random as the noise in the bias increases (n_R^∞ decreases; b-d, respectively), and the transition time distribution widens (e, error bars indicate the 5th and 95th percentiles for the B state occupation $\langle h_B(t) \rangle$).

sponse, we adjust γ such that the integrated bias strength $\langle \int \gamma n_R(t) dt \rangle \equiv 1$ remains fixed, while changing the value of μ_R ($\mu_R^{-1} = 100, 1, 0.1$). This leads to $\gamma = \mu_R / n_R^{\text{pulse}}$.

Fig. 10 shows that the toggle switch possesses an optimal pulse duration, even though the integrated bias strength remains constant. For moderately long pulses, the occupation of the B -state after $t = 500$, increases with decreasing pulse duration. This can be understood in terms of a non-linear threshold behavior of the switch. As the pulse duration μ_R^{-1} is decreased, the initial bias strength $\gamma n_R^{\text{pulse}} = \mu_R$ increases, enhancing the switching probability. However, as the pulses become even shorter the switching probability decreases again. This decrease




Figure 10. NS-FFS results for the toggle switch, under the influence of pulses of biasing molecules R of varying duration μ_R^{-1} but constant total efficiency $\langle \int \gamma n_R(t) dt \rangle \equiv 1$. The time-dependent bias $\langle \gamma n_R \rangle$ is shown in (a). The response of the system to pulses of durations $\mu_R^{-1} = \{100, 1, 0.1\}$ is shown in (b,c,d), respectively. The crossing probability, taken to be the B state occupation $\langle h_B(t = 500) \rangle$ exhibits a maximum at a pulse duration around $\mu_R^{-1} = 1$ (e)

is a dynamical effect: since the switch cannot respond to changes in n_R which occur on a timescale shorter than its own intrinsic kinetic time scale k^{-1} , it acts as a low-pass filter. In this sense the toggle switch can be said to be robust against both strong transient perturbations and persistent weak perturbations. These results clearly show that the response of a genetic switch to a perturbation (or signal) is a dynamic property which depends not only on the (integrated or peak) pulse strength of the perturbation but also on its pulse shape.

V. DISCUSSION

A. General features

The NS-FFS scheme has a number of characteristic features. First, like FFS, NS-FFS does not perturb the given dynamical equations of the system (i.e. no bias-

ing force is applied). Instead, NS-FFS generates a biased ensemble of unbiased trajectories by proliferating those that move in a preferred direction and terminating those that do not, with appropriate reweighting. This means that NS-FFS, like FFS, is suitable for systems for which the dynamical equations cannot easily be biased and reweighted (e.g. because they do not obey detailed balance). It also makes NS-FFS highly suitable for implementation as a wrapper around existing simulation code.

Second, NS-FFS generates trajectories with a complete history from $t = 0$ on. This means that no assumption of memory loss is made when a trajectory passes between interfaces, as in some other rare event simulation methods^{5,6}. Perhaps more significantly, NS-FFS also does not assume loss of memory on reentry to the A state. Most existing rare-event methods for stationary systems, such as TPS, TIS, and FFS^{3,4,12}, rely on the assumption that trajectories which re-enter A equilibrate rapidly and can be treated as new, independent trajectories when they eventually re-exit A . For most stationary systems where $\tau_A \ll \tau_{AB}$ this is a reasonable assumption, but for non-stationary systems, such as those with non-Markovian or time-inhomogeneous macroscopic switching dynamics²⁰, correlated entrance and re-exit from a basin can make a significant contribution to the time-dependent quantity of interest. Thus it is an essential feature of NS-FFS that memory loss is not assumed for $t > 0$, even if the system re-enters the A basin.²⁷

Third, once the NS-FFS algorithm has reached its steady state, each interface bin emits one trajectory per started tree on average, so that trajectories are sampled uniformly over the range \mathcal{R} (see figs. 5, 6, 8). Some fluctuations around this uniform sampling average are tolerated in exchange for narrow weight distributions at each bin. These features are a direct consequence of the basic flatPERM branching rule¹⁹, which does not implement a negative feedback on (unweighted) trajectory numbers. NS-FFS is thus a ‘weak’ flat-histogram method.

Fourth, we note that the effectiveness of a simulation scheme depends not only on its ability to generate many samples but also on the independence of these samples. Clearly, in NS-FFS, after a branching event, child trajectories remain correlated for a certain time. This suggests that one should allow further branching of the children only after a refractory time of the order of the typical decorrelation time. We did not observe this modification to produce any significant improvement for the systems studied here. This is presumably because these systems are sufficiently stochastic that branched trajectories anyway decorrelate rapidly between interfaces. In contrast, it turned out to be crucial to control the exponential growth of trajectory trees in the initial phase of a simulation since these generate highly correlated samples, which delay convergence of the crossing weight histogram. This is simply and efficiently accomplished by using weight limits, as described above.

B. Progress coordinate vs. time based branching

The λ -if and t -if interface setups (sec. IIIB) differ in that the former equalizes number fluxes of trajectories in λ -direction across the region \mathcal{R} while the latter equalizes their number density. Nevertheless, it is of course possible to use either scheme to measure any quantity in a given physical system; the choice of setup will affect only the efficiency of the calculation. We now briefly discuss how the observable of interest and the computational overhead associated with branching can affect the choice of the most suitable setup.

Target observable Suppose one wishes to measure the time-dependent propensity $k_{AB}(t)$ for exit over some final level λ_L of the progress coordinate at the boundary of state B , over some time interval $[t_1, t_I]$ (cf. Fig. 1). If re-crossings back from B can be safely neglected, we may place an absorbing boundary at λ_L . It is then natural to use the λ -if setup, placing λ -interfaces at levels $\{\lambda_l\}_{l=1..L}$, and count forward crossings only. In the limit of slow escape over λ_L such that the survival probability $\int p(\lambda, t) d\lambda \simeq 1$ over the time of observation, the observable k_{AB} coincides with the positive flux over the last interface λ_L ²⁰. Since the relative error in the positive flux is equalized over all preceding interfaces, the λ -if setup will generate uniform sampling for $k_{AB}(t)$ at the final interface over the time interval of interest.

Alternatively, one may be interested in a potential of mean force $-\log p(\lambda, t)$ over a region $\mathcal{R} = [\lambda_1, \lambda_L] \times [t_1, t_I]$ in λ - t space. In that case the t -if setup is the more natural choice, since it generates a uniform relative error in $p(\lambda, t)$ over all bins. One then obtains an estimate of $-\log p$ with uniform absolute error over the region \mathcal{R} . In particular, saddle points and basins are sampled with equal frequency.

Branching overhead The λ -if and t -if setups differ in the relative overhead of coordinate evaluation and branching. Clearly, the detection of crossings requires the evaluation of λ , which incurs some computational overhead o_λ per evaluation. The branching move itself and the maintenance of the tree structure add a second type of overhead o_b , which we also take to be constant per branching event.

In the t -if setup, λ evaluation and branching are coupled and happen once every time-interface spacing τ_t ; the overhead is $(o_\lambda + o_b)/\tau_t$ per unit simulated time. In contrast, in the λ -if setup, they are decoupled: interface crossings are checked at intervals τ_λ while the trajectory is being propagated, whereas trajectory branching happens only if a crossing is detected. The λ -evaluation interval τ_λ is an independent adjustable parameter (but must not be too long or the algorithm will fail to detect crossings). If interface crossings happen every τ_b on average, the overhead is $o_\lambda/\tau_\lambda + o_b/\tau_b$ per simulated time.

If λ evaluation is expensive ($o_\lambda \gg o_b$), it is useful to increase τ_λ or τ_t , respectively, as far as possible without degrading the weight statistics. In the opposite case that λ evaluation is cheap ($o_\lambda \ll o_b$), the λ -if setup may be

advantageous, since this setup gives the option to check λ often ($\tau_\lambda \ll \tau_b$), at small cost. Especially in systems where excursions towards the B state tend to be short-lived, it is advantageous to check for interface crossings significantly more often than they actually occur, since this increases the chance of detecting and capitalizing on short-lived forward excursions. The dependence of NS-FFS performance on parameters will be addressed more fully in a future publication.

C. Variants and extensions of the algorithm

Relation to stationary FFS If NS-FFS is used to simulate a system which is in stationary state, early and late barrier crossings are equivalent; crossing statistics can then be improved by binning early and late crossings together. This can be achieved within the λ -if setup: trajectories are started from a stationary distribution at $t = 0$, and a single time bin is defined ($I = 1$), leading to quicker convergence of the, now one-dimensional, crossing weight histogram $H = \{H_{l1}\}_{1 \leq l \leq L}$. This stationary version of NS-FFS is similar but not identical to the ‘branched growth’ variant of conventional FFS¹⁷. In branched-growth FFS, trees have a fixed number of children at each λ interface crossing and are terminated exclusively at the next interface or when the system returns to A ; the stationary NS-FFS can be seen as an adaptive generalization of this scheme.

Multiple progress coordinates In problems with multiple alternative transition pathways, finding a single good progress coordinate can be challenging. Use of a poor progress coordinate fails to distinguish between trajectories which are likely and unlikely to result in a transition, making successful biasing of transition paths impossible. To some degree, NS-FFS already alleviates this problem compared to TIS or FFS: the extra dimension of time acts as a second progress coordinate which allows us to discriminate between early and late crossings. In a class of systems which show distinct ‘slow’ and ‘fast’ pathways of a reaction between A and B , an NS-FFS simulation would be able to separately enhance these pathways.

If different transition pathways share a common time scale then time as an additional progress coordinate is not useful in itself, but one may be able to find a small set of progress coordinates which successfully separate different pathways. In this case, one can make a straightforward generalization of the NS-FFS scheme to multiple progress coordinates. As in any multi-dimensional scheme, this will come at the cost of more bookkeeping and possibly slower convergence of the weight histogram H due to the increased number of bins, see app. C.

Parallel version The NS-FFS algorithm lends itself to a parallel implementation. Each trajectory can be simulated independently until an interface is crossed. At this point, the shared histogram H_{li} is read and updated. After branching, n parallel simulations for the children are spawned. The communication between simulation pro-

cesses is restricted to histogram updates; depending on available bandwidth these updates could also be cached and applied in groups, without biasing the sampling. As the global histogram converges, updates become unnecessary and the simulation gradually becomes trivially parallel.

Adaptive generalizations As shown in app. A, branching/pruning events may be introduced at will as long as weight is conserved on average (Eq. 1). This includes complete freedom of: adaptive updates of the bin boundaries or the interface placement; inserting or removing interfaces; smoothing of bin counts within or across interfaces; or pre-filling of crossing histograms based on prior knowledge.

All of these options should allow for further performance improvement in particular situations, to be explored in future work. In particular, it is interesting to ask if an optimal interface and bin arrangement can be found iteratively, as has been proposed for FFS²⁸. A promising direction might be to monitor the local sampling noise and adapt the interface arrangement in response to it.

VI. CONCLUSION

In this article, we have introduced an enhanced sampling scheme, called NS-FFS, which is conceptually simple allows the efficient sampling of rare events in non-Markovian and non-stationary systems.

The NS-FFS algorithm builds on two widely used ingredients: a flat-histogram branched growth algorithm closely related to PERM^{18,19}, and the concept of phase-space interfaces⁴ to monitor progress towards a transition. NS-FFS is a generalization of FFS¹², and is straightforward to implement, especially when one does not want to store the trajectories themselves. We have demonstrated the correctness of the method, and given several simple example applications which highlight both the effectiveness of the method and the relevance of intrinsically time-dependent rare events.

A host of physical, chemical and biophysical problems are amenable to NS-FFS simulations. These include the computation of time-dependent transition rates in systems with time-dependent external driving²⁰ such as the signal-induced flipping of genetic switches studied here, crystal nucleation during a temperature quench or protein unfolding under force; non-exponential switching time distributions in processes that can be coarse-grained as switches with memory, such as the switching of the bacterial flagellar motor^{20,29}; and escape probabilities from a non-equilibrium distribution in a metastable initial state within a prescribed ‘window of opportunity’, like the flipping of genetic switches induced by transient pulses.

Acknowledgments

We thank Daan Frenkel and Peter Bolhuis for many useful discussions. This work is part of the research program of the “Stichting voor Fundamenteel Onderzoek der Materie (FOM)”, which is financially supported by the “Nederlandse organisatie voor Wetenschappelijk Onderzoek (NWO)”. RJA is supported by a Royal Society University Research Fellowship and by EPSRC under grant number EP/I030298/1.

Appendix A: Reweighting criterion

In this section we sketch a proof that the condition of weight conservation on average over branching outcomes, Eq. 1, is sufficient for unbiased sampling in a general branched-tree simulation scheme. This amounts essentially to careful bookkeeping. No assumptions are made about stationarity or loss of memory.

1. Statement of the problem

Consider a stochastic process $x(t)$ which is started at $t = 0$ with a value of $x(0)$ drawn from some initial distribution. The process is fully characterized by all of its m -point joint probability density functions (pdfs)

$$p(x_m, t_m; \dots; x_1, t_1) = \langle \delta(x_{t_m} - x_m) \cdots \delta(x_{t_1} - x_1) \rangle, \quad (\text{A1})$$

for the trajectory to pass by the sequence of sample points x_1, x_2, \dots, x_m at times $0 \leq t_1 < t_2 < \dots < t_m$. Here x_{t_m} is a shorthand for $x(t_m)$. In a brute-force simulation, the joint pdf is estimated by an average over S independent runs,

$$p(x_m, t_m; \dots; x_1, t_1) \simeq \frac{1}{S} \sum_{s=1}^S \delta(x_{t_m}^s - x_m) \cdots \delta(x_{t_1}^s - x_1). \quad (\text{A2})$$

(we note that the right hand side of this equation is singular; the approximate equality is implied when Eq. A2 is integrated over finite regions).

We now introduce a single branching event at an intermediate time t' , and let $m' = \max\{m | t_m < t'\}$. Upon branching, n' statistically identical copies of the system with independent futures are generated. That is, for a given history up to time t' , each of the copies has the same conditional pdf $p(x_m, t_m; \dots; x_{m'+1}, t_{m'+1} | x_{m'}, t_{m'}; \dots; x_1, t_1)$ to visit future points of phase space, but future points of different children are mutually independent. No Markov assumption about the system is being made. The task is now to assign correct weights to the child branches in order to guarantee unbiased sampling.

2. Unbiased sampling with one branching event

An obvious choice for the weights is to conserve the total in- and outgoing weight at the branching point and to treat children equally. This rule gives a weight $1/n'$ for each child trajectory from t' on.

However, this rule disallows pruning, i.e. $n'=0$. To enable pruning, it is necessary to relax strict weight conservation at the branch point. To do this, child numbers $n' = 0, 1, \dots, n_{\max}$ are drawn at random with probabilities $b(n')$, $\sum_{n'=0}^{n_{\max}} b(n') = 1$. We then assign weights w'_c to the child branches $c = 1 \dots n'$ (if any). The *expected* total weight of trajectories passing through a sequence of points is then

$$\begin{aligned} W(x_m, t_m; \dots; x_1, t_1) &= \left\langle \sum_{c=1}^{n'} w'_c \delta(x_{t_m}^c - x_m) \cdots \delta(x_{t_{m'+1}}^c - x_{m'+1}) \right. \\ &\quad \left. \times \delta(x_{t_{m'}} - x_{m'}) \cdots \delta(x_{t_1} - x_1) \right\rangle \\ &= \left\langle \sum_{c=1}^{n'} w'_c \right\rangle \langle \delta(x_{t_m} - x_m) \cdots \delta(x_{t_1} - x_1) \rangle \\ &= \left\langle \sum_{c=1}^{n'} w'_c \right\rangle p(x_m, t_m; \dots; x_1, t_1), \quad (\text{A3}) \end{aligned}$$

where m' is defined as above, sums running from $n' = 1$ to $n' = 0$ vanish by definition, and we used the fact that children are identical. Clearly the condition of weight conservation on average,

$$1 = \left\langle \sum_{c=1}^{n'} w'_c \right\rangle = \sum_{n'=1}^{n_{\max}} b(n') \sum_{c=1}^{n'} w'_c, \quad (\text{A4})$$

is necessary and sufficient for unbiased sampling, since then $W \equiv p$, i.e. the reweighting has corrected the bias.

Since there is no reason to treat child branches differently, we set $w'_c = w' = r(n)w$ where $w = 1$ is the parent weight. Eq. A4 now reduces to

$$1 = \langle n' r(n') \rangle = \sum_{n'=1}^{n_{\max}} b(n') n' r(n'), \quad (\text{A5})$$

Eq. A5 constrains the choice weight factors r for a given branching distribution b . For instance, if $n_{\max} = 2$, and $(b(n))_{n=0,1,2} = (.5, .2, .3)$ then the choices $(r(n))_{n=1,2} = (0, 5/3)$, $(5, 0)$ and $(5/4, 5/4)$, are all unbiased.

Generalizing Eq. A2, we can then estimate p from a branched simulation of S independent trees as

$$\begin{aligned} p(x_m, t_m; \dots; x_1, t_1) &\simeq \frac{1}{S} \sum_{s=1}^S r(n'_s) \sum_{c_s=1}^{n'_s} \delta(x_{t_m}^{c_s} - x_m) \cdots \delta(x_{t_{m'+1}}^{c_s} - x_{m'+1}) \\ &\quad \times \delta(x_{t_{m'}}^s - x_{m'}) \cdots \delta(x_{t_1}^s - x_1) \quad (\text{A6}) \end{aligned}$$

The first line on the rhs contains sample points after branching into n'_s children (if any), and the second line those before branching (if any); the weight factors r satisfy eq. A5. Note that repeated simulations average over not only the system but also the branching randomness. For instance, the total weight of all trajectories at a given time is conserved only on average over branching outcomes.

So far we have shown that a single stochastic branching move at time t' and re-weighted sampling according to Eq. A6 does not introduce a bias if the child number probabilities $b(n')$ and child weight factors $r(n')$ obey the condition of weight conservation on average, Eq. A5. We note that to verify this numerically, one would need to generate many trees always with branching time t' and count joint hits of all bins around the points x_1, \dots, x_m at the times t_1, \dots, t_m ; if the final time is larger than the branching time, hits are re-weighted the factor $r(n_s)$ appropriate for the respective branch number n_s .

3. Adaptive branching probabilities

We now show that the branching probabilities may be adapted according to an arbitrary protocol. To see this, let b depend on an arbitrary parameter α such that $\sum_{n'} b_\alpha(n') = 1$, and choose r_α accordingly, such that $\sum_{n'} b_\alpha(n') n' r_\alpha(n') = 1$ for all values of α . Then, following the preceding discussion, a branching event is unbiased for any α . We may take α to be any random variable, depending on a set of bias control parameters β via some density $\rho(\alpha|\beta)$. Since Eq. A3 now reads

$$\begin{aligned} W(x_m, t_m; \dots; x_1, t_1) &= \int d\alpha \rho(\alpha|\beta) \sum_{n'} b_\alpha(n') n' p(x_m, t_m; \dots; x_1, t_1) r_\alpha(n') \\ &= \int d\alpha \rho(\alpha|\beta) p(x_m, t_m; \dots; x_1, t_1) \\ &= p(x_m, t_m; \dots; x_1, t_1), \end{aligned} \quad (\text{A7})$$

unbiased sampling is still guaranteed. We thus have complete freedom to introduce dependency of the branching probabilities on arbitrary extra information, including the past or future system state, or the history of the simulation.

4. Multiple branching events

Finally, we generalize the arguments above to multiple independent and non-synchronized branching events on different branches of a tree. The idea is that introducing a new branching event at a time t'' on an existing child branch is bias-free, as long as the new b and r fulfill Eq. A5. We then argue by induction over tree generations.

Each new branching event with n^l children adds a weight factor $r(n^l)$, so that the instantaneous weight along a trajectory passing by the branching events $(n^1, t^1; n^2, t^2; \dots; n^k, t^k)$ becomes

$$w(t; n^1, t^1; \dots; n^k, t^k) = \prod_{l: t^l < t} r(n^l); \quad (\text{A8})$$

The system density estimate Eq. A6 now takes the form of a hierarchical sum over all weighted branches existing at the final time. This is simplest to write down using a recursive definition. We index a particular branch (trajectory segment) anywhere in the tree by the sequence of the child indices starting from the root, $\gamma = (c, c', \dots, c^k)$, and let $|\gamma| = k + 1$ denote its nesting depth. The zeroth child index $c^0 = c \equiv s$ is the tree index. The points along the trajectory from the root up to and including γ are denoted x_t^γ . Let

$$\pi^\gamma(x_m, t_m; \dots; x_1, t_1) = \prod_{\tilde{m}=1}^m \delta(x_{t_{\tilde{m}}}^\gamma - x_{\tilde{m}}), \quad (\text{A9})$$

if γ has no children, and define recursively

$$\pi^\gamma(x_m, t_m; \dots; x_1, t_1) = r_\gamma \sum_{c=1}^{n_\gamma} \pi^{(\gamma, c)}(x_m, t_m; \dots; x_1, t_1), \quad (\text{A10})$$

if γ has n_γ children with weight factor r_γ . This recursion terminates since the simulation has a finite maximal nesting depth K . Note that $\pi^\gamma(x_m, t_m; \dots; x_1, t_1)$ counts the weighted hits of γ and all its descendants to the points $(x_m, t_m; \dots; x_1, t_1)$. The weight is effectively given by Eq. A8.

All S simulated trees together can then be represented as the descendants of the empty path $()$. If we define $r_{()} = 1/S$, $n_{()} = S$ and $|\cdot| = 0$, the estimate for the joint m -point density is

$$p(x_m, t_m; \dots; x_1, t_1) \simeq \langle \pi^{()}(x_m, t_m; \dots; x_1, t_1) \rangle. \quad (\text{A11})$$

Since Eq. A9 coincides with the brute force estimate of the density, and Eq. A10 is unbiased by the arguments in sec. A2, it follows by induction over K that the density estimate Eq. A11 is unbiased.

Eqs. (A9, A10, A11) directly translate into a histogramming algorithm, which recursively parses a trajectory tree while binning weighted counts. Binning can be carried out on-line. This is straightforward if only the one-point density $p(x, t)$ is desired; online updates of m -point densities would require $m - 1$ nested inner loops over the tree for each simulation step.

As a corollary, arbitrary m -point observables A can be estimated along the same lines. One just has to replace π by π_A in Eq. A10 and A9 by

$$\pi_A^\gamma(t_m; \dots; t_m) = A(x_{t_m}, \dots, x_{t_1}), \quad (\text{A12})$$

if γ has no children. For instance, a two-point auto-correlation function would correspond to $A(x_{t_2}, x_{t_1}) = x_{t_2} x_{t_1} - \langle x_{t_2} \rangle \langle x_{t_1} \rangle$.

Appendix B: Variance of weighted bin counts

Here, we write the variance of the total weight $W = \sum_{a=1}^N w_a$ accumulated in a given bin B_{li} in terms of the statistics of the trajectory weights w_a and the number N of trajectories which has reached bin B_{li} . We recall the ‘law of total variance’:

$$\langle A^2 \rangle - \langle A \rangle^2 = \langle \langle A^2 | \mathcal{C} \rangle - \langle A | \mathcal{C} \rangle^2 \rangle + \langle (\langle A | \mathcal{C} \rangle)^2 - \langle A | \mathcal{C} \rangle^2 \rangle \quad (\text{B1})$$

Here the inner expectation values are conditioned on some event \mathcal{C} , and the outer expectations average over \mathcal{C} . We will denote the conditional variance as $\langle \delta X^2 | \mathcal{C} \rangle \equiv \langle X^2 | \mathcal{C} \rangle - \langle X | \mathcal{C} \rangle^2$ for any observable X . Eq. B1 then becomes $\langle \delta A^2 \rangle = \langle \langle \delta A^2 | \mathcal{C} \rangle \rangle + \langle \delta \langle A | \mathcal{C} \rangle^2 \rangle$.

Consider an idealized NS-FFS simulation in which trajectories arriving at B_{li} are uncorrelated. Specifically, we assume that the incoming trajectory weights $\{w_a\}$ are mutually independent and independent of N , and that arrivals are a Poisson process. For the mean total crossing weight we then obtain $\langle W \rangle = \langle N \rangle \langle w_a \rangle$. The total weight variance is given by

$$\begin{aligned} \langle \delta W^2 \rangle &= \langle \langle \delta W^2 | N \rangle \rangle + \langle \delta \langle W | N \rangle^2 \rangle \\ &= \langle N \rangle \langle \delta w_a^2 \rangle + \langle w_a \rangle^2 \langle \delta N^2 \rangle \\ &= \langle N \rangle (\langle \delta w_a^2 \rangle + \langle w_a \rangle^2); \end{aligned}$$

the relative variance of the collected weight becomes

$$\frac{\langle \delta W^2 \rangle}{\langle W \rangle^2} = \frac{1}{\langle N \rangle} \left[1 + \frac{\langle \delta w_a^2 \rangle}{\langle w_a \rangle^2} \right].$$

In the more general case, we now assume that correlations between branches increase the noise while respecting the same scaling with the count number, and write

$$\frac{\langle \delta W^2 \rangle}{\langle W \rangle^2} \simeq \frac{\alpha_N}{\langle N \rangle} \left[1 + \alpha_w \frac{\langle \delta w_a^2 \rangle}{\langle w_a \rangle^2} \right] \quad (\text{B2})$$

where $\alpha_N > 1$ if trajectories arrive in bunches and $\alpha_w > 1$ if their weights are correlated. The crossing flux j_{li} is estimated as W_{li}/S and thus has the same noise, Eq. B2, as W_{li} .

The noise in bin weights can be split up further. The incoming weights $\{w_a\}$ are distributed with a mean and variance which result from both the inter-bin variance between starting bins $B_{l'i'}$ and from the intra-bin variances of outgoing weights from within $B_{l'i'}$. We have

$$\begin{aligned} \langle w_a \rangle &= \langle \langle w_a | l'i' \rangle \rangle, \text{ and using Eq. B1,} \\ \langle \delta w_a^2 \rangle &= \langle \langle \delta w_a^2 | l'i' \rangle \rangle + \langle \delta \langle w_a | l'i' \rangle^2 \rangle; \end{aligned} \quad (\text{B3})$$

here the first term is the mean intra-bin variance within originating bins, and the second term is the inter-bin variance. Plugging in Eq. B3 we can write the noise in W as

$$\frac{\langle \delta W^2 \rangle}{\langle W \rangle^2} = \frac{\alpha_N}{\langle N \rangle} \left[1 + \alpha_w \left\{ \frac{\langle \langle \delta w_a^2 | l'i' \rangle \rangle}{\langle w_a \rangle^2} + \frac{\langle \delta \langle w_a | l'i' \rangle^2 \rangle}{\langle w_a \rangle^2} \right\} \right]. \quad (\text{B4})$$

As the simulation progresses, crossing flux estimates converge, so that all trajectories leaving $B_{l'i'}$ are eventually assigned the same weight. The intra-bin term $\langle \langle \delta w_a^2 | l'i' \rangle \rangle / \langle w_a \rangle^2 \propto \langle \delta j_{l'i'}^2 \rangle / \langle j_{l'i'} \rangle^2$ thus vanishes as $N \rightarrow \infty$. The inter-bin term $\langle \delta \langle w_a | l'i' \rangle^2 \rangle / \langle w_a \rangle^2$ reflects the non-uniform transition probabilities between bins and persists also in steady state.

In order to balance noise contributions, it seems reasonable to choose bin size and interface spacing such that $\alpha_w \langle \delta \langle w_a | l'i' \rangle^2 \rangle \simeq \langle w_a \rangle^2$ in an equilibrated simulation.

Appendix C: Multi-dimensional progress coordinates

A multi-coordinate NS-FFS simulation can be set up as follows. First, find K progress coordinates $\{\lambda^k\}_{k=1\dots K}$, with corresponding sets of levels $\{\lambda_0^k < \dots < \lambda_{L_k}^k\}_{k=1\dots K}$. Denote the interval $(\lambda_{l-1}^k, \lambda_l^k) = \Lambda_l^k$. Define a subset of $K' \leq K$ progress coordinates $\{\lambda^{k'}\}_{k'=1\dots K'}$. Only the first K' sets of interfaces will trigger branching events; for these interfaces, define bins:

$$\begin{aligned} B_{l_1 l_2 \dots l_{k'} \dots l_K}^{k'} &= \\ \{(x, t) | \lambda^{k'}(x, t) = \lambda_{l_{k'}}^{k'} \text{ and } \lambda^k(x, t) \in \Lambda_{l_k}^k \text{ for } k \neq k'\} & \quad (\text{C1}) \end{aligned}$$

Then, proceed as before: branching moves are triggered on the first K' interfaces, based on the corresponding crossing weights $H_{l_1 l_2 \dots l_{k'} \dots l_K}^{k'}$.

In this setting, time is treated as another progress coordinate. To recover the NS-FFS setups discussed in sec. III in this setting, let $K = 2, K' = 1$. For the λ -if setup, set $\lambda^1 = \lambda, \lambda^2 = t, L_1 = L$ and $L_2 = I$; for the t -if setup, set $\lambda^1 = t, \lambda^2 = \lambda, L_1 = I$ and $L_2 = L$.

Appendix D: Piecewise linear potential

We solve the Fokker Planck equation associated with Eq. 7 in Laplace space,

$$sp - p_0 + \partial_x (\text{sgn}(x) a D p - D \partial_x p) = 0, \quad (\text{D1})$$

where $p = p(x, s) = \int_0^\infty p(x, t) e^{-st} dt$ is the Laplace transformed density, and the initial condition $p_0(x) = p(x, t = 0_-) \equiv 0$. Particles are injected at $t = 0$; the boundary conditions for the total flux $j(x, t) = \text{sgn}(x) a D p(x, t) - D \partial_x p(x, t)$ at the left boundary read

$$\begin{aligned} j(x = -1, t) &= \delta(t) \text{ or in Laplace space,} \\ j(x = -1, s) &= 1. \end{aligned} \quad (\text{D2})$$

They incorporate both the reflecting boundary and the injection of unit probability at $x = -1$ at $t = 0$.

At the right boundary, we consider either reflecting (referred to as r/r) or absorbing (r/a) conditions:

$$j(x = 1, s) = 0, \text{ or} \quad (\text{D3})$$

$$p(x = 1, s) = 0, \text{ respectively.} \quad (\text{D4})$$

Eqs. D1, D2 and D3/D4 can be solved by using the ansatz $p(x, s) = e^{\frac{1}{2}ax}\tilde{p}(x, s)$, and joining solutions in the regions $x < 0$ and $x > 0$. After straightforward but lengthy algebra, the solution in the r/r case can be written as

$$p(x, s) = e^{-\frac{1}{2}a(1-|r|)} \times \frac{qa \sinh(q(1-|r|)) - 2q^2 \cosh(q(1-r)) - \theta(-r)a^2 \sinh(qr) \sinh(q)}{2sa \sinh^2(q) - 4sq \sinh(q) \cosh(q)} \quad (\text{D5})$$

where $q = \sqrt{s/D + a^2/4}$ and θ is the unit step function. For the r/a case we obtain

$$p(x, s) = e^{-\frac{1}{2}a(1-|r|)} \times \frac{qa \sinh(q(1-r)) - \theta(-r)\frac{a^2}{2}(\cosh(q(1+r)) - \cosh(q(1-r)))}{2sa \sinh^2(q) + aDq^2}. \quad (\text{D6})$$

Both Green's functions have poles only on the non-positive real s -axis. The barrier crossing time can be extracted by solving for the largest negative pole at $-s_{AB}$; in the r/r case, $s_{AB} = k_{AB} + k_{BA} = 2k_{AB}$ while in the r/a case, $s_{AB} = k_{AB}$. If the barrier is high, we may expand the relevant denominators in eqs. D5, D6 for $\frac{|s|}{Da^2} \ll 1$. We find in both cases (r/r and r/a) that the waiting time scales exponentially with the barrier:

$$\tau_{AB} = s_{AB}^{-1} = \frac{2e^a}{a^2D} + O(1/a). \quad (\text{D7})$$

In the r/a case we can also evaluate the exit flux through the absorbing boundary,

$$j(x=1, s) = j_{AB}(s) = \frac{4q^2}{a^2 - (a^2 - 4q^2) \cosh(2q)} \quad (\text{D8})$$

The equilibration time within basin A can be estimated as a diffusion time for covering the thermally accessible range of x ; alternatively, a more accurate pre-factor can be obtained by solving Eq. D1 with Eq. D2 as above but replacing the W-shaped potential by a uniformly increasing ramp potential $U = ax$ of the same slope. Evaluating the slowest relaxation time now gives the relaxation time in A ,

$$\tau_A = \frac{4}{a^2D}. \quad (\text{D9})$$

Finally, the crossing time τ_C is the mean first passage time for diffusion through region C ; that is, from the boundary of region A at $-x_C$, up the barrier and down on the other side until reaching the boundary of region B at x_C , without returning to A . If we disregard trajectories that cross $x = 0$ more than once, τ_C is the sum of the mean first passage times for the two segments, with negative and positive constant drift, respectively. The non-intuitive but well-known result is that diffusion of successful transition paths up the barrier takes as long as down the barrier (see e.g.³⁰). In the drift-dominated regime, the mean first passage time is controlled by the drift velocity aD . One obtains

$$\tau_C = 2 \times \frac{x_C}{aD} \lesssim \frac{2}{aD}. \quad (\text{D10})$$

REFERENCES

- ¹C. H. Bennett. Molecular dynamics and transition state theory: The simulation of infrequent events. In *Algorithms for Chemical Computations*, volume 46 of *ACS Symposium Series*, pages 63–97. American Chemical Society, June 1977. 0.
- ²David Chandler. Statistical mechanics of isomerization dynamics in liquids and the transition state approximation. *The Journal of Chemical Physics*, 68(6):2959, 1978.
- ³Peter G Bolhuis, David Chandler, Christoph Dellago, and Phillip L Geissler. Transition path sampling: throwing ropes over rough mountain passes, in the dark. *Annu Rev Phys Chem*, 53:291–318, 2002. TPS Review.
- ⁴Titus S. van Erp, Daniele Moroni, and Peter G. Bolhuis. A novel path sampling method for the calculation of rate constants. *The Journal of Chemical Physics*, 118(17):7762, 2003. TIS.
- ⁵Daniele Moroni, Peter G Bolhuis, and Titus S van Erp. Rate constants for diffusive processes by partial path sampling. *J Chem Phys*, 120(9):4055–4065, Mar 2004.
- ⁶Anton K Faradjian and Ron Elber. Computing time scales from reaction coordinates by milestoning. *J Chem Phys*, 120(23):10880–10889, Jun 2004.
- ⁷Weinan E, Weiqing Ren, and Eric Vanden-Eijnden. String method for the study of rare events. *Physical Review B*, 66(5):052301, 2002.
- ⁸G. A. Huber and S. Kim. Weighted-ensemble brownian dynamics simulations for protein association reactions. *Biophys J*, 70(1):97–110, Jan 1996.
- ⁹B. W. Zhang, D. Jasnow, and D. M. Zuckerman. The "weighted ensemble" path sampling method is statistically exact for a broad class of stochastic processes and binning procedures. *Journal of Chemical Physics*, 132(5), 2010.
- ¹⁰Matthias Heymann and Eric Vanden-Eijnden. Pathways of maximum likelihood for rare events in nonequilibrium systems: application to nucleation in the presence of shear. *Phys Rev Lett*, 100(14):140601, Apr 2008.
- ¹¹Aryeh Warmflash, Prabhakar Bhimalapuram, and Aaron R Dinner. Umbrella sampling for nonequilibrium processes. *J Chem Phys*, 127(15):154112, Oct 2007.
- ¹²RJ Allen, D Frenkel, and PR ten Wolde. Simulating rare events in equilibrium or nonequilibrium stochastic systems. *J. Chem. Phys.*, 124(2):024102, JAN 14 2006.
- ¹³Ivan Osorio, Hitten Zaveri, Mark G. Frei, and Susan Arthurs. *Epilepsy: The Intersection of Neurosciences, Biology, Mathematics, Engineering and Physics*. CRC Press, April 2011.
- ¹⁴Girol M Stiel, Jordi Garcia-Ojalvo, Louisa M Liberman, and Michael B Elowitz. An excitable gene regulatory circuit induces transient cellular differentiation. *Nature*, 440(7083):545–550, Mar 2006.
- ¹⁵Joshua T Berryman and Tanja Schilling. Sampling rare events in nonequilibrium and nonstationary systems. *J Chem Phys*, 133(24):244101, Dec 2010.
- ¹⁶G. E. Crooks and D. Chandler. Efficient transition path sampling for nonequilibrium stochastic dynamics. *Phys Rev E Stat Nonlin Soft Matter Phys*, 64(2 Pt 2):026109, Aug 2001.
- ¹⁷Rosalind J Allen, Chantal Valeriani, and Pieter Rein ten Wolde. Forward flux sampling for rare event simulations. *Journal of Physics: Condensed Matter*, 21(46):463102 (21pp), 2009.
- ¹⁸P. Grassberger. Pruned-enriched rosenbluth method: Simulations of theta polymers of chain length up to 1,000,000. *Physical Review E*, 56(3):3682–3693, September 1997.
- ¹⁹Thomas Prellberg and Jarosław Krawczyk. Flat histogram version of the pruned and enriched Rosenbluth method. *Phys Rev Lett*, 92(12):120602, Mar 2004.
- ²⁰Nils B Becker and Pieter Rein ten Wolde. Switching events in non-stationary systems. submitted to JCP.
- ²¹The term 'free energy' is used in the sense of a negative logarithm of the density; this remains meaningful also for non-equilibrium steady states.
- ²²Hsiao-Ping Hsu and Peter Grassberger. A review of monte

- carlo simulations of polymers with PERM. *Journal of Statistical Physics*, 144:597–637, July 2011.
- ²³JAC Weideman and LN Trefethen. Parabolic and hyperbolic contours for computing the bromwich integral. *Mathematics of Computation*, 76(259):1341, 2007.
- ²⁴Patrick B Warren and Pieter Rein ten Wolde. Enhancement of the stability of genetic switches by overlapping upstream regulatory domains. *Phys Rev Lett*, 92(12):128101, Mar 2004.
- ²⁵P. B. Warren and P. R. ten Wolde. Chemical models of genetic toggle switches. *Journal of Physical Chemistry B*, 109(14):6812–6823, April 2005.
- ²⁶Mark Ptashne. *A Genetic Switch, Third Edition: Phage Lambda Revisited*. Cold Spring Harbor Laboratory Press, 3rd edition, April 2004.
- ²⁷We note that if the microscopic dynamics of the system has memory in itself, the history for $t < 0$ may need to be specified as part of the initial condition²⁰.
- ²⁸Ernesto E. Borrero and Fernando A. Escobedo. Optimizing the sampling and staging for simulations of rare events via forward flux sampling schemes. *J Chem Phys*, 129(2):024115, Jul 2008.
- ²⁹Siebe B. van Albada, Sorin Tănase-Nicola, and Pieter Rein ten Wolde. The switching dynamics of the bacterial flagellar motor. *Mol Syst Biol*, 5:316, 2009.
- ³⁰Sidney Redner. *A guide to first-passage processes*. Cambridge Univ. Press, Cambridge [u.a], 1. publ. edition, 2001.