

TOPICAL REVIEW

Forward Flux Sampling for rare event simulations

Rosalind J Allen¹, Chantal Valeriani¹ and Pieter Rein ten Wolde²

1. SUPA, School of Physics and Astronomy, The University of Edinburgh, Mayfield Road, Edinburgh EH9 3JZ, UK.

2. FOM Institute for Atomic and Molecular Physics, Science Park 113, 1098 XG Amsterdam, The Netherlands

E-mail: rallen2@ph.ed.ac.uk

Abstract. Rare events are ubiquitous in many different fields, yet they are notoriously difficult to simulate because few, if any, events are observed in a conventional simulation run. Over the past several decades, specialised simulation methods have been developed to overcome this problem. We review one recently-developed class of such methods, known as Forward Flux Sampling. Forward Flux Sampling uses a series of interfaces between the initial and final states to calculate rate constants and generate transition paths, for rare events in equilibrium or nonequilibrium systems with stochastic dynamics. This review draws together a number of recent advances, summarises several applications of the method and highlights challenges that remain to be overcome.

<i>Forward flux sampling</i>	2
1 Introduction	3
2 Background	4
2.1 “Bennett-Chandler” type methods	4
2.2 Transition Path Sampling	5
2.3 Transition Interface Sampling	6
2.4 Milestoning	7
2.5 The Weighted Ensemble method	7
2.6 The Finite Temperature String Method	7
3 Forward flux sampling	8
3.1 FFS variants	10
3.1.1 “Direct” FFS	10
3.1.2 Branched Growth	11
3.1.3 The “Rosenbluth-like” method	12
3.2 Requirement for stochastic dynamics	13
3.3 Pruning	14
4 Computational efficiency: prediction and optimisation	14
4.1 Analytical expressions for the efficiency	14
4.2 Optimising the efficiency	17
4.2.1 Optimising the number of trial runs	17
4.2.2 Optimising the interface positions	18
4.2.3 Efficiency gains	19
5 The order parameter and the committor	19
5.1 Defining the order parameter	20
5.2 The committor and the reaction coordinate	21
5.3 Using the committor to optimise the order parameter in FFS	22
6 Computing stationary distributions	23
6.1 Obtaining stationary distributions from FFS simulations	24
6.2 Forward flux / umbrella sampling	26
6.3 Nonequilibrium umbrella sampling in multiple dimensions	26
7 Applications	27
7.1 Genetic switch flipping	27
7.2 Homogeneous crystal/bubble nucleation	31
7.3 Nucleation in a sheared Ising model	32
8 Challenges and future directions	34

1. Introduction

“Rare events” are fluctuation-driven transitions that have a low probability of occurring, but often have important consequences when they do occur. Examples are ubiquitous, ranging from large-scale events such as earthquakes, global climate changes, financial crashes and telecommunications network failures, to smaller-scale processes typical of soft condensed matter and biological physics, such as activated chemical reactions, nucleation phenomena, protein conformational changes, switching in biochemical networks and translocation through pores. Computer simulation has an important role to play in understanding rare events, especially as they are often difficult to study experimentally. However, rare events are notoriously difficult to simulate, simply because in the typical simulation time few, if any, events happen. To address this issue, specialised techniques for simulating rare events have been developed in various different contexts over many years. Excellent reviews have already been published on this subject in the fields of condensed matter, chemical and biological physics [1, 2, 3, 4, 5, 6]. In this short topical review article, we focus only on one recently developed class of techniques, known as Forward Flux Sampling (FFS). Although this class of methods is still rather young (at least in this field), it has been applied to a variety of different problems, and several variants and improvements to the methodology have recently been proposed. A number of potential problems have also been highlighted. This article aims to bring together these developments, together with practical advice on using the methods and suggestions for directions of future research. In the late stages of preparation of this article, we became aware of an almost simultaneous review, also focusing on FFS, by Escobedo *et al* [7]. Although some duplication of material between these articles is inevitable, we hope to present a complementary perspective.

Typically, when studying a rare event process, one wishes to know how often the event happens, or equivalently the rate constant k_{AB} for transitions from an initial state A to a final state B . If the transition occurs between two time-invariant steady states and is itself is fast compared to k_{AB}^{-1} , the rate constant k_{AB} will be time-invariant. The distribution function $F(T)$ for the time taken for an equilibrated configuration in the A state to “escape” to the B state is then given by:

$$F(T) = k_{AB} e^{-k_{AB}T} \quad (1)$$

One is usually also interested in the mechanism by which the rare event process happens. For example, for a crystal nucleation problem, one might wish to know the crystal structure and shape of the growing nucleus, or for a protein folding problem, in what order the secondary structure elements form. Information on the mechanism can be obtained by sampling the transition path ensemble (TPE), which is the ensemble of trajectories corresponding to transitions from A to B . However, extracting simple and intuitive conclusions from these transition paths can sometimes be difficult.

The FFS methodology discussed here was originally developed for simulations of rare events in nonequilibrium systems, although it can also be used for equilibrium systems. In this review, we consider “equilibrium” systems to be those whose dynamical

rules obey detailed balance [regardless of whether they are actually in a stationary state]. Detailed balance has the consequence that for these systems, the stationary phase space probability distribution is given by the Boltzmann distribution [1], and the system dynamics is time reversible. In contrast, nonequilibrium dynamical systems do not obey detailed balance, their stationary phase space distribution is not known *a priori* and their dynamics are not time reversible. Nonequilibrium systems present a host of important and interesting rare event processes. However, these systems pose particular challenges for rare event simulation methods, as we shall discuss.

2. Background

In this section, we present a brief overview of rare event simulation methods in the area of condensed matter, chemical and biological physics. Our aim is to provide the background information necessary for the subsequent discussion of the FFS technique, rather than to give a comprehensive review. Consequently, some important methods are omitted completely, or only discussed very briefly, for which we apologise.

2.1. “Bennett-Chandler” type methods

“Bennett-Chandler”, or “reactive flux” methods are based on the transition state theory (TST) expression for the rate constant [8, 9, 10]. In TST, phase space is partitioned by a dividing surface between “reactant” and “product” regions. The rate constant is given by:

$$k_{AB}^{TST} = \frac{\langle |\dot{q}| \rangle}{2} \frac{e^{-W(q^*)/(k_B T)}}{\int_{-\infty}^{q^*} dq e^{-W(q)/(k_B T)}} \quad (2)$$

where q is an order parameter that separates the reactant and product regions and measures the progress of the system between these regions, q^* defines the dividing surface and $W(q)$ is the reversible work needed to move the system from state A to a value q of the order parameter. This is proportional to the free energy, projected onto the coordinate q . The exponential term describes the equilibrium (Boltzmann) probability of finding the system at the dividing surface relative to the A state, while the term $\langle |\dot{q}| \rangle / 2$ is the average velocity of the system from reactants to products across the dividing surface. Eq.(2) assumes that all crossings of the dividing surface contribute to the rate constant, while in reality a single trajectory can cross the dividing surface many times. The “reactive flux” formalism therefore corrects the TST expression with a “transmission coefficient” κ , which is less than unity [11, 12, 3]:

$$k_{AB} = \kappa(t) k_{AB}^{TST} \quad (3)$$

The transmission coefficient $\kappa(t)$ replaces the factor $\langle |\dot{q}| \rangle / 2$ in Eq.(2) by the average initial velocity of trajectories, initiated from an equilibrium distribution at q^* , which are in the product basin after time t . For times intermediate between the molecular timescale and the timescale for transitions between the reactant and product basins, $\kappa(t)$

is independent of time; for this reason we do not include an explicit time dependence for k_{AB} in Eq.(3). In Bennett-Chandler type methods, one chooses an order parameter q and computes the free energy profile $W(q)$ using a method such as Umbrella Sampling [13, 14, 15]. The transmission coefficient κ is then computed by initiating a large number of trajectories from an equilibrium distribution restricted to $q = q^*$ (usually taken to be the maximum of $W(q)$), and counting the fraction of these that end up in the product state. Bennett-Chandler-type methods are conceptually simple, easy to implement, and have been widely used. However, because of the assumption of the Boltzmann distribution inherent in Eq.(2), these methods are not suitable for nonequilibrium systems. These methods also tend to be rather sensitive to poor choices of the order parameter, since this will result in a small value of κ which is hard to compute accurately.

2.2. Transition Path Sampling

Transition Path Sampling (TPS) methods [16, 17, 3, 2] focus directly on sampling the transition path ensemble (TPE) using a Monte-Carlo procedure in trajectory space. A single trajectory connecting the reactant and product regions of phase space is generated, and this is used to produce new trajectories. Several methods are available for generating new trajectories, of which probably the most widely used is “shooting”. Here, a time-slice from the initial trajectory is then selected, and a small change is made, usually to the momentum coordinates. A new trajectory is then generated by integrating the system dynamics forward and backward in time from this slightly altered phase space point. If the new trajectory still joins the reactant and product basins, it is accepted into or rejected from a collection of computed transition paths with a probability that depends on its path weight, the weight of a path of length n steps with phase space coordinates $\{x\}$ being given by:

$$\mathcal{P}[\{x\}] = \rho(x_0) \prod_{i=0}^{n-1} p(x_i \rightarrow x_{i+1}) \quad (4)$$

where $\rho(x_0)$ is the phase space probability density for the initial point in the path and $p(x_i \rightarrow x_{i+1})$ is the probability of making a simulation step from x_i to x_{i+1} . In practice, for Molecular Dynamics simulations in the NVE ensemble, one can simply accept all generated paths that connect the reactant and product basins. However, because Eq.(4) requires knowledge of the phase space distribution ρ_0 , TPS is not suitable for nonequilibrium systems (although a TPS method for nonequilibrium systems has been proposed [18]).

Computation of the rate constant k_{AB} in TPS is based on the correlation function $C(t)$ [3]:

$$C(t) = \frac{\langle h_A(x_0)h_B(x_t) \rangle}{\langle h_A(x_0) \rangle} \quad (5)$$

where $h_A(x)$ is unity in the reactant basin and zero elsewhere, and $h_B(x)$ is unity in the product basin and zero elsewhere. $C(t)$ is the probability of finding the system in the product basin at time t , given that at time 0 it was in the reactant basin. For times

longer than the molecular timescale, $C(t) \approx k_{AB}t$. In practice, one computes $C(t)$ in two stages: a TPS simulation between reactant and product regions, and an ‘‘umbrella sampling’’ procedure in which the end points of transition paths are constrained to lie in a series of windows between A and B, defined by an order parameter. More information about TPS, and about the many improvements to the method which have been made, are given in Refs [16, 17, 3, 2, 4, 6]. TPS has the advantage that it samples paths without the need for an order parameter, although an order parameter is needed to compute the rate constant.

2.3. Transition Interface Sampling

Transition Interface Sampling (TIS) [19, 5] is a variant of TPS in which the rate constant is calculated differently. In TIS, phase space is divided up by a series of non-intersecting interfaces, defined by an order parameter λ , such that the reactant region is defined by $\lambda < \lambda_A = \lambda_0$ and the product region by $\lambda > \lambda_B = \lambda_n$. The expression for k_{AB} used in TIS is then [19, 5, 20]:

$$k_{AB} = \frac{\bar{\Phi}_{A,n}}{h_A} = \frac{\bar{\Phi}_{A,0}}{h_A} P(\lambda_n|\lambda_0) \quad (6)$$

where $\bar{\Phi}_{A,n}$ is the steady-state flux of trajectories leaving the A state and reaching interface λ_n (*i.e.* the B state), and h_A is a history-dependent function that is unity if a trajectory was more recently in A than in B , and zero otherwise. The right-hand side of Eq.(6) expresses the fact that the flux of trajectories that leave A and cross λ_n is equal to the flux of those leaving A and crossing λ_0 , multiplied by the probability $P(\lambda_n|\lambda_0)$ that a trajectory that crosses λ_0 , coming from A , will subsequently reach λ_n before returning to A . The flux $\bar{\Phi}_{A,0}$ is easy to calculate, since trajectories coming from A cross λ_0 frequently. However, the probability $P(\lambda_n|\lambda_0)$ is small and thus difficult to calculate. This difficulty is overcome by expressing $P(\lambda_n|\lambda_0)$ as [19]:

$$P(\lambda_n|\lambda_0) = \prod_{i=0}^{n-1} P(\lambda_{i+1}|\lambda_i) \quad (7)$$

where the product is over all interfaces and $P(\lambda_{i+1}|\lambda_i)$ is the conditional probability that a trajectory that comes from A and crosses λ_i for the first time will subsequently reach λ_{i+1} instead of returning to A . We note that expression (7) does not involve a Markovian approximation, because the probabilities $P(\lambda_{i+1}|\lambda_i)$ are conditional on the history of the trajectories reaching λ_i . Expressions (6) and (7) are known as the ‘‘effective positive flux’’ formulation of the rate constant.

In TIS, the flux $\bar{\Phi}_{A,0}$ is computed using a ‘‘brute-force’’ simulation in the A basin. TPS is then used to sample the ensemble of transition paths from the reactant basin to λ_i (using as an initial path one of the successful paths to λ_{i-1}). The fraction of transition paths ultimately reaching λ_{i+1} , as opposed to λ_A , in the ensemble of paths from A to λ_i , is an estimate for $P(\lambda_{i+1}|\lambda_i)$. It is important to recognise that in TIS, the interfaces are simply used as a convenient way of dividing the transition paths into sections. The order

parameter need not correspond to the true reaction coordinate and it is not assumed that the system loses its “memory”, or becomes uncorrelated, between one interface and the next. However, a version of TIS in which one does assume decorrelation between interfaces, Partial Path Transition Interface Sampling (PPTIS) [21], is more efficient for diffusive transitions. A number of improvements to TIS have been developed in recent years, including computation of free-energy barriers [22], swapping partial paths between ensembles at different interfaces [5, 23, 24] and sampling transitions to multiple final states [25].

2.4. Milestoning

The “Milestoning” method [26, 27] also uses a series of interfaces between the initial and final states, defined by an order parameter λ . In contrast to TIS, milestoning does assume memory loss between interfaces. Short simulation trajectories are initiated from quasi-equilibrium (or first hitting point [28]) distributions at interface λ_i , and continued until they reach the adjacent interfaces λ_{i-1} or λ_{i+1} . From these trajectories, first passage time distributions for transitions between interfaces are obtained, and these can be used to compute the time evolution of the system. In contrast to the other methods discussed above (and FFS), Milestoning does not assume that the transition is a simple two-state process with exponential kinetics. A new variant of this method, “Markovian milestoning”, has recently been proposed [29]: here, the interfaces are defined by Voronoi polyhedra (see Section 5.1) and the kinetic information is obtained by running trajectories confined to these interfaces by reflecting boundary conditions.

2.5. The Weighted Ensemble method

The Weighted Ensemble method of Huber and Kim [30] is also an interface-based method, and is closely related to FFS. This method divides the phase space region between the reactant and product states into a series of bins, and simulates a collection of “walkers”, each of which carries a probability weight, and which either merge or divide as they progress between bins, so as to maintain the number of walkers in each bin. By monitoring the flux of walkers across the interfaces the transition rate constant can be efficiently computed.

2.6. The Finite Temperature String Method

The Finite Temperature String method (FTS) is different in concept to the above methods, since it focuses on the “principal curve” between A and B [31]. This is the path that follows the averaged position of the system, projected onto a series of hyperplanes perpendicular to the path itself. For systems with overdamped Langevin dynamics, the free energy along the principal curve can be directly related to the committor function [31]. The FTS method defines a string of “beads”, or configurations, between A and B, and iteratively refines the positions of the beads until the string corresponds to

the principal curve. In the original version of FTS, this was achieved by performing constrained simulations on the hyperplanes perpendicular to the string. However, a simplified version has recently been published [32] in which one instead defines Voronoi polyhedra around the beads (see Section 5.1) and simulates multiple copies of the system, each one constrained to lie inside one of the Voronoi polyhedra. The beads then evolve towards the average configuration obtained in the simulation for that polyhedron.

3. Forward flux sampling

Forward flux sampling (FFS) methods [33, 34] were developed to simulate rare events in nonequilibrium systems with stochastic dynamics. Nonequilibrium systems are ubiquitous in condensed matter, chemical and biological physics. However, their lack of detailed balance and the consequent absence of a Boltzmann-like stationary distribution function and lack of time reversal symmetry mean that Bennett-Chandler-type methods, TPS, TIS, Milestoning, and most versions of the string method, are not suitable for these systems. FFS is thus one of only a few methods available for simulating rare events in nonequilibrium systems. For equilibrium systems, FFS provides an alternative to the above methods. The development of FFS was inspired by TIS. However, we have subsequently become aware of the prior existence of a similar class of methods used in telecommunications modelling, known as RESTART [35, 36, 37, 38, 39].

FFS, like TIS, uses a series of interfaces between the initial and final states to calculate the transition rate and to sample the transition path ensemble. These interfaces are defined by an order parameter λ : the initial (A) state is defined by $\lambda < \lambda_A = \lambda_0$, the final (B) state by $\lambda > \lambda_B = \lambda_n$, and the remaining interfaces are defined by intermediate values of λ : $\{\lambda_i \dots \lambda_{n-1}\}$. The method requires that $\lambda_{i+1} > \lambda_i$ for all i , and that any trajectory from A to B passes through each interface in turn. This places no restriction on the trajectories, which are free to loop back to recross previous interfaces any number of times. Like TIS, FFS uses the effective positive flux expression for the rate constant, Eqs. (6) and (7). However, FFS differs fundamentally from TIS in the manner in which the conditional probabilities $P(\lambda_{i+1}|\lambda_i)$ and the transition paths are computed. In FFS, the system dynamics are integrated forward in time only, eliminating the requirement for detailed balance.

Broadly speaking, FFS works by “capitalising” on fluctuations of the system dynamics in the direction of the order parameter. When the system undergoes a fluctuation that reaches the first interface, its configuration is stored. This stored configuration is then used as the starting point for repeated “trial runs”, to evaluate the probability that the system will reach the next interface. These trial runs are continued until the system reaches either the next interface (a “success”), or returns to the A state (a “failure”). The end points of successful trials are used to initiate new trial runs, to the subsequent interface. The result is that the system is driven in a ratchet-like manner from the initial to the final state, without imposing any bias on the microscopic dynamics. The probabilities $P(\lambda_{i+1}|\lambda_i)$ of Eqs. (6) and (7) are obtained from the

fraction of successful trial runs at each interface, and these can be multiplied by the flux of trajectories crossing the first interface to obtain the rate constant k_{AB} . A correctly weighted collection of transition paths is obtained by tracing back trial runs from the final state to the initial state [34]. Because each trial run starts from the final point of a previous trial run, the correct system dynamics is preserved along the whole transition path. We note that although FFS does assume that transitions between A and B are uncorrelated and that the rate k_{AB} is time-invariant, there is no requirement for the B state to be stable. Although FFS has generally been used for systems with stable A and B states, it can also be used to predict the probability of rare fluctuations from a stable A state, along a chosen order parameter, to an arbitrarily positioned end point λ_n [40]. It is also important to note that FFS is a static sampling technique, in which each new transition path is generated from scratch. This is in contrast to dynamic methods such as TPS and TIS, in which new transition paths are generated by modifying already existing ones. The advantage of static methods is that they generate uncorrelated samples, making them likely to explore a wider range of path space. However, static methods may also waste computational effort by repeatedly sampling blind alleys.




Figure 1. Schematic illustration of the initial simulation in the A state. Configurations corresponding to the points in the trajectory marked with crosses are stored.

Within the FFS framework, various different protocols have been proposed for firing trial runs and storing configurations at the interfaces. All these variants begin with a simulation in the A state, illustrated schematically in Figure 1. This simulation is used to compute the flux $\bar{\Phi}_{A,0}/\bar{h}_A$ of trajectories out of the A state [Eq.(6)], as well as to obtain a sample of configurations corresponding to crossings of interface λ_0 , to act as starting points for the trial run procedure. The system is initiated in the A state, and (after an initial equilibration period) the time evolution of the order parameter λ is monitored. When the system crosses λ_0 coming from A (*i.e.* in the direction of increasing λ), a counter is incremented, and the system configuration is stored. The simulation is then continued until N_0 configurations have been stored. The flux $\bar{\Phi}_{A,0}/\bar{h}_A$ is obtained by dividing the number of crossings N_0 by the total simulation time (which includes time

spent on fluctuations away from the A state but does not include time spent in the B state [34]).

The FFS variants differ in the way the probability $P(\lambda_n|\lambda_0)$ is computed. In our original presentation of FFS, Refs [33] and [34], we described three different variants. These are outlined below. Since Ref [34], more advanced versions of some of these algorithms have been proposed [41, 42, 40]. These advances are discussed in Sections 4, 5 and 6.

3.1. FFS variants




Figure 2. Schematic illustration of the DFFS method. An ensemble of branched transition paths is generated simultaneously by firing M_i trial runs from randomly chosen configurations at each interface in turn. The different colours denote the trial runs fired in the different stages of the procedure (in the order red, blue, green, purple).

3.1.1. “Direct” FFS The original version of FFS [33, 34] has subsequently been termed “Direct-FFS”, or DFFS [41]. In this algorithm, many transition paths are generated simultaneously, using the following procedure (illustrated in Figure 2):

- (i) Carry out a simulation in the A basin to generate a collection of N_0 configurations corresponding to crossings of interface λ_0 , as well as an estimate of the flux $\overline{\Phi}_{A,0}/\overline{h}_{A,0}$.
- (ii) Choose a configuration from this collection at random and use it to initiate a trial run which is continued until it either reaches λ_1 or returns to λ_0 . If λ_1 is reached, store the end point of the trial run. Repeat this M_0 times, each time choosing a random starting configuration from the collection at λ_0 . Compute $P(\lambda_1|\lambda_0)$ from the fraction of successful trial runs.
- (iii) Repeat step (ii) using the stored configurations at λ_1 to initiate M_1 trial runs to λ_2 or back to λ_0 . Generate a new collection of configurations at λ_2 from the end points of successful trials. Estimate $P(\lambda_2|\lambda_1)$ from the fraction of successful trials.
- (iv) Repeat until λ_n is reached.

Statistical errors can be computed by repeating the DFFS procedure several times or using analytical expressions (see Section 4). DFFS is straightforward to implement and its computational efficiency is rather robust to the choice of parameters (see Section 4), although it does require the storage of many configurations at each interface, which may be an issue for large-scale simulations. In order to extract transition paths from a DFFS simulation, one needs to record the connectivity history of all the trial runs. This allows one to piece together *a posteriori* complete transition paths from the full set of stored trial runs. Finally, it is important to note that the transition paths generated by DFFS are branched: many paths may start from a single configuration at λ_0 .




Figure 3. Schematic illustration of the BG algorithm. Branched transition paths are generated by firing k_i trials from the end points at λ_i of all successful trials from λ_{i-1} . The colours denote different branched paths, which are generated sequentially.

3.1.2. Branched Growth The “Branched Growth” (BG) algorithm [34] generates branched transition paths from A to B one at a time, rather than simultaneously as in DFFS. The algorithm proceeds as follows (see Figure 3):

- (i) Evaluate $\overline{\Phi}_{A,0}/\overline{h}_A$ and generate configurations at λ_0 using a simulation in the A basin.
- (ii) For the first configuration at λ_0 , fire k_0 trial runs, which are continued until λ_1 or λ_0 is reached. Store the end configurations of all successful trial runs.
- (iii) From each of these stored points at λ_1 , initiate k_1 trial runs to λ_2 or back to λ_0 . Store the end points of successful trials.
- (iv) Iterate this procedure over all subsequent interfaces until B is reached, or until no trials are successful at a given interface. This generates a branching tree of paths all starting from the same configuration at λ_0 . Estimate $P(\lambda_{i+1}|\lambda_i)$ as the total number of trials to reach B, divided by the total possible number $\prod_{i=0}^{n-1} k_i$.
- (v) Repeat steps (ii) to (iv) for subsequent configurations at λ_0 and average the estimate for P_B over many path generations (note that any zero values should be included in the average).

The BG method can easily be coded as a recursive algorithm, and has the potential advantage that it does not require storage of large numbers of configurations at each interface, and that extraction of transition paths is simpler than for DFFS (since these are generated one at a time). The BG method is however rather sensitive to parameter choice (see Section 4). If the number of trial runs per interface is too large, the method generates highly branched transition paths, so that sampling the later interfaces is computationally expensive. If too few trials are chosen per interface, few paths succeed in reaching the later interfaces. However, Borrero and Escobedo have proposed a method for automatic optimisation of the parameters [42], which is discussed in Section 4. The same authors have used the BG method as the basis for the FFS-LSE method [41], discussed in Section 5.3.




Figure 4. Schematic illustration of the RB method. Unbranched transition paths are generated one at a time by firing k_i trials from one randomly chosen end point at λ_i from the successful trials fired from λ_{i-1} .

3.1.3. The “Rosenbluth-like” method The Rosenbluth-like (RB) FFS variant allows the sequential generation of unbranched transition paths [34]. It draws on a close analogy between FFS and the sampling of polymer configurations in Monte Carlo simulations. In a polymer simulation, one seeks to grow a new polymer chain, monomer by monomer, in an environment crowded with other polymers. This is analogous to the interface-by-interface growth of a transition path in FFS. The RB FFS variant is an application of the Rosenbluth polymer sampling method [43, 1] to rare event simulations. The procedure is as follows:

- (i) Evaluate $\overline{\Phi}_{A,0}/\overline{h}_A$ and generate configurations at λ_0 using a simulation in the A basin.
- (ii) For the first configuration at λ_0 , fire k_0 trial runs to λ_1 or λ_0 . Store the end points of successful trials.
- (iii) Choose at random one of the stored configurations at λ_1 . Use this as the starting point for k_1 trial runs to λ_2 (or back to λ_0).

- (iv) Iterate this procedure over all interfaces until B is reached, or until no trials are successful.
- (v) Repeat steps (i) to (iv) for successive configurations at λ_0 .
- (vi) Compute $P(\lambda_{i+1}|\lambda_i)$ for each interface using a weighted average as described below.

In the RB method, the paths that are generated do not all have equal statistical weight. The weight $w_{i,b}$ of path b from A to λ_i is given by

$$w_{i,b} = \prod_{j=0}^{i-1} S_{j,b}/k_j \quad (8)$$

where $S_{j,b}$ is the number of successful trials fired at interface j during the generation of path b . To compute the probabilities $P(\lambda_{i+1}|\lambda_i)$, a weighted average is needed for each interface:

$$P(\lambda_{i+1}|\lambda_i) = \frac{\sum_b w_{i,b} S_{i,b}/k_i}{\sum_b w_{i,b}} \quad (9)$$

Here, the index b labels a specific path leading from A to a configuration at interface λ_i , and $S_{i,b}$ is the number of successful trials fired from that configuration to λ_{i+1} . When sampling over many transition paths, both the numerator and denominator of Eq.(9) will become large. In this case, Eq.(9) may become unsatisfactory, and one may prefer to reweight the paths using a Metropolis acceptance/rejection procedure [1], or alternatively a “waste-recycling” scheme [44] at each interface. This is described in detail in Ref [34]: however, we prefer simply to use Eq.(9), since the approaches described in Ref.[34] are rather complicated to code. When analysing the properties of the transition paths generated in the RB method, it is also necessary to include the weighting factor w_n .

The major advantage of the RB method is that the resulting transition paths are unbranched and sequentially generated, making them easy to extract and analyse. The RB method is also easy to code as a recursive algorithm.

3.2. Requirement for stochastic dynamics

FFS requires the dynamics of the system to be stochastic, since for deterministic dynamics all trial runs fired from a given configuration at interface λ_i would be identical. However, this requirement allows for a wide range of possibilities, including kinetic Monte Carlo simulations of chemical reaction networks [33, 45, 46, 47], lattice models with Monte Carlo spin flips [48, 49, 50] and particle-based Brownian dynamics or Monte Carlo simulations [34, 51, 52]. FFS is not suitable for use with completely deterministic Molecular Dynamics (MD) algorithms; however, it has been successfully applied to MD simulations by including a weak coupling to a stochastic Lowe-Andersen thermostat, which preserves the momentum of the system [53, 54, 55]. Given that MD trajectories with infinitesimally different initial conditions diverge within a few picoseconds due to the Lyapunov instability, this is unlikely to constitute a severe perturbation to the system dynamics.

3.3. Pruning

FFS as described above requires that trial runs fired from interface λ_i be integrated until they reach λ_{i+1} or return *all the way back to A*. If transition paths are short, the computational cost of integrating failed trials back to λ_0 is likely to be rather low. However, in some cases, such as diffusive barrier crossings or intermediate metastable states, it may be expensive to integrate all the way back to λ_0 . In these cases one can use a pruning scheme in combination with any of the above FFS variants. Here, trial runs from λ_i are integrated only as far back as some pre-defined value $\lambda_p < \lambda_i$; typically, $\lambda_p = \lambda_{i-1}$. With some probability p_p , a trial run which reaches λ_p is terminated and considered to have failed. If, with probability $(1 - p_p)$, the trial run is not terminated, then its statistical weight is increased by a factor $f_p = 1/(1 - p_p)$. This requires minor modifications to be made to the three FFS algorithms described above, to include differential weights for configurations at λ_i . These are described in detail in Ref. [34]; however, we did not find that pruning produced much improvement in computational efficiency for the examples tested (a model genetic switch and a simple representation of polymer translocation) [34]. This approach may nevertheless prove useful for other problems.

4. Computational efficiency: prediction and optimisation

An essential requirement for a rare event simulation method is that it should provide the rate constant and transition path ensemble more efficiently than brute force simulation. Defining, quantifying and optimising the computational efficiency of such methods is therefore an important task. FFS involves a number of parameters: the number and position of the interfaces, the number of trials fired from each interface, the number of configurations stored at the first interface (for DFFS), as well as the choice of which FFS variant to use and the definition of the order parameter. The computed rate constant and transition paths should not (in principle) depend on any of these choices, but they will affect the efficiency. We use the simple definition for the computational efficiency \mathcal{E} [56]:

$$\mathcal{E} = \frac{1}{\mathcal{C}\mathcal{V}} \quad (10)$$

where \mathcal{C} is the computational cost (in simulation steps) of calculating the rate constant, and \mathcal{V} is the statistical variance in the result, normalised by the square of its mean. A slightly different, but equivalent, expression was used by Van Erp in his analysis of the efficiency of TPS/TIS in comparison to Bennett-Chandler-type methods [57].

4.1. Analytical expressions for the efficiency

Analytical expressions for the efficiency of a method [57, 58] are useful for several reasons. Firstly, they allow one to estimate, before beginning a lengthy calculation, how much effort will be required to obtain a desired level of accuracy. Secondly, they allow the

estimation of error bars on a computed result, where it is not feasible to repeat the calculation. Thirdly, one can use the analytical expressions to optimise the efficiency of the method with respect to parameter choice. It is possible to derive approximate expressions for the efficiency of the FFS methods discussed above, as a function of the number of interfaces, the probability of success at each interface, and the number of trials at each interface. Here, we simply present the key results; more details including derivations can be found in Ref. [58].

For the subsequent discussion, it is important to note that in the BG and RB variants, the parameter defining the number of trials is k_i , the number of trials per configuration at λ_i . For DFFS, the relevant parameter is M_i , the *total* number of trials fired at interface i . To simplify our discussion, however, we shall use the notation k_i also for DFFS, but in this context we define it as $k_i \equiv M_i/N_0$. We shall also simplify our notation for the probabilities, defining $p_i \equiv P(\lambda_{i+1}|\lambda_i)$ and $q_i \equiv 1 - p_i$.

The computational cost \mathcal{C} can be estimated by assuming that the average length of a trial run from λ_i to λ_j is linearly proportional to $|\lambda_i - \lambda_j|$, so that the average cost C_i of a trial run fired from λ_i is [58]:

$$C_i = Q[p_i(\lambda_{i+1} - \lambda_i) + q_i(\lambda_i - \lambda_A)] \quad (11)$$

where Q is a constant. Eq.(11) can be used to write down expressions for \mathcal{C} for the three FFS variants discussed in Section 3; these expressions differ because the variants differ in the average number of trials fired per starting point at λ_0 . The results are:

$$\mathcal{C}^{DFFS} = R + k_0 C_0 + \sum_{i=1}^{n-1} \left[k_i C_i \prod_{j=0}^{i-1} (1 - q_j^{N_0 k_j}) \right] \approx R + \sum_{i=1}^{n-1} k_i C_i \quad (12)$$

$$\mathcal{C}^{BG} = R + k_0 C_0 + \sum_{i=1}^{n-1} \left[k_i C_i \prod_{j=0}^{i-1} p_j k_j \right] \quad (13)$$

and

$$\mathcal{C}^{RB} = R + k_0 C_0 + \sum_{i=1}^{n-1} \left[k_i C_i \prod_{j=0}^{i-1} (1 - q_j^{k_j}) \right] \quad (14)$$

where the cost is defined per starting configuration at λ_0 and R is the cost of generating such a starting configuration. These expressions take into account the fact that if no trials are successful at a given interface, the FFS algorithm will not make it to later interfaces.

The relative variance \mathcal{V} in the computed rate constant is assumed to arise only from the computation of P_B and not from the initial flux calculation. This is justified as long as the initial flux is large enough, and the initial simulation run in the A basin is long enough. The key assumption made in calculating \mathcal{V} is that trial runs at subsequent interfaces are uncorrelated. This allows us to treat the number of successful trial runs from interface i as a binomially distributed random variable, with parameter p_i . Taking into account the details of the different sampling protocols, we arrive at:

$$\mathcal{V}^{DFFS} = \sum_{i=1}^{n-1} \frac{q_i}{p_i k_i} \left[\frac{1}{\prod_{j=0}^{i-1} (1 - q_j^{N_0 k_j})} \right] \approx \sum_{i=1}^{n-1} \frac{q_i}{p_i k_i} \quad (15)$$

and

$$\mathcal{V}^{BG} = \sum_{i=1}^{n-1} \frac{q_i}{\prod_{j=0}^i p_j k_j}. \quad (16)$$

An equivalent expression was also derived in Ref. [58] for the RB method with Metropolis reweighting, but it is rather complicated. A much simpler expression can be derived if the p_i values are instead computed using Eq.(9); this will be discussed in a forthcoming publication [59].

Expressions (15) and (16) assume that the probability of success p_i is the same for all trial runs fired from interface i . In reality, however, some configurations at λ_i will have higher probability of success than others. This can be included in the expressions for \mathcal{V} by assuming that the p_i s have an intrinsic, “landscape variance” U_i . This leads to minor modifications to the results: for details see Ref. [58]. Interestingly, the three FFS variants cope differently with this landscape variance. Because the DFFS and BG methods produce branched paths, they sample many configurations at interfaces $i > 0$ as the number of trials k_i becomes large. This makes them insensitive to the values of the landscape variance U_i for $i > 0$. By contrast, in the RB method, where the paths are not branched, only one configuration is sampled per interface per path, so that all the landscape variance values U_i contribute to the total variance \mathcal{V} .




Figure 5. Predicted computational efficiency \mathcal{E} for the three FFS variants DFFS (solid lines), BG (dotted lines) and RB (with Metropolis acceptance/rejection; dashed lines), for a hypothetical rare event problem with evenly spaced interfaces and $p_i = p$, with $P_B = 10^{-8}$, $Q = R$, $U_i = 0$ and $N_0 = 1000$. (a): \mathcal{E} as a function of the number of trials k , with 5 interfaces. (b): \mathcal{E} as a function of the number of interfaces n , for $k = 25$. [Reproduced with permission from Ref [58]].

The efficiency \mathcal{E} is obtained by substituting Eqs. (12-16) into Eq.(10). Figure 5 shows \mathcal{E} plotted as a function of the number of trials k [panel (a)] and the number of interfaces n [panel (b)] for a hypothetical rare event problem with evenly spaced interfaces and $p_i = p$ for all interfaces, with $P_B = 10^{-8}$, $Q = R$ and $N_0 = 1000$. The striking result is that while the BG method is rather sensitive to the choice of either n or k , both the DFFS and RB methods are insensitive to both these parameters, as long as k and n are large enough. This implies that the computational cost associated with having many interfaces or firing many trial runs is balanced by a proportional gain in

statistical accuracy. Of course, this analysis cannot be taken to extremes: a very large number of interfaces will have associated overhead costs, as well as leading to strong correlations between successive interfaces, neither of which is taken into account in this analysis.




Figure 6. Predicted and measured efficiency \mathcal{E} for the Maier-Stein system, simulated with overdamped Langevin dynamics [for details see Ref [58]]. Lines: predicted efficiency. Solid: DFFS, dotted: RB (with Metropolis acceptance/rejection), dashed: BG. Symbols: measured efficiency. Circles: DFFS, triangles: RB (Metropolis), squares: BG. (a): \mathcal{E} as a function of number of trials k (b): \mathcal{E} as a function of number of interfaces n . [Reproduced with permission from Ref [58]].

Figure 6 shows the predicted efficiency \mathcal{E} , compared with the actual value measured for FFS simulations of the two-dimensional nonequilibrium rare event problem proposed by Maier and Stein [60, 61, 62]. This consists of a particle moving with overdamped Langevin dynamics in a force field that is not the gradient of a potential field (for details see Ref. [58]). The values of R , Q , $\{p_i\}$ and $\{U_i\}$ were extracted from the FFS simulations and used as inputs to the analytical expressions. Not only do the results for the Maier-Stein system closely resemble the trends in Figure 5 for the hypothetical problem, but the agreement between the analytical results and the simulations is extremely good. However, such good agreement cannot be relied on in all cases: for a model genetic switch, where correlations between successive interfaces are more likely, differences of up to a factor of 10 between the analytical predictions for \mathcal{E} and the simulation results were observed [58].

4.2. Optimising the efficiency

Borrero and Escobedo [42] have shown how these analytical expressions can be used to optimise the parameters in FFS simulations, for the DFFS and BG schemes. They adopt two complementary approaches: (i) optimising the number of trial runs $\{k_i\}$ for a fixed set of interfaces, and (ii) optimising the positioning of the interfaces $\{\lambda_i\}$, for fixed $\{k_i\}$.

4.2.1. Optimising the number of trial runs For a given set of interfaces, the optimum values for the number of trials $\{k_i\}$ [or for DFFS, $\{M_i\} \equiv \{N_0 k_i\}$], can be found by

minimising analytically the variance \mathcal{V} with respect to the $\{k_i\}$ (or $\{M_i\}$). The cost \mathcal{C} is constrained to a fixed value using a Lagrange multiplier. This leads to implicit expressions for the optimum $\{k_i\}$ ((or $\{M_i\}$) for the BG and DFFS schemes [42]. For DFFS, under the assumption that the M_i values are large, so that $q_i^{M_i} \approx 0$, this expression is

$$M_i = \frac{P_B}{\sqrt{(\alpha)}} \left(\frac{1-p_i}{p_i} \right)^{1/2} \left(\frac{C_i}{Q} \right)^{-1/2} \quad (17)$$

where α is the Lagrange multiplier that sets the cost, and C_i and Q are as defined in Section 4.1, so that $C_i/Q = [p_i(\lambda_{i+1} - \lambda_i) + q_i(\lambda_i - \lambda_A)]$. A simple practical procedure is then prescribed to obtain the optimum $\{M_i\}$ set, for a fixed cost, for DFFS:

- (i) Set one of the M_i values (e.g. M_0).
- (ii) Compute the other M_i s using:

$$\frac{M_{i+1}}{M_i} = \left(\frac{p_i(\lambda_i - \lambda_A)}{p_{i+1}(\lambda_{i+1} - \lambda_A)} \right)^{1/2} \quad (18)$$

(this assumes that p_i is small; if not the expression is slightly more complicated).

- (iii) Compute the cost associated with this $\{M_i\}$ set from Eq.(12).
- (iv) Iterate to obtain a set of M_i s corresponding to the desired cost.

For the BG method, the equivalent expression to Eq.(17) is more complicated but a similar principle applies; for details see Ref [42].

4.2.2. Optimising the interface positions For a fixed set of M_i (or k_i values), the efficiency can be optimised with respect to the positions of the interfaces $\{\lambda_i\}$, for $0 < i < n$. Borrero and Escobedo assume that the computational cost is fixed by the $\{M_i\}$ (or $\{k_i\}$) [42], and minimise the variance \mathcal{V} with respect to the probabilities $\{p_i\}$, with the constraint that $P_B = \prod_{i=0}^{n-1} p_i$ remain constant. This leads to the intuitive result that for optimum interface placement, there should be a constant flux of partial path trajectories across all interfaces. This implies that the product $M_i p_i$ (for DFFS), or $k_i \prod_{j=0}^{i-1} k_j p_j$ (for BG), should be constant across all interfaces. Since the $\{M_i\}$ (or $\{k_i\}$) are fixed, this specifies the optimum values for the probabilities $\{p_i\}$, which can be achieved by a suitable placement of the interfaces. To translate between $\{p_i\}$ values and the interface positions $\{\lambda_i\}$, one needs an interpolation function $f(\lambda_i)$, for which one choice is:

$$f(\lambda_i) = \frac{\sum_{j=0}^{i-1} \ln p_j}{\sum_{j=0}^{n-1} \ln p_j} \quad (19)$$

The optimisation procedure then consists of:

- (i) Run FFS with an, as yet, non-optimal set of interfaces $\{\lambda_i\}$, to compute the function $f(\lambda_i)$.
- (ii) Compute the optimum $\{p_i\}$ set by demanding constant flux across all interfaces.

- (iii) Use the function $f(\lambda_i)$ to determine the interface placement corresponding to these optimal $\{p_i\}$ values.
- (iv) Iterate the procedure if necessary.

A separate methodology for optimising the position of the starting interface λ_0 has been proposed by Velez-Vega *et al* [55].

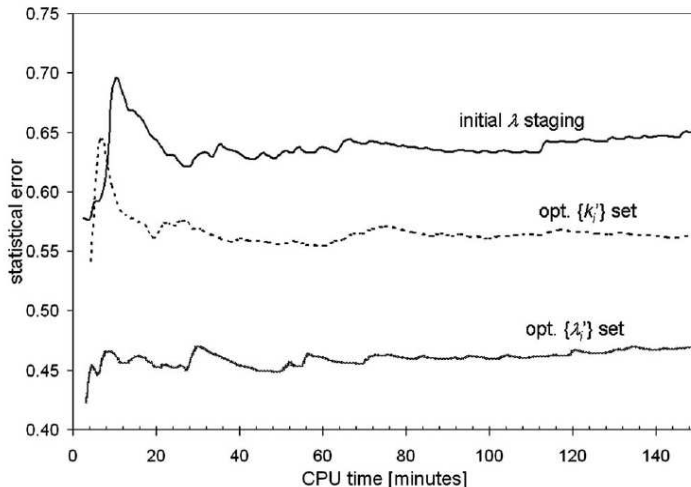


Figure 7. Statistical error in the estimated rate constant for BG FFS simulations of a two dimensional test potential with 4 interfaces, for an initially unoptimised set of interfaces (“initial λ staging”), for the same set of interfaces with optimised number of trials (“opt. $\{k'_i\}$ set”), and for an optimised set of interface positions (“opt. $\{\lambda'_i\}$ set”). Smaller statistical errors are achieved for the same CPU time with the optimised parameters. [Reproduced with kind permission from Ref [42]].

4.2.3. Efficiency gains Borrero and Escobedo demonstrate their efficiency optimisation procedure for a simple two dimensional test potential, for the flipping of a model genetic switch (see Section 7) and for the folding of a lattice protein model [42], obtaining impressive results, as illustrated in Figure 7. For these examples, a single iteration proved to be enough to converge the optimisation for either the number of trials or the interface positions.

5. The order parameter and the committor

FFS relies on the definition of the order parameter λ , which must be some coordinate of the system that increases during the transition from A to B. FFS does not assume that λ corresponds to the true “reaction coordinate”, which is the actual route through phase space followed by the transition path ensemble. However, choosing a good order parameter (one which is close to the true reaction coordinate) will increase the computational efficiency, while a poor choice of order parameter will lead to wasted

effort, as many of the configurations generated at interface λ_i will have little chance of reaching λ_{i+1} . This may even lead to incorrect results if the number of paths sampled is small [63]. In this section, we first discuss how to define the order parameter in interface-based methods such as FFS. We then discuss the committor function, which contains information about the reaction mechanism, and which corresponds to the “ideal” choice of order parameter. Finally, we briefly review several methods which have recently been developed for extracting the reaction coordinate from measured committor values, focusing particularly on the FFS-LSE method of Borrero and Escobedo [41].

5.1. Defining the order parameter

For some rare event problems, it is easy to define a good order parameter. For example, for crystal nucleation processes, one typically chooses the number of particles in the system that are “crystalline” [64], for a polymer translocation problem one can use the number of translocated monomers [34], and for a bistable chemical reaction in which the transition is between states rich in molecules of chemical species A and B, one can use the difference between the number of A and B molecules [33, 34, 45, 46, 47]. However, in other cases the choice of order parameter is less obvious. For example, for hydrophobic polymer collapse the solvent coordinates as well as those of the monomers can play an important role in the transition [65]. Complex reaction coordinates are also a common feature of protein folding problems [66]. Figure 8 illustrates a simple two-dimensional potential energy landscape for which the Z-shaped reaction coordinate cannot be described by either of the two coordinates of the system (x and y), or by any linear combination of these coordinates [67].

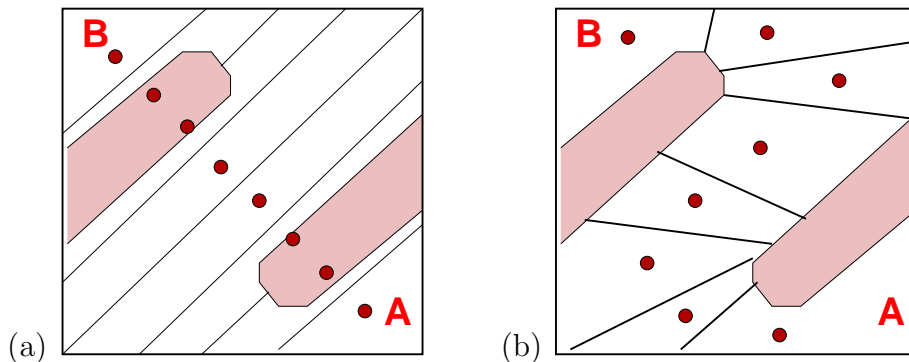


Figure 8. A two-dimensional potential energy landscape for which the reaction coordinate is Z-shaped and cannot be described by a simple linear combination of the x and y coordinates of the system. (a): Partitioning of space with Voronoi polyhedra around a linear string of beads between the A and B states; this gives interfaces which follow the reaction coordinate poorly. (b): Partitioning into interfaces using a curved string of beads; this gives a much better set of interfaces.

For complex cases such as the one illustrated in Figure 8, the Voronoi tessellation approach pioneered by Vanden-Eijnden and Venturoli [32] can prove very useful. Here,

one defines a path connecting A and B by a set of configurations, or “beads”. The path need not be describable in terms of a single collective coordinate. Any configuration of the system can be classified according to which of the beads it lies closest to: this is equivalent to partitioning space into a set of Voronoi polyhedra around the beads. Interfaces can then be defined as the planes in phase space across which the “closest bead identity” changes, or equivalently the edges of the Voronoi polyhedra. Figure 8 shows the interfaces defined by this method for our two-dimensional example, for a linear set of beads [panel (a)] and for a set of beads chosen to lie along the curved reaction coordinate [panel(b)] [59]: it is clear that Voronoi tessellation provides a very convenient and easy-to-implement way to translate a string of beads into a set of interfaces. It is important to note, however, that for highly multidimensional problems it is likely to be necessary to project the phase space onto a small number of order parameters before carrying out the Voronoi tessellation, otherwise the resulting interfaces may be highly convoluted. This approach does not therefore entirely eliminate the need for order parameters. In addition, of course, one still needs to find a suitable set of beads, for example by an iterative technique such as the finite temperature string method (FTS) [31, 32]. In Section 6.3, we discuss briefly how this is done in the context of nonequilibrium Umbrella sampling [67].

5.2. The committor and the reaction coordinate

The committor function $P_B(x)$ is defined as the probability that a trajectory initiated from configuration x will reach the final state B before the initial state A. Along a transition path, the committor function increases from zero to unity. Configurations along the transition paths for which $P_B = 0.5$ have special significance: the collection of these configurations is known as the “transition state ensemble”, or TSE (although alternative definitions are also possible [68, 69]). Analysis of the TSE configurations can provide insight into the reaction mechanism. If the probability distribution for a given order parameter, evaluated over the TSE configurations, is highly peaked, it is likely that this order parameter closely corresponds to the reaction coordinate, whereas a broad or bimodal distribution indicates that other order parameters are needed to fully describe the transition mechanism [2, 3]. Scatter plots for the TSE configurations as functions of various collective coordinates can also provide insight into the reaction mechanism, as discussed in Section 7. To extract committor values for configurations along the transition paths, one typically fires a large number of trial runs from each configuration to estimate the probability of these reaching B rather than A. This is a computationally expensive procedure (although some effort can be saved if one is only interested in the TSE). In recent work, however, Borrero and Escobedo [41] have shown that committor values can be extracted on-the-fly from FFS simulations, by making intelligent use of the information already obtained on the number of successful trials to interface λ_{i+1} for each configuration at interface λ_i . This is discussed in Section 5.3.

The committor function P_B is in some sense the ideal reaction co-ordinate, since it

by definition correlates with the progress of the transition. However, P_B is a complex function of all the coordinates of the system. To obtain scientific insight, one needs to be able to project this function onto a small set of physically meaningful collective coordinates. Hummer *et al* [68, 69] proposed a variational method for optimising reaction coordinates, based on evaluating the projection onto the order parameter of the probability function $p(TP|x)$ that a configuration x forms part of a transition path. Ma and Dinner [70] proposed a method in which one uses “representative” configurations corresponding to different values of the committor P_B to determine the functional dependence of the committor on each of a chosen set of collective coordinates. An optimisation procedure (in this case a genetic algorithm) can then be used to find the best combination of these coordinates to represent the committor. In related work, Peters *et al* [71, 72] also proposed a method for determining the optimum combination of collective coordinates to represent the committor. In their method, the committor values are obtained on-the-fly using a version of transition path sampling called Aimless Shooting. A simple model for the reaction coordinate (eg a linear combination of collective coordinates) is assumed and the parameters of the model optimised using Bayesian likelihood maximisation.

5.3. Using the committor to optimise the order parameter in FFS




Figure 9. Isocommittor lines obtained using FFS-LSE for a simulation on the two dimensional potential energy landscape represented by the contour plot. The numbers indicate the committor values. [Reproduced with kind permission from Ref [41]].

In recent work, Borrero and Escobedo have proposed a method (related to that of Peters *et al* [71, 72]) in which information on the committor is extracted directly from FFS simulations, and used to optimise the choice of order parameter [41]. This

method is known as FFS-LSE (least square estimation). Borrero and Escobedo show that for the branched growth variant of FFS (see Section 3.1), the committor value P_{Bj}^i for configuration j at interface λ_i is given by:

$$P_{Bj}^i = p_j^i(\lambda_{i+1}|\lambda_i) \frac{\sum_{m=1}^{N_j^i} P_{Bm}^{i+1}}{N_j^i} = \frac{\sum_{m=1}^{N_j^i} P_{Bm}^{i+1}}{k_i} \quad (20)$$

where k_i is the number of trials per configuration at interface i , N_j^i is the number of successful trials from configuration j at interface i and $p_j^i(\lambda_{i+1}|\lambda_i) = N_j^i/k_i$ is the probability of success for configuration j at interface i . This equation states that the committor value for configuration j is the probability that a trial run fired from this configuration will reach the next interface λ_{i+1} , multiplied by the average committor value for its “daughter” configurations at λ_{i+1} . If information on the connectivity between configurations at successive interfaces, as well as on the number of successful trials for each configuration, is stored during an FFS run, committor values for these configurations can be extracted at the end of the run with no additional computational effort. Although Eq.(20) was derived for the BG algorithm, it should be possible to apply a similar approach to the other FFS variants.

This procedure produces a set of configurations with associated committor values over the whole range of the order parameter λ . For these configurations, one evaluates a set of m candidate collective coordinates q . This data is then used to fit a parametrised functional form for the dependence of the committor on the q :

$$P_B(q) = \sum_{k=1}^m \beta_k q_k + q^T \mathbf{A} q + \beta_0 + \epsilon \quad (21)$$

where the β values and the matrix \mathbf{A} are optimised by least squares fitting to the data set obtained from FFS, and ϵ is the sum of the squares of the errors, to be minimised. The resulting functional form for $P_B(q)$ is the optimal choice for the FFS order parameter $\lambda(q)$. Choosing the order parameter $\lambda(q)$ that most closely matches the committor function (within the constraints of the fitting function) should lead to interfaces that are perpendicular to the transition paths, resulting in an efficient computation of the rate constant. Borrero and Escobedo found that a single iteration was sufficient to converge on the optimal $\lambda(q)$, for several test cases including the flipping of a model genetic switch and the folding of a lattice protein model [41]. Figure 9 shows the isocommittor lines obtained using this procedure, for a test simulation on the two dimensional potential surface represented by the contour plot. It might be interesting in future to combine this approach with a bead-string description of the order parameter as discussed in Section 5.1.

6. Computing stationary distributions

For many rare event problems, one is interested not only in the rate constant and transition paths, but also in the steady-state probability distribution ρ , as a function of one or more order parameters q . Knowledge of the steady-state distribution allows

one to compute, for example, averages of experimentally measurable observables for comparison with experiments. For systems with stable A and B states, $-\ln \rho(q)$ takes the form of a “barrier” with a peak separating the two stable states. If the dynamics of the system obeys detailed balance, this distribution is directly related to the free energy function F : $F(q) \sim -\ln \rho(q)$. For these equilibrium systems, $\rho(q)$ can be computed by umbrella sampling [1], in which one divides the range of q into a series of windows, runs a separate simulation in each window and uses the Boltzmann distribution to reassemble the probability distributions from each window into the unbiased $\rho(q)$. In this section, we first discuss how $\rho(q)$ can be extracted from FFS simulations, for equilibrium or nonequilibrium systems [73], and a related method for computing $\rho(q)$ with FFS, for equilibrium systems only [40]. We also briefly discuss a method which allows $\rho(q)$ to be computed for nonequilibrium systems in multiple dimensions [74, 67].

We note that for these computations there is no requirement for a stable B state. In addition to computing barrier heights, these methods can also be used to explore the probability distribution for excursions of the system from a stable A state, along an order parameter, without the presence of a stable B state. In this case, $-\ln \rho(q)$ will be an increasing function, whose detailed shape contains information on the physics of fluctuations away from the A state.

6.1. Obtaining stationary distributions from FFS simulations

The stationary distribution function $\rho(q)$ can be obtained on-the-fly during an FFS simulation, in a similar approach to that of Moroni *et al* for computing $\rho(q)$ in PPTIS simulations [22]. To achieve this, one computes histograms for the q values visited by all trial runs (failed and successful), fired from interface i . In the case of a stable B state, it is necessary to run the FFS calculation in both directions (A→B and B→A) [73].




Figure 10. Illustration of categories of trajectories contributing to the stationary distribution $\rho(q)$. Trajectories 1 and 2 originate in A and are sampled by an FFS simulation from A to B. Trajectories 3 and 4 originate in B and are sampled by an FFS simulation from B to A. [Reproduced with permission from Ref [73]].

The principle underlying the calculation of $\rho(q)$ with FFS is illustrated in Figure 10. Trajectories visiting a particular region of phase space can be grouped according to

their origin in either the A or the B state. The stationary distribution function can be written as the sum of the contributions, ψ_A and ψ_B respectively, of these two groups of trajectories:

$$\rho(q) = \psi_A(q) + \psi_B(q) \quad (22)$$

where ψ_A and ψ_B can be written as:

$$\psi_A = \rho_A \Phi_A \tau_+(q; \lambda_0) \quad \psi_B = \rho_B \Phi_B \tau_-(q; \lambda_n) \quad (23)$$

Here, ρ_A is the steady-state probability of finding the system in A, Φ_A is the flux of trajectories out of the A state and $\tau_+(q; \lambda_0)$ is the average time spent with a value q of the order parameter, for a trajectory which originates from λ_0 . The equivalent definitions hold for ρ_B and Φ_B , while $\tau_-(q; \lambda_n)$ is the equivalent average time, for a trajectory which originates from λ_n . These averages must be taken over all trajectories leaving λ_0 (or λ_n), whether or not they eventually reach B (or A), with the correct statistical weights.

In an FFS simulation from A to B, Φ_A is computed during the initial simulation in the A basin. The probabilities ρ_A and ρ_B can be obtained once the rate constants k_{AB} and k_{BA} have been computed, since in steady state $\rho_A k_{AB} = \rho_B k_{BA}$ and $\rho_A + \rho_B = 1$ (assuming transitions are fast compared to the time spent in the A and B basins). The function $\tau_+(q; \lambda_0)$ can be obtained if we measure during the FFS simulation the average time $\pi_+(q; \lambda_i)$ spent with order parameter value q , for a trial run fired from λ_i . This average should be computed over all configurations in all trial runs fired from λ_i , regardless of whether they succeed in reaching λ_{i+1} , and should include any differential weighting factor applied to the trial runs (e.g. in the case of the RB method or if pruning is used). $\tau_+(q; \lambda_0)$ is then given by:

$$\tau_+(q; \lambda_0) = \pi_+(q; \lambda_0) + \sum_{i=1}^{n-1} \pi_+(q; \lambda_i) \prod_{j=0}^{i-1} P(\lambda_{j+1} | \lambda_j) \quad (24)$$

Eq.(24) expresses the average time $\tau_+(q; \lambda_0)$ spent at order parameter value q as the sum of contributions from partial trajectories (trial runs) originating at each interface, weighted by the probability $\prod_{j=0}^{i-1} P(\lambda_{j+1} | \lambda_j)$ of observing such a partial trajectory in a brute-force simulation. As the interface index i increases, this probability decreases, but $\pi_+(q; \lambda_i)$ continues to be well-sampled. This is because FFS allows good sampling of regions of the phase space which are rarely visited in a brute-force simulation. A similar procedure can be applied in the B to A direction to obtain Φ_B , k_{AB} and $\tau_-(q; \lambda_n)$. A detailed description of the practical implementation is given in Ref [73]. We note that if the B state is very much less stable than the A state (or not stable at all), one can make the approximation $\rho_B \approx 0$, so that $\psi_B \approx 0$ in Eq.(22) and only a simulation in the A to B direction is required. It is also important to point out that q need not correspond to the FFS order parameter λ : one may even obtain the distribution as a function of multiple coordinates $\rho(q_1, q_2)$, simply by computing multidimensional histograms $\pi_{\pm}(q_1, q_2; \lambda_i)$ during the FFS simulation. However, for order parameters very different from λ , the advantage of enhanced sampling of rarely visited regions of the phase space will be lost.




Figure 11. Stationary distribution $\rho(x)$ (solid line) obtained using FFS, compared to the normalised Boltzmann distribution (circles) for a symmetric double well potential $V(x) = ax - bx^2 + cx^4$ with $a = 0.25$, $b = 2$, $c = 1$, simulated using overdamped Brownian dynamics with diffusion constant $D = 0.01$, $k_B T = 0.1$ and $dt = 0.05$. The dotted and dashed lines show $\psi_A(x)$ and $\psi_B(x)$ respectively [Reproduced with permission from Ref [73]].

Figure 11 shows $\rho(x)$ for a simple one-dimensional potential, computed using FFS and compared with the expected Boltzmann distribution. Similarly convincing results were obtained for the nucleation barrier in a two-dimensional Ising model [73], and for the nonequilibrium case of the flipping of a model genetic switch [73, 46].

6.2. Forward flux / umbrella sampling

The method described in Section 6.1 requires FFS simulations in both the forward and backward directions, if the steady-state population of the B state is significant. An approach developed recently by Borrero and Escobedo, known as Forward Flux / Umbrella Sampling (FFS-US) [40], aims to avoid this requirement, for systems where detailed balance is obeyed. In FFS-US, histograms obtained in a FFS simulation in the direction $A \rightarrow B$ are combined with histograms obtained with conventional umbrella sampling [1], using the interfaces as hard walls. For the umbrella sampling, trajectories are initiated inside each window using configurations obtained in the FFS simulation. The umbrella sampling histograms correct for the bias that arises in the FFS histograms due to the fact that one simulates only in the $A \rightarrow B$ direction.

6.3. Nonequilibrium umbrella sampling in multiple dimensions

A method has recently been proposed by Warmflash *et al* [74], and extended by Dickson *et al* [67], which allows for efficient computation of multidimensional steady state distributions for nonequilibrium rare event problems. Although this approach is distinct from FFS, we feel it is useful to include a brief description of it in this review,

since it is one of the few methods that give access to the steady state distribution for nonequilibrium systems.

In the method proposed by Warmflash *et al* [74], the phase space region of interest is divided into a lattice using one or several order parameters. Separate simulations are run in each lattice box simultaneously, during which one counts the numbers of simulation trajectories which attempt to transfer between boxes. When this happens, “weight” is transferred between boxes, and the trajectory is reinserted at the boundary of the same box with a configuration drawn from a self-consistently determined statistical distribution. To achieve this, one simulates at the same time a second lattice of boxes, with a grid offset relative to the first lattice. The second lattice provides configurations corresponding to interface crossings that can be used in the first lattice, and *vice versa* [for details see Refs [74] and [67]]. The stationary distribution is finally obtained from the simulated probability distributions within each box, multiplied by the weight computed for that box. This method, unlike FFS, is suitable for problems with slow dynamics in multiple dimensions. However, for many dimensions its implementation is likely to become complicated. The more recent work by Dickson *et al* [67] makes several modifications. Here, the order parameter instead consists of a string of beads, with the boxes now being defined by Voronoi polyhedra (see Section 5.1). A similar approach is used, except that one now simulates two strings, each offset from the other, and the algorithm for reinserting trajectories is somewhat different. Importantly, however, one can now use the computed distributions at the interfaces to iteratively improve the positioning of the beads, in an approach inspired by the string method [75, 76]. Although Refs [74] and [67] focus on the stationary distribution function $\rho(q)$, it seems likely that this approach could also give access to rate constants, in a similar manner to the Weighted Ensemble method of Huber and Kim [30].

7. Applications

Forward flux sampling has been applied to quite a number of different equilibrium and nonequilibrium rare event problems, with a variety of simulation techniques including Metropolis Monte Carlo, Molecular Dynamics, Brownian Dynamics and kinetic Monte Carlo. Transitions studied include nucleation in a variety of different contexts [48, 49, 50, 77, 78, 79, 80, 81, 82, 51, 83, 73, 84, 53], genetic switch flipping [33, 45, 46, 47], changes in DNA configuration [85], droplet coalescence [86], polymer translocation [34, 87, 88] and protein conformational changes [89, 55]. It is not our purpose here to discuss these applications in detail (an excellent review of biomolecular applications is given in [7]). Instead, we present a brief overview of three applications with which we have been involved, highlighting the contributions made using FFS as well as particular methodological challenges.

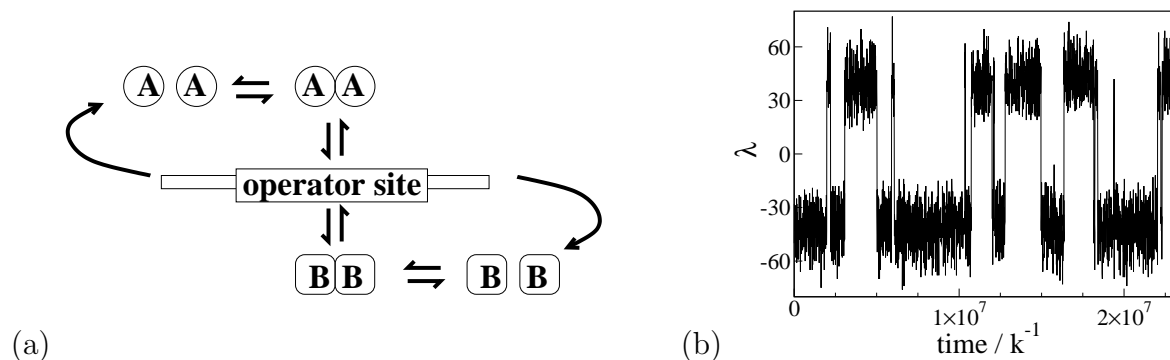


Figure 12. (a) Schematic illustration of the model genetic switch (b) A typical brute-force simulation trajectory. The order parameter $\lambda \equiv N_A - N_B$ is plotted as a function of time, where N_A and N_B are the total number of A and B molecules respectively. [Reproduced with permission from Ref [34]].

Reactions		Forward rate constant	Backward rate constant
$2A \rightleftharpoons A_2$	$2B \rightleftharpoons B_2$	$5k$	$5k$
$O + A_2 \rightleftharpoons OA_2$	$O + B_2 \rightleftharpoons OB_2$	$5k$	k
$O \rightarrow O + A$	$O \rightarrow O + B$	k	-
$OA_2 \rightarrow OA_2 + A$	$OB_2 \rightarrow OB_2 + B$	k	-
$A \rightarrow \emptyset$	$B \rightarrow \emptyset$	$0.25k$	-

Table 1. Reaction scheme for the model genetic switch. The unit of time is k^{-1} .

7.1. Genetic switch flipping

Gene regulatory networks control the behaviour of biological cells. In these networks, genes encode protein molecules which in turn control the expression of other genes. The resulting networks of interactions between genes allow cells to perform the computations necessary for survival and proliferation. Of particular interest are gene regulatory networks with multiple stable states, corresponding to alternative cellular developmental outcomes. A classical example is the bistable bacteriophage λ switch, which controls the transition between lysogeny (quiescent integration into the host cell) and lysis (replication and killing of the host cell) for phage λ , a virus which infects the bacterium *Escherichia coli* [90]. We studied a highly simplified representation of the gene network controlling this switch. In this model switch, two genes A and B mutually repress one other (see Figure 12a). If gene A is turned on, protein A is produced. This protein can dimerise and in the dimer form it can bind to the DNA and prevent the production of protein B from gene B . There is thus a stable state with a high concentration of protein A and a low concentration of protein B . If, however, due to a fluctuation, protein A dissociates from the DNA, then gene B can be expressed. The newly synthesised protein B can dimerise and bind to the DNA, preventing the expression of gene A .

This can ultimately lead to the flipping of the switch into the alternative stable state with a high concentration of protein B. Figure 12b shows a typical trajectory for a brute-force simulation of this model switch. The system undergoes infrequent but rapid flips between the A-rich and B-rich states. These random flips are driven by intrinsic biochemical “noise”: fluctuations due to the stochasticity of chemical reactions. Recent experiments have shown that this noise can play an important role in gene regulation [91]; yet it is known that the bacteriophage λ genetic switch is extremely stable, with a spontaneous flipping rate of less than once in every 10^9 generations [92]. This raises the question of what the mechanisms are that govern that stability of bistable genetic switches in the presence of biochemical noise. We have addressed this question using FFS.

We were interested how the detailed rules for DNA binding affect the switch stability: in particular, the comparison between a switch in which a dimer of protein A excludes the binding of B dimers to the DNA (and vice versa), with a switch in which both dimers can bind simultaneously [93, 94, 33]. The former is termed an “exclusive switch” and the latter a “general switch”. The exclusive switch is described by the set of chemical reactions given in Figure 12a. In this reaction scheme, O represents a DNA regulatory sequence (operator) adjacent to two divergently transcribed genes A and B , which code respectively for proteins A and B. These can dimerise to form A_2 and B_2 . Genes A and B can each produce proteins with the same rate, but whether they do so depends on the state of the operator O . When an A_2 dimer is bound to O , the production of B is blocked, and likewise, when a B_2 dimer is bound to O , the production of A is blocked. Proteins can also vanish (in the monomer form), modelling degradation and dilution in a cell. The reaction set for the general switch is identical except for the addition of the extra reactions $OA_2 + B_2 \rightleftharpoons OA_2B_2$ and $OB_2 + A_2 \rightleftharpoons OA_2B_2$ (*i.e.* the operator can bind both dimers simultaneously): when dimers are bound to the operator, neither protein can be produced.

We simulated this model using brute force simulation and DFFS, using the order parameter $\lambda \equiv N_A - N_B$, where N_A and N_B are the total number of A and B molecules respectively. Our results showed that the rules for operator binding can have a dramatic effect on the stability of the switch: for typical parameter values, the exclusive switch is orders of magnitude more stable than the general switch [93, 94, 33]. To understand this result, we extracted transition paths from our FFS simulations. Figures 13a and 13b show typical switching pathways for the general and exclusive switch, plotted in the $N_A - N_B$ plane [33]. The switching pathways are very different: during a typical flipping trajectory, the general switch passes through a region where the copy numbers N_A and N_B are both nearly zero, in contrast to the exclusive switch. To characterise the switching pathways further, we extracted configurations from the transition state ensemble (TSE). These configurations are plotted in the $N_A - N_B$ plane in Figures 13c and 13d, where the colour coding illustrates the operator state. These plots show that for the general switch, the TSE is dominated by configurations in which both proteins are bound to the operator and the production of both genes is repressed, which explains

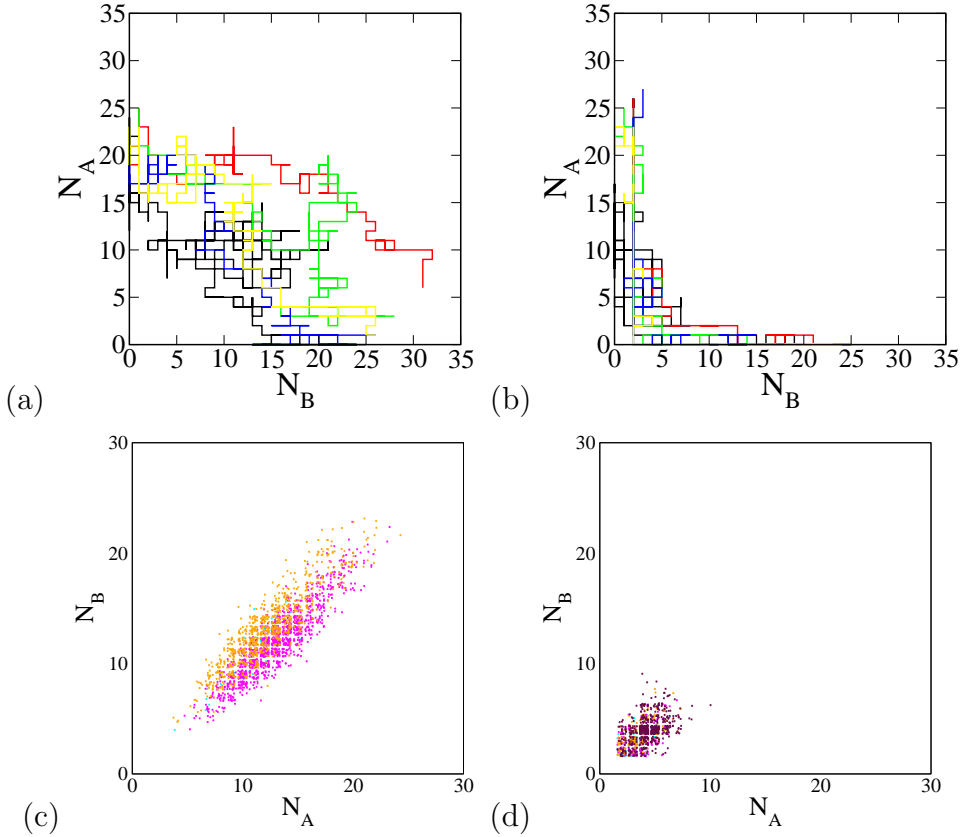


Figure 13. Top: Five randomly chosen transition paths, plotted in the $N_A - N_B$ plane, for (a) the exclusive and (b) the general switch. Each transition path is shown in a different colour. Bottom: TSE configurations plotted in the $N_A - N_B$ plane, and colour coded according to operator state, for (c) the exclusive switch and (d) the general switch. Cyan: O , gold: OA_2 , magenta: OB_2 , violet: OA_2B_2 . [Panel (a) reproduced with permission from Ref [33]].

why at the top of the “barrier”, N_A and N_B are both nearly zero. In the exclusive switch, however, these states are not allowed (since the proteins mutually exclude each other’s binding). Thus, in the transition paths, N_A and N_B never become nearly zero simultaneously. This explains why the exclusive switch is more stable than the general switch. Due to rare fluctuations, copies of the minority species will occasionally be produced. In the general switch, these can immediately bind the operator, leading to the repression of the majority species, which is the critical step that ultimately leads to the flipping of the switch. In the exclusive switch, however, a newly synthesised dimer of the minority species probably cannot bind the operator, since it is likely to be blocked by binding of the majority species. The system then has to wait for another fluctuation whereby the majority species dissociates from the operator before a flipping event can be initiated.

The model described above is a highly simplified representation of a real genetic switch. Recently, we have simulated a much more detailed model of the bacteriophage λ switch [47]. This system comprised over 500 chemical reactions, which were

computationally expensive to simulate, requiring the use of FFS as well as dynamical coarse-graining of some of the chemical equilibria (although importantly not of the operator binding reactions). Again using DFFS with the same choice of order parameter as for the simple model switch, we were able to reproduce the extreme stability of the switch, as observed experimentally. These simulations also revealed a key role for a DNA looping interaction (not present in the simple model), in maintaining the extraordinary stability of the switch.

These results show that rare event simulation methods such as FFS can successfully be applied to study rare events in complex biochemical networks whose dynamics is intrinsically out of equilibrium. Multistability has been found to play an important role in many different biological contexts, ranging from cell differentiation, apoptosis, the immune system, to the cell cycle. We hope that FFS will prove a useful tool for modelling these processes.

7.2. Homogeneous crystal/bubble nucleation

Phase transitions occurring by homogeneous nucleation are a widespread and important class of rare event processes. When a liquid is cooled below its melting point, the liquid state becomes metastable with respect to the crystal. The supercooled system can spend a long time in the liquid state before undergoing a rapid transformation into the thermodynamically stable solid phase. Similarly, a liquid that is heated above its boiling point can undergo a “cavitation” transition in which bubbles of the gaseous phase are formed. Since these nucleation processes are rare events, brute-force simulations are often impractical. However, rare event methods such as FFS can be used to calculate rate constants and transition pathways. Nucleation in “quasi-equilibrium” systems whose dynamics obeys detailed balance has been tackled using a range of different rare event simulation methods, including PPTIS [95], Metadynamics [96] committor analysis combined with two-dimensional Umbrella sampling scheme [97], Mean First Passage Time calculations [98, 99] and order-parameter-based Monte Carlo simulation [100]. We refer to these as “quasi-equilibrium” rather than “equilibrium” systems because for nucleation problems the initial state is always metastable with respect to the final state. For such systems, FFS provides a complementary approach to these methods. FFS also provides the potential for studying nucleation phenomena in out-of-equilibrium systems whose dynamics does not obey detailed balance (for example, with applied external shear, as discussed in the following section).

Over the last few years, FFS has been used to study crystal nucleation in both covalent [80, 81] and ionic quasi-equilibrium systems [82, 51, 83]. For these systems, it is possible to compare the results of FFS to those obtained by other rare event methods. In particular, in Ref.[82], we studied the nucleation of crystalline sodium chloride from the melt at moderate super-cooling. The nucleation rate was computed with a Bennett-Chandler procedure, in which Umbrella Sampling was used to compute the free energy barrier, followed by the firing of trajectories from the top of the barrier to obtain the

transmission coefficient. We also computed the nucleation rate using FFS. In both cases, the order parameter was taken to be the size of the largest solid cluster. Both methods yielded the same nucleation rate, to within the statistical error bars. In later work, [73], we used FFS to compute the nucleation free energy barrier for a two-dimensional Ising system, and compared the results to those of Umbrella Sampling. Both methods gave (to within error bars) the same free energy barrier height and shape. FFS has also recently been used to study vapour-crystal nucleation [84].

Recently, we combined FFS with Molecular Dynamics simulations to study bubble nucleation (cavitation) [53], and obtained results that differed from those of Umbrella Sampling. We computed the nucleation rate using FFS and also analysed the transition path ensemble. Our results showed that cavitation starts with compact bubbles rather than ramified ones, as has previously been suggested by Umbrella Sampling [101]. The FFS method does allow for the formation of ramified structures, but these pathways are kinetically unfavourable. Such kinetic effects cannot be observed in the Umbrella Sampling scheme [101], which assumes a local thermal equilibrium. Our FFS simulations also indicated a strong correlation between local temperature fluctuations in the liquid and subsequent bubble formation. Kinetic effects in homogeneous nucleation have also been investigated in the context of the crystallisation of a binary mixture of oppositely charged colloids, interacting via Yukawa potentials [51, 83]. Here, FFS simulations suggested that, for some thermodynamic conditions, the growth of the crystal does not follow the minimum free-energy path. This indicates a lack of ergodicity of the fluid on the timescale of crystal growth. This effect could not have been observed with a technique such as Umbrella Sampling, which assumes quasi-equilibrium transition paths.

These studies constitute a promising start for FFS in shedding new light on nucleation processes. For systems where the nucleation paths can be expected to be in local thermal equilibrium, FFS results agree well with those of other methods (such as Umbrella Sampling), whereas for several cases where kinetic effects are important, FFS has revealed unexpected behaviour. Plenty of scope remains for further investigation of such effects.

7.3. Nucleation in a sheared Ising model

Nucleation processes under shear are a class of nonequilibrium rare event problems to which FFS should be able to make a valuable contribution, since they are scientifically and technologically important and in many cases remain poorly understood. As a test case, we used FFS to study nucleation in a sheared two-dimensional Ising model [48, 49]. The lattice was sheared by periodically randomly selecting a row and shifting this and all higher rows to the right by one lattice site. This imposes a linear “velocity” profile. The spins were simulated with Metropolis Monte Carlo dynamics. We used DFFS, with the number of up spins as our order parameter, to compute the rate of nucleation of the stable state with a majority of up spins, starting from the “metastable” state with

a majority of down spins, in the presence of a weak external field favouring the up state [in the absence of shear, the free energy barrier to nucleation was $\approx 22k_B T$] [73]. This system is a poor model for most experimental systems, because particle transport is not modelled, and because the velocity profile is imposed rather than being determined by the system itself. Nevertheless, the computed nucleation rate showed interesting behaviour as a function of shear, with a peak at intermediate shear rate, as shown in Figure 14a.




Figure 14. (a): Nucleation rate I as a function of shear rate $\dot{\gamma}$ for the two dimensional sheared Ising model, computed using DFFS. The nucleation rate peaks at intermediate shear rate. For details of the parameters, please see Ref. [48]. (b): Scatter plot of the local density ρ_x (ρ_y) of up spins surrounding the largest cluster of up spins in the x (y) direction, versus the number of spins N_c in the largest cluster, for configurations in the transition state ensemble, at shear rate $\dot{\gamma} = 0.06$. 95% confidence intervals for the Pearson correlation coefficient r are also shown. The negative correlation observed for ρ_x , but not for ρ_y , demonstrates that coalescence along the x direction, driven by shear, plays a role in the nucleation mechanism. [Reproduced with permission from Ref [48]].

Understanding the physical mechanisms underlying the nonmonotonic trend in Figure 14a proved to be more difficult than calculating the rate itself. Simply comparing TSE configurations at different shear rates did not explain the effect of shear on the nucleation rate. We therefore resorted to devising modified shear algorithms, to test various hypotheses. For example, an algorithm with randomised shear direction (to eliminate shear-induced cluster breakup) removed the decrease in nucleation rate at high shear rate [48]. We also devised an unusual way of analysing the TSE, to test the hypothesis that shear-induced cluster coalescence was important in the enhancement of nucleation at low shear rates. We constructed an order parameter which we expected to be coupled to coalescence: the local density ρ_x of up spins close to the largest cluster in the x direction. Since the shear is applied in the x direction, we postulated that, if coalescence is important, then configurations with large values of ρ_x would have greater tendency to nucleate than those with small ρ_x . However, we did not expect such an effect for ρ_y . This could be tested by making scatter plots of ρ_x versus the largest cluster size N_c , for TSE configurations (which all have the same committor value $P_B = 0.5$). If coalescence is important, both N_c and ρ_x should contribute to the committor, so

that TSE configurations with large N_c will tend to have small ρ_x , and *vice versa*. We therefore expected negative correlation between ρ_x and N_c , but not between ρ_y and N_c , in the presence of shear only. This was indeed observed, as shown in Figure 14b, allowing us to conclude that shear-induced cluster coalescence is an important factor, at least in this model. The contrast between the behaviour of this model and Classical Nucleation Theory was discussed in a follow-up work [49], in which we investigated the effect of the external field strength.

The example discussed here is highly simplified in comparison to realistic sheared nucleation problems. However, even from this simple example, it is clear that nucleation problems under shear hold much potential, as well as presenting new challenges, for rare event simulation methods such as FFS.

8. Challenges and future directions

The development of rare event simulation methods in general, and FFS in particular, is far from complete, and many challenges remain. Some challenges, relating to computational efficiency, parameter optimisation and the choice of order parameter have already been addressed, resulting in the optimised versions of FFS [41, 42] discussed in this review. These improved algorithms should extend the method’s applicability to more computationally expensive and/or scientifically complex problems. It is likely that further fruits will be gained by combining advances made in different areas. For example, the Voronoi tessellation approach (Section 5.1) could be combined with committor-based order parameter optimisation schemes (Section 5.3). New FFS variants could also be developed by exploiting the analogy with polymer sampling discussed in Section 3.1. New developments will be driven by new users and new applications and to encourage this we believe that FFS should be implemented as soon as possible in widely used simulation packages.

One important issue is the possible presence of intermediate metastable states between A and B. These are a common feature of protein folding problems, and are likely to occur in many other contexts as well. With current FFS methods, partial paths which get stuck in these intermediate states will prove extremely expensive. This should prove a fruitful area for methodological development, especially if it is possible automatically to detect intermediate metastable states. Such methods could for example be used for “landscape exploration” in cases where the location of the B state is not known *a priori*. This would be useful for a variety of problems including the dynamics of glasses and protein conformational changes.

A second, related issue concerns systems with multiple alternative reaction channels. These are problematic for many rare event simulation methods. Dynamic sampling schemes like TPS and TIS, in which new transition paths are generated from old ones, have difficulty finding previously unexplored reaction channels (although replica exchange approaches can help [102, 24]). Since FFS is a static scheme in which new paths are generated from scratch, it should in principle be able to explore all

reaction channels. However, the choice of order parameter is likely to be crucial. For example, Sear [63] demonstrated that DFFS can fail to give the correct rate constant for a nucleation model in a two dimensional landscape where the dynamics evolves much more slowly in one dimension than in the other, if the order parameter involves only the fast coordinate. This is because the system does not fully explore the phase space along the slow coordinate and thus misses important transition paths. Juraszek *et al* [54] also reported that DFFS produced an incorrect rate constant for a protein folding problem where there were several possible folding pathways. This may have been due to undersampling of the ensemble of configurations at the first interface [54]; order parameter optimisation might be advantageous for such problems [55, 42], although it is not immediately clear whether one can optimise the order parameter to allow sampling of two reaction channels simultaneously. Further investigation of the performance of FFS-like methods for problems with multiple reaction channels is clearly needed, and progress in this direction has recently been initiated in the context of both TPS [25] and FFS [40].

Another important topic concerns how the method explores path space. In FFS, the segments of a transition path, once laid down, cannot be changed. If the nascent path turns out to be unfavourable, it will fail to reach the B state and the only option is to start again with a new path [for the BG or RB variants]. By contrast, in TPS or TIS, new paths are generated by shooting forwards and backwards in time, so that the initial segments of a path can relax in phase space. This offers likely advantages in complex free energy landscapes [24]. However, TPS and TIS are not suitable for nonequilibrium problems. It may be possible to devise FFS-like methods which sample complex landscapes more efficiently, by more intelligent sampling of the configurations at the interfaces, or by self-consistent determination of the order parameter [41].

Other directions of great interest are the idea of using time itself as an order parameter [5], and the potential for borrowing ideas from related methods in which trajectories are pruned or enriched continuously depending on the value of the Lyapunov exponent [103]. Finally, we note that the above discussion has assumed that the rate of transitions from A to B is constant in time (*i.e.* that our rare event is a Poisson process). This is not always the case [104, 105]. It remains to be seen whether FFS-like methods can be devised for non-Poissonian rare event problems.

Acknowledgments

RJA was funded by the Royal Society of Edinburgh and by the Royal Society. CV was funded by the EPSRC under grant EP/E030173 and by the European Union via a Marie Curie Individual fellowship. PRtW's contribution 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)". The authors acknowledge the contributions of Daan Frenkel and Patrick Warren to the development of the FFS method, and helpful discussions with Peter Bolhuis, Christoph

Dellago, Marco Morelli, Eduardo Sanz and Sorin Tănase-Nicola. We are grateful to Fernando Escobedo for providing us with a preprint of Ref. [7] and for the use of several figures from his work.

References

- [1] D. Frenkel and B. Smit. *Understanding Molecular Simulation. From Algorithms to Applications*. Academic Press, Boston, second edition, 2002.
- [2] P. G. Bolhuis, D. Chandler, C. Dellago, and P. L. Geissler. Transition path sampling: Throwing ropes over rough mountain passes, in the dark. *Annu. Rev. Phys. Chem.*, 53:291–318, 2002.
- [3] C. Dellago, P. G. Bolhuis, and P. L. Geissler. Transition path sampling. *Adv. Chem. Phys.*, 123:1–78, 2002.
- [4] C. Dellago and P. G. Bolhuis. Atomistic approaches in modern biology: from quantum chemistry to molecular simulations. pages 291–317, Berlin, 2007. Springer Verlag.
- [5] T. S. van Erp and P. G. Bolhuis. Elaborating Transition Interface Sampling methods. *J. Comp. Phys.*, 205:157–181, 2005.
- [6] C. Dellago and P. G. Bolhuis. Advanced computer simulation approaches for soft matter sciences III. pages 167–233. Springer, 2008.
- [7] F. A. Escobedo, E. E. Borrero, and J. C. Araque. Transition path sampling and forward flux sampling. Applications to biological systems. *J. Phys: Condens. Matter*, 21:333101, 2009.
- [8] H. Eyring. The activated complex in chemical reactions. *J. Chem. Phys.*, 3:107–115, 1935.
- [9] E. Wigner. The transition state method. *Trans. Faraday Soc.*, 34:29–40, 1938.
- [10] J. Horiuti. On the statistical mechanical treatment of the absolute rate of chemical reaction. *Bull. Chem. Soc. Japan*, 13:210–216, 1938.
- [11] C. H. Bennett. In R. Christofferson, editor, *Algorithms for Chemical Computations, ACS Symposium, Series No.46*, Washington, D.C., 1977. American Chemical Society.
- [12] D. Chandler. Statistical mechanics of isomerization dynamics in liquids and transition state approximation. *J. Chem. Phys.*, 68:2959–2970, 1978.
- [13] G. M. Torrie and J. P. Valleau. Monte-Carlo free-energy estimates using non-Boltzmann sampling - application to subcritical Lennard-Jones fluid. *Chem. Phys. Lett.*, 28:578, 1974.
- [14] J. S. van Duijneveld and D. Frenkel. Computer simulation study of free-energy barriers in crystal nucleation. *J. Chem. Phys.*, 96:4655–4668, 1992.
- [15] P. R. ten Wolde, M. J. Ruiz-Montero, and D. Frenkel. Numerical calculation of the rate of crystal nucleation in a Lennard-Jones system at moderate undercooling. *J. Chem. Phys.*, 104:9932–9947, 1996.
- [16] C. Dellago, P. G. Bolhuis, F. S. Csajka, and D. Chandler. Transition path sampling and the calculation of rate constants. *J. Chem. Phys.*, 108:1964–1977, 1998.
- [17] C. Dellago, P. G. Bolhuis, and D. Chandler. On the calculation of reaction rate constants in the transition path ensemble. *J. Chem. Phys.*, 110:6617–6625, 1999.
- [18] G. E. Crooks and D. Chandler. Efficient path sampling for nonequilibrium stochastic dynamics. *Phys. Rev. E*, 64:026109, 2001.
- [19] T. S. van Erp, D. Moroni, and P. G. Bolhuis. A novel path sampling method for the calculation of rate constants. *J. Chem. Phys.*, 118:7762, 2003.
- [20] J. B. Anderson. Predicting rare events in molecular dynamics. *Adv. Chem. Phys.*, 91:381–431, 1995.
- [21] D. Moroni, P. G. Bolhuis, and T. S. van Erp. Rate constants for diffusive processes by partial path sampling. *J. Chem. Phys.*, 120:4055–4065, 2004.
- [22] D. Moroni, T. S. van Erp, and P. G. Bolhuis. Simultaneous computation of free energies and kinetics of rare events. *Phys. Rev. E*, 71:056709, 2005.

- [23] T. S. van Erp. Reaction rate calculation by parallel path swapping. *Phys. Rev. Lett.*, 98:268301, 2007.
- [24] T. S. van Erp. Efficient path sampling on multiple reaction channels. *Comp. Phys. Commun.*, 179:34–40, 2008.
- [25] J. Rogal and P. G. Bolhuis. Multiple state transition path sampling. *J. Chem. Phys.*, 129:224107, 2008.
- [26] A. K. Faradjian and R. Elber. Computing time scales from reaction coordinates by milestoneing. *J. Chem. Phys.*, 120:10880–10889, 2004.
- [27] A. M. A. West, R. Elber, and D. Shalloway. Extending molecular dynamics timescales with milestoneing: Example of complex kinetics in a solvated peptide. *J. Chem. Phys.*, 126:145104, 2007.
- [28] E. Vanden-Eijnden, M. Venturoli, G. Ciccotti, and R. Elber. On the assumptions underlying milestoneing. *J. Chem. Phys.*, 129:174102, 2008.
- [29] E. Vanden-Eijnden and M. Venturoli. Markovian milestoneing with Voronoi tessellations. *J. Chem. Phys.*, 130:194101, 2009.
- [30] G. A. Huber and S. Kim. Weighted-ensemble Brownian dynamics simulations for protein association reactions. *Biophys. J.*, 70:97–110, 1996.
- [31] W. E, W. Ren, and E. Vanden-Eijnden. Finite temperature string method for the study of rare events. *J. Phys. Chem. B*, 109:6688, 2005.
- [32] E. Vanden-Eijnden and M. Venturoli. Revisiting the finite temperature string method for the calculation of reaction tubes and free energies. *J. Chem. Phys.*, 130:194103, 2009.
- [33] R. J. Allen, P. B. Warren, and P. R. ten Wolde. Sampling rare switching events in biochemical networks. *Phys. Rev. Lett.*, 94:018104, 2005.
- [34] R. J. Allen, D. Frenkel, and P. R. ten Wolde. Simulating rare events in equilibrium or nonequilibrium stochastic systems. *J. Chem. Phys.*, 124:024102, 2006.
- [35] H. Kahn and T. E. Harris. Estimation of particle transmission by random sampling. In *National Bureau of Standards Applied Mathematics Series, Vol. 12*, pages 27–30, 1951.
- [36] M. Villén-Altamirano and J. Villén-Altamirano. RESTART: A method for accelerating rare event simulations. In J. W. Cohen and C. D. Pack, editors, *Queuing, Performance and Control in ATM, Proceedings of the 13th International Telegraphic Congress*, 1991.
- [37] A. J. Bayes. Statistical techniques for simulation models. *Australian Computer Journal*, 2:180–184, 1970.
- [38] M. Villén-Altamirano, A. Martínez-Marrón, J. L. Gamo, and F. Fernández-Cuesta. Enhancement of the accelerated simulation method RESTART by considering multiple thresholds. In J. Lebetoulle and J. W. Roberts, editors, *The fundamental role of teletraffic in the evolution of telecommunications networks: Proceedings of the 14th International Telegraphic Congress*, pages 797–810, 1994.
- [39] M. Villén-Altamirano and J. Villén-Altamirano. Analysis of RESTART simulation: Theoretical basis and sensitivity study. *European Transactions on Telecommunications*, 13:373–385, 2002.
- [40] E. E. Borrero and F. A. Escobedo. Simulating the kinetics and thermodynamics of transitions via forward flux / umbrella sampling. *J. Phys. Chem. B*, 113:6434–6445, 2009.
- [41] E. E. Borrero and F. A. Escobedo. Reaction coordinates and transition pathways of rare events via forward flux sampling. *J. Chem. Phys.*, 127:164101, 2007.
- [42] E. E. Borrero and F. A. Escobedo. Optimizing the sampling and staging for simulations of rare events via forward flux sampling schemes. *J. Chem. Phys.*, 129:024115, 2008.
- [43] M. N. Rosenbluth and A. W. Rosenbluth. Monte Carlo calculation of the average extension of molecular chains. *J. Chem. Phys.*, 23:356, 1955.
- [44] D. Frenkel. Speed-up of Monte Carlo simulations by sampling of rejected states. *Proc. Natl. Acad. Sci. USA*, 101:17571, 2004.
- [45] M. J. Morelli, R. J. Allen, S. Tănase-Nicola, and P. R. ten Wolde. Eliminating fast reactions in stochastic simulations of biochemical networks: A bistable genetic switch. *J. Chem. Phys.*,

- 128:045105, 2008.
- [46] M. J. Morelli, S. Tănase-Nicola, R. J. Allen, and P. R. ten Wolde. Reaction coordinates for the flipping of genetic switches. *Biophys. J.*, 94:3413–3423, 2008.
 - [47] M. J. Morelli, P. R. ten Wolde, and R. J. Allen. DNA looping provides stability and robustness to the bacteriophage λ switch. *Proc. Natl. Acad. Sci. USA*, 106:8101–8106, 2009.
 - [48] R. J. Allen, C. Valeriani, S. Tănase-Nicola, P. R. ten Wolde, and D. Frenkel. Homogeneous nucleation under shear in a two-dimensional Ising model: Cluster growth, coalescence, and breakup. *J. Chem. Phys.*, 129:134704, 2008.
 - [49] R. J. Allen, C. Valeriani, and S. Tănase-Nicola. Nucleation in a sheared Ising model: effects of external field. *Prog. Theor. Phys. Supp.*, 175:144–153, 2008.
 - [50] A. J. Page and R. P. Sear. Heterogeneous nucleation in and out of pores. *Phys. Rev. Lett.*, 97:065701, 2006.
 - [51] E. Sanz, C. Valeriani, D. Frenkel, and M. Dijkstra. Evidence for out-of-equilibrium crystal nucleation in suspensions of oppositely charged colloids. *Phys. Rev. Lett.*, 99:055501, 2007.
 - [52] J. S. van Zon and P. R. ten Wolde. Green’s-function reaction dynamics: a particle-based approach for simulating biochemical networks in time and space. *J. Chem. Phys.*, 123:234910, 2005.
 - [53] Z.-J. Wang, C. Valeriani, and D. Frenkel. Homogeneous bubble nucleation driven by local hot spots: A molecular dynamics study. *J. Phys. Chem. B*, 113:3776–3784, 2009.
 - [54] J. Juraszek. *Proteins in Action; Simulations of Conformational changes in small proteins*. PhD thesis, Universiteit van Amsterdam, 2008.
 - [55] C. Velez-Vega, E. E. Borrero, and F. A. Escobedo. Kinetics and reaction coordinate for the isomerization of alanine dipeptide by a forward flux sampling protocol. *J. Chem. Phys.*, 130:225101, 2009.
 - [56] G. C. A. M. Mooij and D. Frenkel. A systematic optimization scheme for configurational bias Monte Carlo. *Mol. Simul.*, 17:41–55, 1996.
 - [57] T. S. van Erp. Efficiency analysis of reaction rate calculation methods using analytical models I: The two-dimensional sharp barrier. *J. Chem. Phys.*, 125:174106, 2006.
 - [58] R. J. Allen, D. Frenkel, and P. R. ten Wolde. Forward flux sampling-type schemes for simulating rare events: Efficiency analysis. *J. Chem. Phys.*, 124:194111, 2006.
 - [59] R. M. S. Anderson, C. Valeriani, and R. J. Allen. (In preparation).
 - [60] R. S. Maier and D. L. Stein. Escape problem for irreversible systems. *Phys. Rev. E*, 48:931–938, 1993.
 - [61] R. S. Maier and D. L. Stein. A scaling theory of bifurcations in the symmetric weak-noise escape problem. *J. Stat. Phys.*, 83:291–357, 1996.
 - [62] D. G. Luchinsky, R. S. Maier, R. Mannella, P. V. E. McClintock, and D. L. Stein. Observation of saddle-point avoidance in noise-induced escape. *Phys. Rev. Lett.*, 82:1806–1809, 1999.
 - [63] R. P. Sear. Nucleation in the presence of slow microscopic dynamics. *J. Chem. Phys.*, 128:214513, 2008.
 - [64] P. R. ten Wolde, M. J. Ruiz-Montero, and D. Frenkel. Numerical calculation of the rate of crystal nucleation in a Lennard-Jones system at moderate undercooling. *J. Chem. Phys.*, 104:9932, 1996.
 - [65] P. R. ten Wolde and D. Chandler. Drying-induced hydrophobic polymer collapse. *Proc. Natl. Acad. Sci. USA*, 99:6539–6543, 2002.
 - [66] P. G. Bolhuis, C. Dellago, and D. Chandler. Reaction coordinates of biomolecular isomerization. *Proc. Natl. Acad. Sci. USA*, 97:5877–5882, 2000.
 - [67] A. Dickson, A. Warmflash, and A. R. Dinner. Nonequilibrium umbrella sampling in spaces of many order parameters. *J. Chem. Phys.*, 130:074104, 2009.
 - [68] G. Hummer. From transition paths to transition states and rate coefficients. *J. Chem. Phys.*, 120:516–523, 2004.
 - [69] R. B. Best and G. Hummer. Reaction coordinates and rates from transition paths. *Proc. Natl. Acad. Sci. USA*, 102:6732–6737, 2005.

- [70] A. Ma and A. R. Dinner. Automatic method for identifying reaction coordinates in complex systems. *J. Phys. Chem. B*, 109:6769, 2005.
- [71] B. Peters and B. L. Trout. Obtaining reaction coordinates by likelihood maximization. *J. Chem. Phys.*, 125:054108, 2006.
- [72] B. Peters, G. T. Beckham, and B. L. Trout. Extensions to the likelihood maximization approach for finding reaction coordinates. *J. Chem. Phys.*, 127:034109, 2007.
- [73] C. Valeriani, R. J. Allen, M. J. Morelli, D. Frenkel, and P. R. ten Wolde. Computing stationary distributions in equilibrium and nonequilibrium systems with forward flux sampling. *J. Chem. Phys.*, 127:114109, 2007.
- [74] A. Warmflash, P. Bhimalapuram, and A. R. Dinner. Umbrella sampling for nonequilibrium processes. *J. Chem. Phys.*, 127:154112, 2007.
- [75] W. E, W. Ren, and E. Vanden-Eijnden. String method for the study of rare events. *Phys. Rev. B*, 66:052301, 2002.
- [76] W. Ren, E. Vanden-Eijnden, P. Maragakis, and W. E. Transition pathways in complex systems: Application of the finite temperature string method to the alanine dipeptide. *J. Chem. Phys.*, 123:134109, 2005.
- [77] R. P. Sear. Heterogeneous and homogeneous nucleation compared: Rapid nucleation on microscopic impurities. *J. Phys. Chem. B*, 110:4985, 2006.
- [78] R. P. Sear. Nucleation at contact lines where fluid-fluid interfaces meet solid surfaces. *J. Phys.: Condens. Matter*, 19:466106, 2007.
- [79] D. P. Sanders, H. Larralde, and F. Leyvraz. Competitive nucleation and the Ostwald rule in a generalized Potts model with multiple metastable phases. *Phys. Rev. B*, 75:132101, 2007.
- [80] L. M. Ghiringhelli, C. Valeriani, E. J. Meijer, and D. Frenkel. Local structure of liquid carbon controls diamond nucleation. *Phys. Rev. Lett.*, 99:055702–4, 2007.
- [81] L. M. Ghiringhelli, C. Valeriani, J. H. Los, E. J. Meijer, A. Fasolino, and D. Frenkel. State-of-the-art models for the phase diagram of carbon and diamond nucleation. *Molecular Physics*, 106:2011–2038, 2008.
- [82] C. Valeriani, E. Sanz, and D. Frenkel. Rate of homogeneous crystal nucleation in molten NaCl. *J. Chem. Phys.*, 122:194501–194506, 2005.
- [83] E. Sanz, C. Valeriani, T. Vissers, A. Fortini, M.E. Leunissen, A. van Blaaderen, D. Frenkel, and M. Dijkstra. Out-of-equilibrium processes in suspensions of oppositely charged colloids: liquid-to-crystal nucleation and gel formation. *J. Phys. Cond. Matt.*, 20:494247, 2008.
- [84] J.A. van Meel, A.J. Page, R.P. Sear, and D. Frenkel. Two-step vapor-crystal nucleation close below triple point. *J. Chem. Phys.*, 129:204505, 2008.
- [85] W. Möbius, R. A. Neher, and U. Gerland. Kinetic accessibility of buried DNA sites in nucleosomes. *Phys. Rev. Lett.*, 97:208102, 2006.
- [86] L. Rekvig and D. Frenkel. Molecular simulations of droplet coalescence in oil/water/surfactant systems. *J. Chem. Phys.*, 127:134701, 2007.
- [87] L. Huang and D. E. Makarov. The rate constant of polymer reversal inside a pore. *J. Chem. Phys.*, 128:114903, 2008.
- [88] J. P. Hernandez-Ortiz and J. J. de Pablo. Hydrodynamic effects on the translocation rate of a polymer through a narrow pore. (Submitted).
- [89] E. E. Borrero and F. A. Escobedo. Folding kinetics of a lattice protein via a forward flux sampling approach. *J. Chem. Phys.*, 125:164904, 2006.
- [90] M. Ptashne. *A Genetic Switch: Gene Control and λ phage*. Cell Press and Blackwell Scientific Publications, Cambridge, MA, 1986.
- [91] A. Raj and A. van Oudenaarden. Nature, nurture or chance: Stochastic gene expression and its consequences. *Cell*, 135:216–226, 2008.
- [92] J. W. Little, D. P. Shepley, and D. W. Wert. Robustness of a gene regulatory circuit. *EMBO J.*, 18:4299–4307, 1999.
- [93] P. B. Warren and P. R. ten Wolde. Enhancement of the stability of genetic switches by overlapping

- upstream regulatory domains. *Phys. Rev. Lett.*, 92:128101, 2004.
- [94] P. B. Warren and P. R. ten Wolde. Chemical models of genetic toggle switches. *J. Phys. Chem. B*, 109:6812, 2005.
- [95] D. Moroni, P. R. ten Wolde, and P. G. Bolhuis. Interplay between structure and size in a critical crystal nucleus. *Phys. Rev. Lett.*, 94:235703, 2005.
- [96] F. Trudu, D. Donadio, and M. Parrinello. Freezing of a Lennard-Jones fluid: from nucleation to spinodal regime. *Phys. Rev. Lett.*, 97:105701, 2006.
- [97] R. Radhakrishnan and B. L. Trout. Nucleation of hexagonal ice (Ih) in liquid water. *J. Am. Chem. Soc.*, 125:7743, 2003.
- [98] J. Wedekind, R. Strey, and D. Reguera. New method to analyze simulations of activated processes. *J. Chem. Phys.*, 126:134103, 2007.
- [99] G. Chkonia, J. Wölk, R. Strey, J. Wedekind, and D. Reguera. Evaluating nucleation rates in direct simulations. *J. Chem. Phys.*, 130:064505, 2009.
- [100] M. Chopra, M. Muller, and J. de Pablo. Order-parameter-based Monte Carlo simulation of crystallization. *J. Chem. Phys.*, 124:134102, 2006.
- [101] V. K. Shen and P. G. Debenedetti. A computational study of homogeneous liquid-vapor nucleation in the lennard-jones fluid. *J. Chem. Phys.*, 111:3581–3589, 1999.
- [102] P. G. Bolhuis. Rare events via multiple channels sampled by path replica exchange. *J. Chem. Phys.*, 129:114108, 2008.
- [103] J. Tailleur and J. Kurchan. Probing rare physical trajectories with Lyapunov weighted dynamics. *Nature Physics*, 3:203, 2007.
- [104] P. Visco, R. J. Allen, and M. R. Evans. Exact solution of a model DNA-inversion genetic switch with orientational control. *Phys. Rev. Lett.*, 101:118104, 2008.
- [105] P. Visco, R. J. Allen, and M. R. Evans. Statistical physics of a model binary genetic switch with linear feedback. *Phys. Rev. E*, 79:031923, 2009.