UvA-DARE (Digital Academic Repository)

Foundations and latest advances in replica exchange transition interface sampling

Cabriolu, R.; Refsnes, K.M.S; Bolhuis, P.G.; van Erp, T.S.

[Link to publication](#)

Citation for published version (APA):
Cabriolu, R., Refsnes, K. M. S., Bolhuis, P. G., & van Erp, T. S. (2017). Foundations and latest advances in replica exchange transition interface sampling. *Journal of Chemical Physics*, *147*(15), [152722]. https://doi.org/10.1063/1.4989844

Foundations and latest advances in replica exchange transition interface sampling

Raffaela Cabriolu,[1,a)] Kristin M. Skjelbred Refsnes,[1,b)] Peter G. Bolhuis,[2,c)] and Titus S. van Erp[1,d)]

[1]*Department of Chemistry, Norwegian University of Science and Technology (NTNU), Høgskoleringen 5, 7491 Trondheim, Norway*

[2]*Van 't Hoff Institute for Molecular Sciences, University of Amsterdam, P.O. Box 94157, 1090 GD Amsterdam, The Netherlands*

Nearly 20 years ago, transition path sampling (TPS) emerged as an alternative method to free energy based approaches for the study of rare events such as nucleation, protein folding, chemical reactions, and phase transitions. TPS effectively performs Monte Carlo simulations with relatively short molecular dynamics trajectories, with the advantage of not having to alter the actual potential energy surface nor the underlying physical dynamics. Although the TPS approach also introduced a methodology to compute reaction rates, this approach was for a long time considered theoretically attractive, providing the exact same results as extensively long molecular dynamics simulations, but still expensive for most relevant applications. With the increase of computer power and improvements in the algorithmic methodology, quantitative path sampling is finding applications in more and more areas of research. In particular, the transition interface sampling (TIS) and the replica exchange TIS (RETIS) algorithms have, in turn, improved the efficiency of quantitative path sampling significantly, while maintaining the exact nature of the approach. Also, open-source software packages are making these methods, for which implementation is not straightforward, now available for a wider group of users. In addition, a blooming development takes place regarding both applications and algorithmic refinements. Therefore, it is timely to explore the wide panorama of the new developments in this field. This is the aim of this article, which focuses on the most efficient exact path sampling approach, RETIS, as well as its recent applications, extensions, and variations. *Published by AIP Publishing.*
https://doi.org/10.1063/1.4989844

I. INTRODUCTION

Rare events encompass many fields of great interest. For example, many phenomena are dominated by rare events in geological processes,[1] in spreading of diseases,[2] in physical-chemical processes,[3] in climate changes,[4] in stock market fluctuations,[5] and in human conflicts.[6] In physics and chemistry, in particular, rare events are widespread. We refer to an event as rare when occurring very infrequently compared to other relaxation processes involved in the phenomenon. In several reactions, the product forms on a long time scale compared to the molecular time scale, i.e., molecular vibrations. Indeed, in such activated events, the *crossing* of the free energy barrier dividing the reactant from the product state happens extremely infrequently. However, once this rare event occurs, the reaction will proceed very rapidly to the product state. In nucleation phenomena, for example, nuclei of the new phase would form and dissolve many times before assuming, due to thermal fluctuations, a large enough size that will deterministically grow up to the product state.[3,7]

Unfortunately, most experiments cannot directly follow the mechanistic of the rare-event processes at the molecular level but, on the other hand, are able to indirectly collect macroscopical evidence of the phenomena to a certain accurate extent. In principle, equilibrium properties of rare events could be successfully investigated using Monte Carlo (MC) simulations that, through *nonphysical moves*, respecting equilibrium statistical mechanics principles take the system from the reactant state to local minima of the free energy landscape[8] and eventually to the products. However, in order to have a dynamical description of the physical mechanism and kinetic details of the transition processes, molecular dynamics (MD) based methods are necessary.

Unfortunately, the disparate length and time scales involved in those processes make the standard MD approach ineffective. The time scale of atomistic MD simulations for realistic dimension ($>10^5$ particles) is typically limited to 10^3 ns and is therefore far below the relevant time range of many processes dominated by rare events. Furthermore, often, the lack of *a priori* knowledge of the physical mechanism makes the results from MD simulations, carried out using coarse-grained models, unreliable, and, hence, a detailed description

---

a)Electronic mail: raffaela.cabriolu@ntnu.no
b)Electronic mail: kristin.m.skjelbred@ntnu.no
c)Electronic mail: p.g.bolhuis@uva.nl
d)Author to whom correspondence should be addressed: titus.van.erp@ntnu.no

at the atomistic scale is the only option. For those reasons, in the last few decades, tremendous efforts have been dedicated to the development of new computational statistical mechanical methods that address the challenges inherent to the rare events. The exploration of new computational methods, which is still progressing rapidly, has provided a wealth of information on rare-event processes but also highlighted the need for more efficient and versatile approaches.

In this article, we briefly discuss the rare-event computational techniques in general and then focus on the exact path sampling approaches such as transition interface sampling (TIS) and replica exchange TIS (RETIS). In particular, since the latter method is the most efficient while still being exact, this approach has been selected as the central theme of this paper. Our intention is to give an explanation of the new developments on path sampling techniques without going through the historical developments nor by going into too much mathematical detail. For more complete and broader overviews, we refer to other books and review articles such as Refs. 3 and 7–11. In particular, the recent book by Peters[3] gives a good introduction to the rare event methodology.

This paper is organized as follows. In Sec. II, we give a brief introduction on the methods to study rare events. In particular, we distinguish between free energy based methods, approximate path sampling methods, and exact path sampling methods. Section III explains the algorithm and theory behind the RETIS method, which is the starting point of this article. In Sec. IV, we give an overview of the computational studies that have applied RETIS. We have dedicated Sec. V to practical and technical features of the path sampling techniques, such as the efficiency of the algorithms related to the placement of the interfaces and to the reaction coordinate (RC). In Sec. VI, we outline the recent developments of the path sampling techniques describing the new shooting moves, analysis of paths, QuanTIS, multiple state TIS, and single replica TIS algorithms. Finally, we end with some concluding remarks in Sec. VII.

II. METHODS FOR RARE EVENTS

Among the different simulation approaches to study rare events, a distinction between *free energy* and *path sampling* methods can be made. The first class of methods provides information on the reaction mechanism through sampling the configuration space. Path sampling methods, on the other hand, are directly focused on the dynamics of the process.

A. Free energy based methods

In free energy based rare event methods, the projection of the Helmholtz free energy barrier onto certain collective variables $q_1, q_2, \ldots$, or functions of phase space coordinates, needs to be computed. This is also called the Landau free energy.[3] In case the barriers are sufficiently low, the sampling of configuration space can be done by either MC or MD. MC naturally samples the canonical distribution at a constant temperature, while MD needs to be combined with a proper thermostat as it would otherwise sample the micro-canonical

distribution at a constant energy. The Landau free energy then follows from $F(q_1, q_2, \ldots) = -k_B T \ln \rho(q_1, q_2, \ldots)/\rho_o$, where T is the temperature, k_b is the Boltzmann constant, and ρ_o is a constant to make the argument of the logarithmic function dimensionless (it can be chosen arbitrarily if only free energy differences matter). The distribution $\rho(q_1, q_2, \ldots)$ follows from the sampling data by simply binning the relevant order parameter space and keeping track of how many times each bin is visited during the simulation run.

However, as in most cases the sampling tends to get trapped in metastable states, configuration based importance sampling is needed. While a wide variety of importance sampling methods have been developed for constructing free energy profiles, almost all approaches can be viewed, in some respect, as a variation of two well-established methods, which are thermodynamic integration (TI)[8] and umbrella sampling (US).[12] Adaptive biasing force methods[13–15] are evolutions of the TI technique, while metadynamics can be viewed as an adaptive US approach.

It should be noted, however, that a free energy surface provides insight that is not necessarily dynamically relevant. For instance, the heights of free barriers between metastable states depend on which order parameters have been chosen. Dynamical information can be obtained if it is combined with some sort of dynamical approach, which is discussed in Sec. II B.

B. Approximate and exact dynamical path sampling methods

There are many different approaches to study the dynamics of rare events. Of course, the most straightforward approach would be MD. Due to the recent development of special purpose computers, such as the Anton machines, the millisecond time scale has come into reach[16] using brute force MD. However, many processes take place at a time scale larger than this and many reactive trajectories are usually necessary to perform decent statistics. In addition, special purpose machine simulations are restricted to a relatively narrow class of force fields since they must be *programmed* into the hardware of the computer chips. *Ab initio* MD simulations, for instance, cannot be treated. Instead of increasing the speed of dynamical exploration via hardware development, one can rely on clever algorithms. These can rely on approximations, but orders of magnitude increase is still possible while getting theoretically the exact same results one would obtain with endlessly long MD runs.

The most common approach is to use the free energy profile as a starting point and to invoke the transition state theory (TST) approximation. In this case, no additional simulation is needed besides the free energy calculation. TST is very successful in low-dimensional systems, and, for deterministic dynamics, it is assumed to be correct if the ideal reaction coordinate is used. However, finding this ideal reaction coordinate for which TST is exact is very difficult and even its existence is questionable.[17] Moreover, for non-deterministic stochastic dynamics, the TST approximation is wrong regardless the chosen reaction coordinate (RC).[3,8]

Other approaches to approximate the dynamics of the rare event are Partial Path TIS (PPTIS)[18] and milestoning.[19] The PPTIS method is a variation of the exact TIS algorithm and has been developed to study efficiently diffusive processes that are characterized by relatively flat, rough, and wide free energy barriers. The trajectories sampled in PPTIS are considerably shorter than those being generated in the TIS or RETIS algorithm. Based on the Markovian assumption that a trajectory loses its memory over the distance between two interfaces, PPTIS is able to compute the same dynamical information as TIS or RETIS by generating considerably shorter paths than those being generated in these exact algorithms. Milestoning,[19] which was developed at the same time independently by Faradjian and Elber, is, except for some algorithmic differences, similar to PPTIS. Still, milestoning's Markovian approximation is stronger than the one in PPTIS as it assumes full memory loss at each interface. On the other hand, by computing time-dependent crossing probability densities, instead of just crossing probabilities, milestoning does not have to rely on the separation of time scales. This can be an advantage when computing other dynamical properties that do not necessarily rely on this separation, such as diffusion constants. As was suggested in Ref. 20 and realized in Ref. 21, the inclusion of spatial memory as in PPTIS and time-dependence as in milestoning can be combined in a single method. It is interesting to note that milestoning and PPTIS can become exact for the case of the ideal reaction coordinate.[22] However, if the reaction coordinate is not ideal and not enough loss of memory is allowed by sufficient separation between the interfaces, results might be quantitatively inaccurate and even qualitatively misleading.[11]

Exact path sampling methods, such as transition path sampling (TPS),[23–25] do not have this issue since they provide results that are correct and independent of the RC that has been chosen. Starting from an initial reactive trajectory of a certain length, the *shooting algorithm* generates new trajectories with the same path length. The new trajectories are accepted according to the Metropolis-Hastings algorithm.[8] Further, the rate constant can be determined by computing a population correlation function via an US approach in which the umbrella potential is applied to the end point of the potential. Significant improvement to the rate evaluation approach was made by the introduction of the TIS algorithm.[26] TIS allows the path length to be flexible, minimizing the number of MD steps required, and shows faster convergence relative to the number of paths generated. Moreover, the US approach is replaced in TIS by path ensembles based on the interface crossing condition. The approach was further improved via the RETIS method[27,28] which is the main focus of this article. Section III is devoted to explain this method in more detail.

Further, we can mention forward flux sampling (FFS) as a member of the exact path sampling approaches. FFS is based on the TIS theoretical framework but uses splitting[29,30] instead of the MC shooting approach. The method emerging from this very much resembles the RESTART algorithm.[31] FFS has the advantage of being able to simulate non-equilibrium processes, but, on the other hand, it cannot be used in combination with a dynamics that is largely deterministic in nature. In addition,

the use of the FFS algorithm comes with a relatively high risk that it provides reactive trajectories that cross the barrier via the wrong mechanism.[11]

Finally, dynamical calculations can also be used in combination with a free energy calculation in order to correct the TST approximation. This is the so-called reactive flux (RF) method that in the literature is also referred to as the Wigner-Keck-Eyring (WKE) method[32] or as the Bennett-Chandler (BC)[8] method, giving credit to different generations of scientists who have contributed to the approach.[33–37] The approach corrects the TST expression for fast re-crossings by releasing many trajectories from the top of the free energy barrier and computing a flux weighted average. It is important to note that although the exact dynamical approaches provide results that do not depend on the choice of the reaction coordinate, the efficiency generally will. RF and in particular FFS are highly sensitive to the choice of the RC, which implies that the approaches become inefficient if the RC is not well chosen. This efficiency issue will further be discussed in Sec. V B.

III. REPLICA EXCHANGE TRANSITION INTERFACE SAMPLING

RETIS, or replica exchange TIS, is an approach that combines the standard TIS algorithm with a swapping algorithm. While replica exchange methods improve the sampling efficiency simulating replicas at different temperatures, RETIS does not require additional simulations, but it just borrows the idea of *exchanging* trajectories between different *path ensembles* instead of *swapping* configurations obtained by simulations at different *temperatures*.

According to non-equilibrium statistical mechanics, to study the rate of a reaction between two well-defined stable states A and B, we could resort to the derivative of a time correlation function $C(t)$ between the reactant and product populations.[3] In molecular simulation, the correlation function $C(t)$ can be expressed as ensemble averages of indicator functions h_A and h_B, and it measures the probability to find the system in the state B at time t provided that it was in A at the initial time 0,

$$C(t) = \frac{\langle h_A(0) h_B(t) \rangle}{\langle h_A \rangle}. \qquad (1)$$

A reaction coordinate (RC) $\lambda(x)$, that is a function of the phase space x, is chosen to well distinguish two limited regions around stable state A and B. If the system is in A or B, $\lambda(x)$ is smaller than λ_A or larger than λ_B, respectively. According to TPS, h_A is 1 if the system is found in A or zero otherwise. In the same way, h_B is 1 if the system is in B or 0 in all the other phase points.

Importantly, correlation functions bridge the macroscopic definition of the rate laws with the microscopic dynamic information extracted by simulations. In general, according to kinetic theory, if there is a clear separation of time scales in the process, there will be a regime where the correlation function grows linearly with the time, or equivalently, its time derivative will show a horizontal plateau that equals the forward reaction rate constant k_{AB},

$$k(t) = \frac{dC(t)}{dt}, \tag{2}$$

$$k_{AB} = k(t') \text{ for } t_{\text{mol}} < t' \ll t_{\text{rxn}}, \tag{3}$$

where t_{mol} is the molecular time scale related to temporary molecular fluctuations and t_{rxn} is the exponential relaxation time within which the reaction occurs (see Fig. 1).

In Eq. (1), the ensemble averages $\langle \ldots \rangle$ should be taken over the initial conditions. The time point $t = 0$ has no absolute meaning since for any generic time correlation function connecting function values of g and h at different times, the following holds: $\langle g(0)h(t) \rangle = \langle g(t')h(t+t') \rangle$ for any t'. Moreover, in the limit $t \to \infty : \langle g(t')h(t+t') \rangle = \langle g \rangle \langle h \rangle$, which is the reason why $C(t)$ converges to $\langle h_B \rangle$ at $t \gg t_{\text{rxn}}$.

If the transition is not a rare event, the most straightforward way to compute $C(t)$ is by averaging over different time slices in an MD simulation, basically shifting the initial time $t = 0$,

$$C(M\Delta t) = \frac{\frac{1}{N-M+1}\sum_{i=0}^{N-M} h_A(i\Delta t)h_B((i+M)\Delta t)}{\frac{1}{N+1}\sum_{i=0}^{N} h_A(i\Delta t)}, \tag{4}$$

where Δt is the MD time step, N is the total number of MD steps performed in the simulation, and $M\Delta t = t$ is the time for which the correlation function is considered.

TPS and TIS/RETIS differ with respect to the characteristic functions used to calculate $C(t)$ from Eq. (1). The characteristic functions in TPS can be both zero at a given time (see Fig. 2). In fact, the nominator in Eq. (4), involving the sum over $h_A(i\Delta t)h_B((i+M)\Delta t)$, is mostly zero. For instance, for $M = 25$, the only terms in Fig. 2 that give 1 are $i = 0$ and $i = 11$.
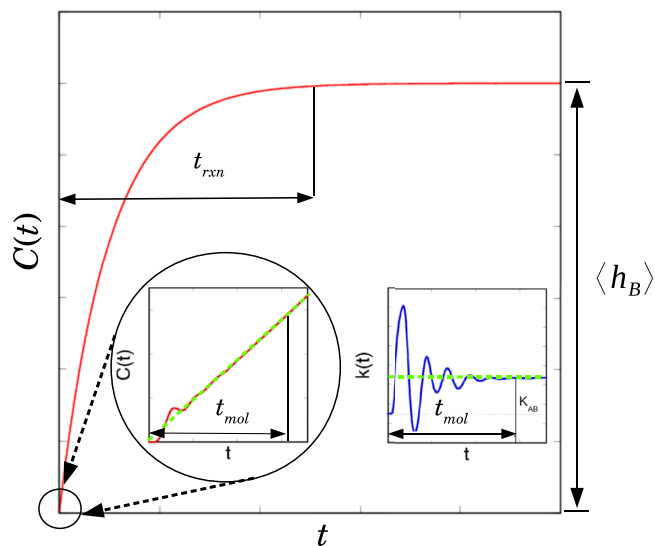
The numerical derivative of $C(t)$ gives

FIG. 1. Correlation function for determining the rate constant [Eq. (2)]. Outer panel shows the long time scale behavior in which the correlation function converges to a horizontal plateau. Insets show $C(t)$ and its time derivative $k(t)$ at the shorter time scale. The rate constant is obtained from the slope of $C(t)$ at a time t' in the region $t_{\text{mol}} < t' \ll t_{\text{rxn}}$. Dashed green lines in the insets show the correlation function and its derivative based on *overall states*.


FIG. 2. Explanation of the state definitions and the correlation function C(t) [Eqs. (1)–(4)] based on a hypothetical MD run that visits both state A (green square) and state B (yellow square). Symbols in the trajectories are configurations at each MD step. The full green and full yellow circles are the trajectory points included in A ($h_A = 1$) and B ($h_B = 1$). At the cross symbols, both the TPS characteristic functions are null, while in TIS, there is a switching between overall state $\mathcal{A}$ (green line) and overall state $\mathcal{B}$ (light brown line). The dashed horizontal line represents the transition state dividing surface along y.

$$k(t) = \frac{dC}{dt} = \frac{\langle h_A(0)\frac{d}{dt}(h_B(t))\rangle}{\langle h_A \rangle} \Rightarrow$$

$$k(M\Delta t) = \frac{1}{\Delta t \langle h_A \rangle}\frac{1}{N-M}\sum_{i=0}^{N-M-1} [h_A(i\Delta t)$$

$$\times \{h_B((i+M+1)\Delta t) - h_B((i+M)\Delta t)\}]. \tag{5}$$

Hence, for $M = 25$, the only non-zero contributions are for $i = 0$ and $i = 10$ where the part within the curly brackets produces -1 and $+1$, respectively.

The positive and negative contributions of the characteristic functions to the average in Eq. (5) give rise to the fluctuating behavior for $t < t_{\text{mol}}$ in $k(t)$. Moreover, k_{AB} requires sufficient data to ensure convergence, and preliminary results might even be negative. Naturally, TPS does not rely on a plain MD simulation to compute $C(t)$ but uses an important sampling technique. Still, the slow convergence due to negative terms is also present in these importance sampling approaches. This effect, in addition to the fixed path length, makes TPS less efficient than TIS and RETIS.

In RETIS, as in TIS, the introduction of the *overall states* $\mathcal{A}$ and $\mathcal{B}$ eliminates the fluctuations. Instead, $C(t)$ based on *overall states* is linear from the start (see Fig. 1). By definition, the overall state $\mathcal{A}$ includes all phase space points lying inside the stable state A and all the phase space points that, based on their history, were more recently in stable state A than in stable state B. In the same way, the overall state $\mathcal{B}$ includes all phase space points lying inside the stable state B and all phase points that were more recently part of stable state B than A. The corresponding characteristic functions $h_{\mathcal{A}}$ and $h_{\mathcal{B}}$ are not very sensitive to the stable state definitions. This is clearly illustrated by the fact that if we would shrink the stable state boundary of A in Fig. 2 such that only MD point 0 is inside stable state A, none of the time slices shown in the figure, that are presently part of $\mathcal{A}$, would actually change to $\mathcal{B}$.

In RETIS, as in the TIS algorithm, the phase space is divided by a set of n interfaces $\lambda_i \in \{\lambda_0, \ldots, \lambda_n\}$, with $\lambda_i < \lambda_{i-1}$. Each interface λ_i is then defined as the multidimensional surface at which the RC assumes exactly the value λ_i. In particular, the configuration with a RC value less than $\lambda_0 = \lambda_A$ belongs to the reactant state A, while it belongs to the product state B if its RC has a value higher than $\lambda_n = \lambda_B$.

Using statistical mechanical arguments, the TIS reaction rate can be written as[11]

$$k_{AB} = f_A P(\lambda_B|\lambda_A) = f_A \prod_{i=0}^{n-1} P_A(\lambda_{i+1}|\lambda_i), \quad (6)$$

where f_A is the flux of trajectories through the initial interface λ_A per unit time and $P(\lambda_B|\lambda_A)$ is a conditional probability. Namely, $P(\lambda_B|\lambda_A) = P_A(\lambda_n|\lambda_0)$ is the probability that a path starting from A, after having crossed λ_A, will cross the interface λ_B before returning to A. This *overall* conditional probability can be conveniently factorized into the probabilities $P_A(\lambda_{i+1}|\lambda_i)$ that have much higher values than the overall probability, reducing the computational cost. In particular, $P_A(\lambda_{i+1}|\lambda_i)$ is the probability that a trajectory starting from A crosses the interface λ_{i+1} after having crossed the interface λ_i without returning to A first. This is a particular case of the generic history dependent conditional crossing probabilities that are being used in all TIS variations (see Fig. 3).

The crossing probability is central to all TIS variations like RETIS. Using a finer grid of sub-interfaces in the analysis, the crossing probability can be depicted as a continuous function $P_A(\lambda|\lambda_0)$. This function is basically the dynamical analog of the free energy profile, but the shape of this curve depends explicitly on the dynamics of the system. For instance, if the equations of motion are governed by the Langevin dynamics, the curve will, unlike the free energy profile, depend on the friction coefficient. $P_A(\lambda|\lambda_0)$ is a strictly decreasing function starting from 1 at λ_A and ending with a horizontal plateau after crossing the barrier (see Fig. 7 in Sec. IV).

Using the Monte Carlo (MC) importance sampling technique, TIS/RETIS generates the ensembles of paths

$\{[0^+], [1^+], \ldots, [(n-1)^+]\}$ that obey the *crossing* condition of the interfaces $\{\lambda_0, \lambda_1, \ldots, \lambda_{n-1}\}$, respectively. The main differences between TIS and RETIS are the swapping move and the $[0^-]$ ensemble that we will discuss now.

In TIS, f_A is determined by a plain MD simulation,

$$f_A = \frac{N_c^+}{T_{\in\mathcal{A}}}, \quad (7)$$

where N_c^+ is the number of positive crossings with interface $\lambda_A = \lambda_0$ and $T_{\in\mathcal{A}}$ is the time spent in $\mathcal{A}$. Since the transition to $\mathcal{B}$ will be a rare event, $T_{\in\mathcal{A}}$ will be in most cases equal to the total simulation time. However, if a spontaneous crossing happens, it is best not to wait until the system returns to state A, but to run another MD simulation initialized from A and take the average in the end.

RETIS, on the other hand, is only based on path sampling simulations. For this purpose, it introduces the path ensemble $[0^-]$ which comprises all paths that start at λ_A, proceed in the opposite direction of the reaction progress, and end at λ_A again. The introduction of the $[0^-]$ ensemble in RETIS allows the calculation of the flux through the average path lengths, $\langle t^{[0^+]}\rangle$ and $\langle t^{[0^-]}\rangle$, of the paths in the $[0^+]$ and $[0^-]$ ensembles,

$$f_A = \frac{1}{\langle t^{[0^+]}\rangle + \langle t^{[0^-]}\rangle}. \quad (8)$$

The equivalence between Eqs. (8) and (7) is explained in Fig. 4.

FIG. 4. Explanation of the path ensembles, the RETIS flux relation Eq. (8), and the crossing probabilities, based on a hypothetical MD run. The RC versus simulation time is shown. **Path ensembles**: The MC scheme in RETIS aims to get the same statistical path distributions as if they would have been cut out of an infinitely long MD simulation. For instance, the blue segments of this MD simulation will give an identical statistical collection of paths as the $[0^-]$ path ensemble obtained by MC moves. The segments that are (partly) red are the same paths collected in the $[0^+]$ ensemble, and the segments that are partly green and partly yellow belong to the $[1^+]$ and $[2^+]$ path ensembles, respectively. The black dashed part originates from state B and is, therefore, not part of any RETIS ensemble. **Flux**: The different path lengths of the blue and (partial) red segments summed up provide $T_{\in\mathcal{A}}$: $\sum_{i=1}^{N_c^+}(t_i^{[0^-]} + t_i^{[0^+]}) = T_{\in\mathcal{A}}$. Substitution of this in Eq. (7) and using $\langle t\rangle = \sum_{i=1}^{n} t_i/n$, we obtain Eq. (8). **Crossing probabilities**: the overall crossing probability $P_A(\lambda_n|\lambda_0)$ equals the fraction of the (partial) red path reaching λ_n or (#paths reaching λ_n/#red paths). Naturally, the following relation holds: (#paths reaching λ_n/#red paths) = (#paths reaching λ_n/#yellow paths) × (#yellow paths/#green paths) × (#green paths/#red paths). This is equivalent to Eq. (6): $P_A(\lambda_n|\lambda_0) = P_A(\lambda_n|\lambda_2) \times P_A(\lambda_2|\lambda_1) \times P_A(\lambda_1|\lambda_0)$. Or, in this particular case, $1/9 = 1 \times 1/3 \times 1/3$.

FIG. 3. Explanation of the generic crossing probability $P\binom{k\,|\,j}{l\,|\,i}$, which is a history dependent conditional probability used in all TIS variations. The condition $|_i^j$ is indicated by red and implies that interface λ_j needs to be crossed in a single MD time step, and, in addition, λ_i should be more recently crossed than λ_j. In other words, it should be a first crossing with λ_j since crossing λ_i. Under this condition, $\binom{k}{l}|$ (indicated by blue) refers to the chance that λ_k will be crossed before λ_l. TIS and RETIS crossing probabilities are based on a special case of this in which $k = j + 1$ and $i = l = 0$: $P_A(\lambda_{j+1}|\lambda_j) = P\binom{j+1\,|\,j}{0\,|\,0}$. PPTIS on the other hand is based on *short memory* crossing probabilities: $p_j^{\pm} = P\binom{j+1\,|\,j}{j-1\,|\,j-1}$, $p_j^{\mp} = P\binom{j-1\,|\,j}{j+1\,|\,j+1}$, $p^{\ddagger} = P\binom{j+1\,|\,j}{j-1\,|\,j+1}$, and $p^{=} = P\binom{j-1\,|\,j}{j+1\,|\,j-1}$.

This figure also shows the essence of the importance sampling in path space. To obtain a single reactive path, MD generates (in Fig. 4) 9 paths in the $[0^-]$ ensemble, 9 paths in the $[0^+]$ ensemble, 3 paths in the $[1^+]$ ensemble, and just one path in the $[2^+]$ ensemble. In real rare event cases, this will be worse: if the overall crossing probability is 10^{-6}, MD needs to sample at least one million paths in the $[0^+]$ and in the $[0^-]$ ensemble. TIS and RETIS typically aim to sample the same number of paths in each ensemble. For instance, if $P_A(\lambda_B|\lambda_A) = 10^{-6}$, we could place 6 interfaces such that $P_A(\lambda_{i+1}|\lambda_i) \approx 0.1$. In that case, we could in principle already get an estimate of the overall crossing probability using just 10 paths (but preferably more) in the $[0^+]$, $[1^+]$, $[2^+]$, $[3^+]$, $[4^+]$, and $[5^+]$ ensembles, basically reducing the minimum number of required paths from 10^6 to just 60.

Essential to sampling the right path distribution (as if cut from an infinite MD simulation) is to utilize proper MC moves that respect detailed balance.[8] In path sampling, the most central MC move is the shooting move.[38] In this move, a time slice of the last accepted path is taken at random. Then, this point is modified by another randomization procedure, for instance, by changing the momenta of this point. Finally, this new point is used to go forward and backward in time using a MD step integrator in order to create a new path.

The detailed-balance relation that must be fulfilled can be written as

$$\frac{P_{\text{gen}}[\mathbf{x}^{(o)} \rightarrow \mathbf{x}^{(n)}]}{P_{\text{gen}}[\mathbf{x}^{(n)} \rightarrow \mathbf{x}^{(o)}]} \frac{P_{\text{acc}}[\mathbf{x}^{(o)} \rightarrow \mathbf{x}^{(n)}]}{P_{\text{acc}}[\mathbf{x}^{(n)} \rightarrow \mathbf{x}^{(o)}]} = \frac{P[\mathbf{x}^{(n)}]}{P[\mathbf{x}^{(o)}]}, \quad (9)$$

where $\mathbf{x}^{(o)}$ and $\mathbf{x}^{(n)}$ denote the old and new paths, respectively, $P[\mathbf{x}]$ is the probability of path $\mathbf{x}$, and $P_{\text{gen}}[\mathbf{x} \rightarrow \mathbf{x}']$ is the probability to generate path $\mathbf{x}'$ starting from $\mathbf{x}$ (to be precise, these should actually be called probability densities).

Following the Metropolis-Hastings scheme, the acceptance rule of the move can be written as

$$P_{\text{acc}}[\mathbf{x}^{(o)} \rightarrow \mathbf{x}^{(n)}]$$
$$= \hat{h}(\mathbf{x}^{(n)}) \min \left[1, \frac{P[\mathbf{x}^{(n)}]}{P[\mathbf{x}^{(o)}]} \frac{P_{\text{gen}}[\mathbf{x}^{(n)} \rightarrow \mathbf{x}^{(o)}]}{P_{\text{gen}}[\mathbf{x}^{(o)} \rightarrow \mathbf{x}^{(n)}]}\right], \quad (10)$$

where $\hat{h}(\mathbf{x}^{(n)})$ is 1 if the new path fulfills the path ensemble's condition (e.g., crossing λ_i for the ensemble $[i^+]$), otherwise it is 0. In the shooting move, the generation probability is a product of different sub-probabilities. These are the probability to select the shooting point, the probability to select new velocities, and the probabilities that path $\mathbf{x}'$ will be created by the MD integrator starting from the modified shooting point. There are many variations possible with respect to the selection of the shooting point and the way the momenta are changed. If the dynamics are stochastic, it is also common to not change the velocities at all and to change only parts of the path by going only forward or only backward in time.

As mentioned above, we want to obtain the same statistical collection of paths as if they were cut out from an endlessly long MD simulation, and, therefore, $P[\mathbf{x}]$ is directly related by the probability that $\mathbf{x}$ can be created by MD. However, as also the generation procedure is based on MD,

nearly all terms contained by $P[\ldots]$ and $P_{\text{gen}}[\ldots]$ cancel.[3] For instance, if the velocities are not changed or completely regenerated using a Maxwell-Boltzmann distribution, the only remaining terms in Eq. (10) are the selection probabilities. Hence, if each time slice of the path has an equal probability to be chosen as the shooting point, Eq. (10) reduces to

$$P_{\text{acc}}[\mathbf{x}^{(o)} \rightarrow \mathbf{x}^{(n)}] = \hat{h}(\mathbf{x}^{(n)}) \min \left[1, \frac{L^{(o)}}{L^{(n)}}\right], \quad (11)$$

where $L^{(o)}$ and $L^{(n)}$ are the path lengths of the old and new paths, respectively. For more general expressions, we refer to Ref. 3.

Equation (11) shows that if the new path becomes relatively too long compared to the previous path, it will be likely rejected. The rejection/acceptance procedure could be carried out analogously to the standard Metropolis procedure: after the new trial path is completed, one takes a random number $\alpha \in [0:1]$ and then accepts the new path if $\alpha < L^{(o)}/L^{(n)}$. However, the efficiency can be improved[26] by drawing the random number α at the start of the MC move and determine a maximum allowed path length before creating the new path

$$L^{\text{max}} = \text{int}[L^{(o)}/\alpha]. \quad (12)$$

Then, the trial move can directly be terminated and rejected whenever it exceeds this maximum, saving a lot of unnecessary MD steps.

To summarize, the shooting move randomly picks a time slice from the old path and makes a slight modification to the phase point, usually randomizing the velocities. From the new phase point, the trajectory is propagated through MD backward and forward in time until crossing λ_A or λ_B. The path (in ensemble $[i^+]$) can be accepted if it starts at λ_A and crosses λ_i at least once while the path length remains within the maximum path length determined at the start of each shooting move. If these conditions are not met, the old path is kept and used again to repeat this procedure. If it is accepted, the old path is replaced by the new one before the procedure continues.

Figure 5 shows the shooting move and all the other different MC moves that are used in RETIS. The time-reversal move does not require any integration because it consists only of changing the time direction of an existing path. Finally, RETIS performs swapping moves (replica exchange) between ensembles. $[i^+] \leftrightarrow [(i+1)^+]$ is accepted if the $[i^+]$-path happened to cross λ_{i+1}. The $[0^+] \leftrightarrow [0^-]$ swap is always accepted and implies exchanging the end- and starting-points of the $[0^-]$ and $[0^+]$ paths. This is sometimes dubbed[39] the *minus move*. After this exchange, new paths will be generated by going forward and backward in time, respectively (see Fig. 5). Although more expensive than the other swapping moves, this move in particular helps increase the ergodicity of the sampling. Overall crossing probabilities constructed from RETIS typically provide much smoother curves that the ones obtained via TIS. In a study on DNA denaturation, using a mesoscopic model,[40] the RETIS approach was found to be a factor 20 more efficient than TIS.[27]
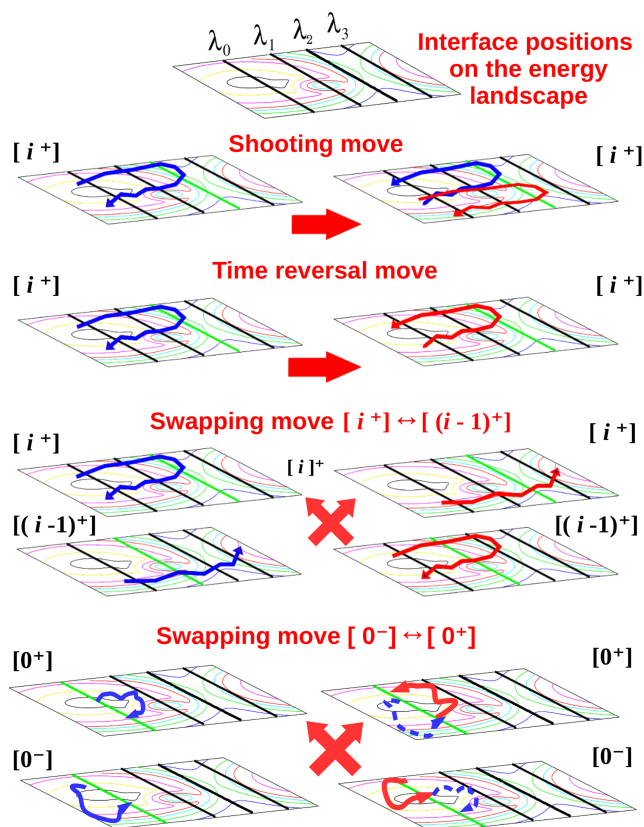
FIG. 5. Schematic picture showing the shooting, time-reversal, and swapping moves. The ensembles are pictured before and after the move, respectively, on the left and right sides. The green interface is the one defining the crossing condition. The blue arrowed lines represent the original paths, while the red ones represent the new paths. Only the shooting and the swapping $[0^-] \leftrightarrow [0^+]$ moves require the generation of MD steps, while the other moves are essentially cost-free.

Although these MC moves formerly obey detailed balance, the numerical trajectories actually might not. In high-dimensional complex systems, the chaotic behavior of the dynamics due to the Lyapunov instability of the dynamics will ensure that any two deterministic MD trajectories with nearly identical starting conditions will ultimately diverge exponentially and might end up at different states A or B. Due to round-off errors, it might be practically impossible to regenerate the old path from the new path. Hence $P_{gen}[\mathbf{x}^{(n)} \rightarrow \mathbf{x}^{(o)}]$ is, strictly speaking, zero in computer generated trajectories if the system is deterministic. Vlugt examined this issue[41] by comparing TPS results using a standard integrator and a bitwise time-reversible integrator and found no difference. The common assumption, even if (to the best of our knowledge) not rigorously proven, is that theorems based on the shadow Hamiltonian[8] derived for simplectic MD integrators also apply for a TPS type of sampling. Hence, the detailed balance relations only have to be satisfied formally as long as the MD integrator is simplectic.

A prerequisite to start the MC sampling is to have an initial path in each ensemble, which needs to be established with some kind of initialization procedure. In some cases, this procedure is not trivial; however, it should be noted that the first path in each ensemble does not have to be a very good one and can even be unphysical. As in any MC method, the first number

of MC cycles is essentially used to equilibrate and is removed from the analysis to compute crossing probabilities and other important quantities. Initial paths could be obtained from high temperature runs or non-dynamical approaches such as nudged elastic band,[42] or by using a shooting algorithm in which the path making less progress than the old one is automatically rejected.

IV. APPLICATIONS OF RETIS

Although RETIS is more efficient than TPS or TIS, the number of studies using the RETIS method is yet somewhat lagging behind due to its non-trivial implementation and the lack of user-friendly pluggable software packages which could provide this technique to scientific users. Recent open-source software projects such as PyRETIS[43,44] and open-path-sampling[45] are presently under development and will hopefully remove this blockade. Also, the naming is not fully standardized or commonly used in the literature as the method was originally referred to as parallel path swapping[27] and the name RETIS was coined in the follow-up paper.[28] In this short overview, we define a RETIS application as a study that uses TIS including the replica exchange moves between the path ensembles.

The RETIS applications published today range from nucleation, chemical reactions, to biological transitions, while the type of dynamics range from Langevin and Brownian motion to classical and even *ab initio* MD. For instance, Lechner *et al.*[46] applied RETIS to study the crystal nucleation of colloidal suspensions, while in another nucleation study by Menzl *et al.*[47] on cavitation in water under tension, the authors compared RETIS to the Bennett-Chandler approach and found that the calculated cavitation rates are in good agreement with classical nucleation theory where the curvature dependence of surface tension is taken into account. Their calculations are also in excellent agreement with inclusion experiments, suggesting that homogeneous nucleation is observed in inclusion, while impurities are the source of heterogeneous nucleation.

There are several examples illustrating the applicability of RETIS on biological systems including isomerization of biological molecules[48] and protein (un)folding[49,50](these studies employed the multiple state and single replica methodologies discussed in Secs. VI D and VI E). A study of the 35-residue-fragment (HP-35) villin headpiece in implicit water demonstrates that this method can be used to study high (un)folding barriers.[49] Further the (un)folding network of the Trp-cage mini-protein in explicit water resulted in a kinetic rate matrix in excellent agreement with IR experiments.[50]

Saroukhani *et al.* applied a range of methods from harmonic transition state theory (HTST) to RETIS in their study on dislocation dynamics in Al-4 wt. %Cu.[51] Their results show that while HTST is not able to predict the rate due to entropic effects, both TIS and RETIS predicted rates in agreement with the MD simulations. Snapshots of the dislocation transition are shown in Fig. 6. The study also shows how the RETIS method is more effective in sampling the different possible reaction mechanisms than TIS, which has a higher tendency to get stuck in a single reaction tube.

FIG. 6. Snapshots of dislocations in Al-4 wt. %Cu overcoming two obstacles at 200 MPa and 300 K. The left image shows the first partial dislocation in the initial configuration, and the right image shows the second partial dislocation in the final configuration as the respective obstacles are overcome. The middle image is a snapshot at the center of the intermediate cells. Reprinted with permission from S. Saroukhani *et al.*, J. Mech. Phys. Solids **90**, 203 (2016). Copyright 2016 Elsevier.

The first realization of a RETIS study using *ab initio* MD was given in Ref. 52, which reports on the water auto-ionization in water clusters. This publication also introduced


FIG. 7. Top: Snapshots of various dissociation mechanisms of the silicate complex in aqueous phase. Hydrogen atoms marked in green and purple are possible candidates for participation in the dissociation process. Na$^+$ ions are colored in blue, and orange oxygen ions belong to the water molecules with a bridging hydrogen bond. Snapshots [(a) and (b)] show mechanisms with hydrogen transfer, while (c) shows a mechanism without hydrogen transfer. Bottom: Crossing probability of the dissociation reaction as a function of the reaction coordinate (Si–O distance). The curve is shown together with the Gibbs free energy curve of Ref. 54 obtained via thermodynamic integration. Reprinted with permission from M. Moqadam *et al.*, Phys. Chem. Chem. Phys. **19**, 13361 (2017). Copyright 2017 Royal Society of Chemistry.

an analysis method for identifying the subtleties of the reaction mechanism that is further discussed in Sec. VI B.

The first real *ab initio* chemical reaction study in aqueous solutions was performed by Moqadam *et al.*[53] This paper analyzes the mechanism and computes the rate of the silicate dimerization reactions in aqueous solutions using a simulation box containing 64 water molecules. The study collected during several months around one hundred thousand trajectories of which several thousands were reactive (20 000 for the dissociation reaction study). Regarding the number of MD steps, this study would correspond to a MD simulation with a total of 1000-1400 ps simulation time. The expected time to observe a *single* reactive event is however as high as 50 000 ps. In other words, the RETIS study gathers the statistical information which would be equivalent to a 20 000 × 50 000 ps = 100 μs plain MD run, far below the reach of *ab initio* molecular dynamics.

As RETIS actually uses the real dynamics of the system without biasing potential energy surfaces or applying constraints, it revealed information on the mechanism that contradicted earlier hypotheses based on free energy studies. For instance, as shown in Fig. 7, the RETIS simulations revealed two possible reaction routes for the dissociation of the silicate ion complex, which is mediated by a proton transfer or not. In this case, the mechanism including proton transfer is highly predominant with a probability of 80%. The fact that RETIS produces reactive trajectories for both mechanisms supports the ergodicity of the method.

Also in another reaction, in which a water group is subtracted from the silicate ion complex, showed two possible mechanisms, where a proton is transferred directly via the silicate complex or via a hydrogen bond network in the solvent. While previous studies have reached contradictory conclusions on which of the two mechanisms prevail, the RETIS simulation generated about two thousand unbiased trajectories showing that the direct proton transfer is only slightly favorable.

V. EFFICIENCY

One way to measure the efficiency of a simulation method is to determine the CPU time needed to obtain the prefixed error. However, CPU time is not a well-defined property as it depends on the hardware, the frequency at which the output is written to the hard disks, and the technical details of the implementation. Yet, as in large-scale classical MD simulations and *ab initio* MD simulations, the calculation of the forces is the most expensive operation, the effective CPU time is generally expressed as the number of force evaluations or, equivalently, the number of MD steps. In this way, algorithms can be compared irrespective of the hardware or technicalities of the implementation. In some studies, such as those dealing with nucleation processes, the calculation of the order parameter is actually more expensive, but this is still fine as the number of evaluations of the order parameter is also proportional to the number of MD steps.

In Ref. 55, the effective CPU time for a relative error of 1 was denoted as τ^{eff}. In practice, this property is computed from a simulation via the following equation:

$$\tau^{\text{eff}} = \text{number of MD steps} \times \text{relative error}^2, \quad (13)$$

where the relative error can be obtained from block averaging or the bootstrap method. The lower the value of τ^{eff}, the more efficient a method is. The *efficiency* is sometimes defined as the inverse of this property.

A. Placement of interfaces

The CPU efficiency time for one of the TIS sub-simulations for computing the crossing probabilities, $p_i = P_A(\lambda_{i+1}|\lambda_i)$, equals[55]

$$\tau_i^{\text{eff}} = \frac{1 - p_i}{p_i} \xi_i \tau_i^{\text{path}} \mathcal{N}_i, \quad (14)$$

where τ_i^{path} is the average path length in path ensemble i, and ξ_i is the ratio between the average length of the trial paths in ensemble i and the physical average path length τ_i^{path}. Generally ξ_i is somewhat lower than 1 as sometimes a path can be rejected after only a few MD steps, e.g., if the backward integration ends in state B. Finally, $\mathcal{N}_i$ is the statistical inefficiency of the sampling. Often, it is assumed that both ξ_i and $\mathcal{N}_i$ are more or less constant for each path ensemble but that the average path length increases for increasing i as the minimal distance, $\lambda_i - \lambda_A$, that needs to be covered in order to fulfill the crossing condition increases as a function of i. If the barrier is steadily increasing at a constant slope, we can assume that placing the interfaces at equidistant separation $\Delta\lambda$ gives equal crossing probabilities $p_i = p$ for each sub-simulation. We assume further that the average path length increases according to a power law $\tau_i^{\text{path}} \propto (\lambda_i - \lambda_A)^g = (i\Delta\lambda)^g$ [more generic: $\xi_i \tau_i^{\text{path}} \mathcal{N}_i \propto (i\Delta\lambda)^g$], where experience shows that g is typically around 1 for steep barriers while it is around 2 for diffusive flat barriers. Based on the error propagation rules $\epsilon_{\text{tot}}^2 = \sum \epsilon_i^2$, one can determine the CPU efficiency of the overall simulation, which is shown[55] to be proportional to

$$\tau^{\text{eff}} \propto \frac{1 - p}{p|\ln p|^2}. \quad (15)$$

The number of interfaces n follows from the overall crossing probability as $p^n = P_A(\lambda_B|\lambda_A)$, and the interface separation is, hence, $\Delta\lambda = (\lambda_B - \lambda_A)/n$. The overall efficiency as a function of p is shown in Fig. 8, normalized to one at its minimum $p = 0.2$. The minimum $p = 0.2$ is a trade-off between two effects. If the separation of interfaces is very large, p will be small and it takes a lot of MD steps to obtain a decent error due to the $\propto 1/p$ dependence in Eq. (14). However, since only few simulations need to be combined, the overall error and, hence, the overall CPU efficiency will not be so much different than the individual errors or CPU efficiencies. On the other hand, if many interfaces are used, it takes a few MD steps per sub-simulation to get low relative errors, but since the total computational cost and the overall error implies summing over the n individual simulations, the overall efficiency will be low. Still, as shown in Fig. 8, there is a wide range of acceptable p values ranging from 0.05 to 0.5 where the overall efficiency is still within 40% of the theoretical optimum. In practice, TIS simulations with relatively few steps could be performed sequentially for $i = 0, 1, \ldots$, and after each simulation, the next interface could


FIG. 8. The overall CPU efficiency time as a function of p [Eq. (15)] for the model system of a continuous increasing barrier with constant slope. Here, $p = P_A(\lambda_{i+1}|\lambda_i)$ equals the crossing probabilities that are being computed in each simulation i, which are assumed to be equal for each i. The CPU efficiency time τ^{eff} is proportional to the number of MD steps required to obtain a predetermined relative error and is here normalized to its minimum at $p = 0.2$.

be defined at the point where the local crossing probability has dropped to 0.2, until the state B is reached. After that, a long TIS or RETIS run can be launched. These runs will probably reveal that some simulations, based on the longer runs, have crossing probabilities significantly different than 0.2. However, as long as these are still within the acceptable range or if there are only a few outside this range, it is probably not worth adapting the interface positions another time at the expense of longer initializations.

Note that the *optimal* $p_i = 0.2$ is merely a rule of thumb rather than an absolute truth as it is based on several assumptions. Also, it should be noted that RETIS might actually have a different optimum. Due to the swapping moves, the different sub-simulations are no longer independent, which implies that the error propagation is more complex depending on covariant terms.[55] In addition, as the RETIS scheme gives trajectories for two ensembles each time the path ensembles [i^+] and [$(i + 1)^+$] can be swapped, i.e., if the trajectory of the ith simulation crosses λ_{i+1}, it might well be more effective to aim for somewhat higher p_i values.

B. Importance of reaction coordinate

The concepts of reaction coordinate (RC), order parameter (OP), or collective variable (CV) have been given slightly different meanings in different communities, which has been a cause of confusion. Some authors would simply use these concepts as synonyms while others would say that there is basically only one RC, which is the ultimate descriptor of the reaction. Generally, this ultimate RC is defined as the committor function though different opinions exist whether this function should be defined in full phase space or configuration space. Basically, the committor defines for each configuration point/phase point a value between zero and one, which corresponds to the chance that a dynamical trajectory starting from

this point ends up in the product state rather than the reactant state.

In other studies, the RC is just the main parameter chosen by the scientist to perform the actual importance sampling, to define the US windows, the constraint planes in TI, or the TIS interfaces, while the order parameter or collective variable could be any other parameter that is used in the post analysis.

Early studies[23–25] have described the TPS method as an approach that does not require a RC but only an order parameter. The interpretation of this statement might be valid or not dependent on which definition one uses for the RC and which type of TPS simulation is meant. If the sole aim of TPS is to gain a statistical collection of reactive pathways, the statement is certainly valid. The order parameter then just defines the stable state regions without quantifying any progress if the system is at the barrier region. In such a simulation, path sampling can be truly viewed as *a blind search in the dark*.[56] If, however, the full reaction rate computation is performed by means of path sampling (TPS, TIS, RETIS, or FFS), the statement is debatable.

In contrast to TST or approximate path sampling approaches, it is true that exact path sampling approaches should give the correct result independent of the choice of the RC. In that sense, one could argue that it does not require an ultimate RC. The same is, however, true for the RF method where it is not common to distinguish between the order parameter and the RC. There is, however, a difference in how the efficiency of the method depends on the chosen reaction coordinate as was shown in Ref. 55.

The main difficulty of finding a proper RC is illustrated in Fig. 9 for the case of NaCl dissociation. Whereas the RC is obvious in the gas phase, simply the distance between the two ions, this is not obvious in aqueous solutions. As shown in Fig. 9, the solvent structure needs to undergo a change during the dissociation process. An *ideal* RC should therefore include some solvent degrees of freedom. The distance alone is not a good descriptor of the reaction since the solvent structure might promote association or dissociation even if the ions are relatively far or close, respectively.[57,58]

We can also refer to Fig. 2 assuming that the *y*-coordinate in this figure is the distance between ions and the *x*-coordinate is some unknown coordinate describing accurately the solvent structure. If the dynamics is only viewed along the *y*-coordinate, it shows that there is a high probability to return to the associated state even if the maximum of the free energy barrier seems to be crossed (MD step 5 in Fig. 2). Reversely, if the transition state surface is crossed from the side of the dissociated state (B), there is again a high probability that the system returns to state *B* (MD step 29 in Fig. 2). Hence, the projection along *y* introduces a memory effect (non-Markovianity), which seems to enhance the likeliness of returning to the state where it comes from.

In free energy based methods, this effect will cause troublesome hysteresis. As all free energy methods provide some kind of force to push the system over the barrier, the system tends to explore a different pathway when pushed forward from the reactant to product state than when pushed from the product state back to the reactant state. This is due to a barrier orthogonal to the chosen RC and each side of this orthogonal barrier should, in principle, be sampled. Even if the TI or US simulation is sufficiently long to sample properly across the orthogonal barriers, the RF approach will still fail since the transmission coefficient becomes negligibly small. These problems would likely disappear when using an ideal RC, which for NaCl should not only depend on the ionic distance but also on the solvent structure.

The quantitative assessment of efficiency versus the *quality* of the RC was discussed in Ref. 55. Analytic expressions of the CPU efficiency time τ^{eff} can only be obtained in highly simplified models. However, to describe the above issues, the model should have at least some minimal complexity in order to make the RC non-trivial. In Ref. 55, such a model was examined, which consists of a single particle in a two-dimensional box. For simplicity, the particle is only allowed to move along the *x*-coordinate. But when colliding against the walls, the particle will obtain randomized velocities from a Maxwell distribution pointing away from the wall and a new random *y*-coordinate (here coined as the "sjoelbak" model). Further, between collisions the particle moves completely adiabatic. Between the two walls, there is a slanted symmetric barrier where the angle θ of the tilt now reflects the quality of *x* as the RC. The model, even if somewhat unphysical, has a dynamics that perfectly obeys the Boltzmann distribution $\rho(x, y) \propto e^{-\beta V(x,y)}$ (with $\beta = 1/k_B T$), has just enough complexity to the innate hysteresis problem, and allows analytic treatment of the overall efficiency using MD, RF, or TIS/RETIS.


FIG. 9. Schematic cartoon illustration of NaCl dissociation. (a) shows the associated state in which the NaCl complex has no net charge and the water molecules in the first solvation shell mainly form hydrogen bonds between themselves. (b) shows the barrier crossing event, which requires the breakage of many hydrogen bonds. Finally, (c) shows the dissociated case where the negatively charged Cl$^-$ and the positively charged Na$^+$ are mainly closely surrounded by the hydrogen atoms and oxygen atoms, respectively.

As a point of reference, the CPU efficiency time of plain MD is exponentially dependent on the height of the barrier and inverse temperature,

$$\tau_{MD}^{eff} \propto e^{\beta \Delta E}, \tag{16}$$

and it is due to this exponential dependence that MD is unfeasible for most relevant processes.

Rare event sampling techniques can improve the scaling, but how good they achieve this depends on the RC. It is informative to compare the cases where the RC is very good and very bad. The parameter to distinguish these two cases is the following:

$$\alpha = 2\beta\Delta E \sin\theta L_y / W, \tag{17}$$

where the width W and length segment L_y are indicated in Fig. 10(b).

In the case that the RC is orthogonal to the rim of the barrier ($\theta = 0 \Rightarrow \alpha = 0$), the system becomes effectively one-dimensional and the TST approximation becomes exact. The computational cost is dominated by the free energy calculation. In Ref. 55, the efficiency time was, hence, calculated for US using rectangular windows. Even if θ is slightly different from zero, then α might become very large. In this case, the hysteresis effect will aggravate the free energy calculation, but the transmission calculation is affected even more. Therefore, we make the following assumption:

$$\tau_{RF}^{eff} = \begin{cases} \tau_{US}^{eff} & \text{if } \alpha \ll 1, \\ \tau_{\kappa}^{eff} & \text{if } \alpha \gg 1. \end{cases} \tag{18}$$


FIG. 10. Illustration of the *shuffleboard* model. (a) The particle (blue disk) can only move along the x-coordinate and it moves frictionless. At each collision with the walls, it will obtain random velocities pointing away from the wall. In addition, it will obtain a new random y-coordinate (via the hands who shuffle the disks like in the Dutch shuffleboard "sjoelbak" game). If the initial kinetic energy is larger than the height of the barrier, it will cross the barrier. Otherwise it will return with the same speed it started with. (b) Top view indicating the parameters θ, W, and L_y. At the top of the free energy along x, the configuration points indicated by the red crosses have the same energy and are, therefore, equally likely. Still, their potential energy is lower than the actual barrier ΔE. The two trajectories (blue lines with arrows) starting from the left wall are not equally likely. The bottom one crosses the barrier implying a higher kinetic energy at the start larger than ΔE. (c) Due to the non-local nature of the paths, one can imagine an effective potential energy surface that is felt by the end point of the paths, which does not have the typical hysteresis shape.

For US using rectangular windows and $\alpha \ll 1$, the following scaling behavior can be found:

$$\tau_{US}^{eff}(\alpha \ll 1) \propto (\beta \Delta E)^2, \tag{19}$$

which is not exponential but only quadratically dependent on $\beta\Delta E$. This makes US orders of magnitude faster if barriers are high. However, for a transmission coefficient with $\alpha \gg 1$, the scaling becomes exponential again,

$$\tau_{\kappa}^{eff}(\alpha \gg 1) \propto \frac{e^{\alpha}}{\sqrt{\alpha}} = \frac{e^{2\beta\Delta E \sin\theta L_y / W}}{\sqrt{2\beta\Delta E \sin\theta L_y / W}}. \tag{20}$$

The reason for this is that there are basically two relatively high probability regions on each side of the orthogonal barrier [see Fig. 10(b)]. Not only does this give rise to hysteresis but it also means that trajectories released from the most likely regions at the *top* of the free energy barrier (projected along x) will almost always result in either $A \rightarrow A$ or $B \rightarrow B$ trajectory, i.e., if the equations of motion are followed backward and forward in time starting from the configuration points indicated by the red crosses with randomized velocities. Hence, the transmission coefficient becomes extremely small implying that there is an astronomical number of trajectories needed to obtain an $A \rightarrow B$ trajectory and to estimate the small transmission coefficient.

In trajectory space, the non-local nature of a path eliminates the hysteresis. The lower path in Fig. 10(b) is not as likely as the upper trajectory. Since the trajectory starts from the left wall and crosses the top of the barrier, the initial kinetic energy has to be very high in order to produce this trajectory. Therefore, this trajectory is less likely than the upper one. As a result, one can say that the global identity of the path starting at A implies that the end point of the trajectory *sees* an effective potential that looks like Fig. 10(c). The effective potential does not have two low energy regions at each side of an orthogonal barrier, which eliminates the hysteresis issue. One can show that the scaling of the CPU efficiency time remains quadratic at all values of θ,

$$\tau_{TIS/RETIS}^{eff} \leq \hat{\tau}_{TIS/RETIS}^{eff} \propto (\beta \Delta E')^2,$$
$$\Delta E' = \Delta E(1 + 2\alpha/\beta\Delta E). \tag{21}$$

Naturally, with this dynamics in which the system can only move along x, the shooting move should imply small position changes in addition to velocity changes.

One should, however, realize that the above analysis does not imply that choosing a reasonable reaction coordinate is not important at all for TIS/RETIS simulations. The efficiency analysis does not include the fact that the acceptance of the MC moves might be lower or show slower decorrelation for an improper reaction coordinate. Still, the relative efficiency compared to RF is expected to increase if the reaction coordinate is not well chosen (using the same reaction coordinate for both methods). The essence of this relative scaling is due to the fact that RF applies an importance sampling to accurately determine the small probability to be at the transition state dividing surface, but no importance sampling scheme is used to compute the transmission coefficient, which can also be very small. TIS/RETIS directly computes the dynamical factor using an importance sampling scheme and since the overall crossing probability is

rather reaction coordinate independent, there is no term that suddenly drops to zero when another reaction coordinate is used. It is, however, possible to use path sampling for computing transmission coefficients instead of crossing probabilities. Examples of such approaches are given in Refs. 59 and 60.

The efficiency scaling for FFS cannot be made for this system as the dynamics is deterministic apart from the wall collisions. However, Ref. 11 shows for a non-deterministic one-dimensional study (corresponding to $\theta = 0$) using Langevin dynamics that FFS can still give misleading results. This is because the system is still effectively two-dimensional in phase space. As shown in Ref. 11, the FFS simulations erroneously suggest that the dynamics is not symmetric with the particle moving up the barrier very slowly and going down very fast.

VI. RECENT DEVELOPMENTS

The TIS/RETIS methodology is versatile, which is not only shown by its applications but also by its spin-off algorithms. In this section, we will describe those recent developments. These are new shooting moves, stone skipping and web throwing, the *predictive power method* as an alternative to the committor function approach for the analysis of the generated paths, and QMMM (quantum mechanic molecular mechanics) in the time domain to combine effectively two different levels of force evaluations (e.g., classical force fields and DFT). Finally, we will discuss in a bit more detail the multiple state TIS (MSTIS) and single replica TIS (SRTIS) that are already well established approaches.

A. New shooting moves

Despite the considerable development from TPS to TIS and RETIS, the principle MC move, shooting, has not changed much during the years. Most suggested variations rely on the velocity modification step and/or on the shooting point selection step.[62–66] However, this has as a result that two consecutive accepted paths always share some almost identical time slices. An approach to decorrelate the paths much faster was suggested in Ref. 61 via the so-called stone skipping and web-throwing moves (Fig. 11). In both the MC moves, each step consists in generating N_{sub} paths that are much shorter in length than the actual paths. After this generation, the final subpath is completed and accepted if it obeys the path ensemble's requirement. Stone skipping is the most efficient move for the increasing part of the potential, while web throwing moves are preferred if both the up-going and down-going parts of the potential are sampled (the hysteresis region). The completed MC move obeys super-detailed balance, which ensures that the sampling is correct. These moves are more expensive than the standard shooting move, but if the final path is accepted, it will be considerably more decorrelated from the previous path than with a shooting move. However, if the final path is not accepted, a lot of MD steps are wasted without any progress. The last point is countered by changing the path weights such that basically all paths can be accepted. The only


FIG. 11. Stone skipping and web throwing moves. Top: illustration of the stone skipping move within the $[i^+]$ ensemble. Subpaths are launched by shooting from the λ_i interface. Subpaths are followed until reaching state B or recrossing λ_i again. In the last case, the end point serves as the initiation point for the next subpath generation. After the completion of a number of subpaths ($N_s = 4$ in this case), the last subpath is completed to become a full path. Bottom: illustration of the web throwing move in the $[j^+]$ ensemble. The subpaths are now being initiated from either λ_j or λ_{sour}. The last interface (surface of unlikely return) ensures that backward integration from the subpaths almost certainly ends up in state A. Like stone skipping, several subpaths are being generated. If they connect λ_j and λ_{sour}, they are used for the next generation. In the end, the last subpath is completed to become a full path. For a detailed description of the moves and the proof of the super-detailed balance relation, we refer to Ref. 61.

possibility of rejection is when the completion of the sub-path produces a full path that ends in state B along both time directions. By minimizing the need to generate full trajectories, the exact RETIS method can become nearly as efficient as the approximate path sampling method such as PPTIS or milestoning. The new shooting moves showed a factor 12 improvement compared to standard shooting in a study for DNA denaturation.[61]

B. Analysis of paths

Systematic approaches to analyze reaction mechanisms in terms of descriptive order parameters have been focused on committor analysis.[56,57,67–71] As mentioned above, the committor is a function of either configuration space or phase space and provides for each configuration/phase point the chance that a MD trajectory starting from this point would end up in state B rather than A. The committor is, by many researchers, considered as the ultimate RC. However, the determination of the committor surfaces is computationally intensive. Generally it is more costly than a rate calculation (an exception is the computation of the average committor, which can directly follow from reweighting of trajectories sampling in RETIS[72]). In addition, it is generally not providing direct applicable physical insight to develop, e.g., new catalysts.

Moreover, in all practical analyses of the committor, only the configuration space is considered, which might lead to missing important dynamical effects. An alternative approach to get that kind of insight was developed in Ref. 52, which does not require any extra simulations besides those already

performed in standard TIS/RETIS. Moreover, it is easier to include velocity dependent parameters in the analysis. The approach views the RC as a simple measure of progress, but then the analysis is focused on the other order parameters by evaluating their usefulness for predicting such progress at early stages in the reaction.

The main function that can be computed based on the data of a RETIS simulation is the predictive power functional $T[q]$. $T[q]$ is a functional that, as the order parameter q itself, can in principle depend on all coordinates and velocities of the particles in the system. For instance, q could be a function that describes the solvent structure (Fig. 9),

The predictive power is given by

$$T_A^{\lambda^c,\lambda^r}[q] = 1 - S_A^{\lambda^c,\lambda^r}[q], \qquad (22)$$

$$S_A^{\lambda^c,\lambda^r}[q] = \frac{1}{\mathcal{P}_A(\lambda^r|\lambda^c)} \int \left(\frac{r^{\lambda^c,\lambda^r}(q) u^{\lambda^c,\lambda^r}(q)}{r^{\lambda^c,\lambda^r}(q) + u^{\lambda^c,\lambda^r}(q)} \right) dq,$$

where λ^c and $\lambda^r > \lambda^c$ are two values for the RC; the first one is used for analyzing the first crossing points for trajectories crossing λ^c and the second one is used to split these trajectories into two groups: the ones reaching λ^r (called partially reactive) and the ones that do not make sufficient progress (unreactive). The first crossing points with λ^c can then be used to construct distributions where $r(q)$ and $u(q)$ give the probabilities that a trajectory crossing through λ^c is crossing this surface at a point q and, respectively, crossing (r) and not crossing (u) λ^r. If there is very little overlap between the $r(q)$ and $u(q)$ distributions, then $S \approx 0$ and $T \approx 1$, implying that the q parameter is very discriminative between partial reactive and unreactive trajectories. The value of q at an early stage (crossing λ^c) can then be used to predict whether the crossing will result in a trajectory reaching λ^r or not. In addition, it can also provide clues on how to enhance the reaction by modifying the external conditions such that crossings with λ^c will preferably occur in the right domain of the q order parameter. Finally, it is worthwhile to note that within the RETIS approach, it is possible to re-weight the trajectories themselves[73,74] in order to obtain the reweighted path ensemble. These trajectories can then be projected on any OP or CV of interest and further analyzed.

C. QuanTIS

The RETIS path ensembles provide a natural approach to split a chemical reaction into different processes that can be treated with different levels of theory. A natural division would be to treat all $[i^+]$ with $i = 0, 1, \ldots, n - 1$ ensembles at the quantum mechanical level (mostly density functional theory) and the $[0^-]$ path ensemble on the level of classical MD. The concept is illustrated in Fig. 12.

The approach QuanTIS resembles QMMM[75] except that QM and MM are now knitted together in the time domain instead of space. Without the replica exchange moves, the approach would, in principle, give the exact crossing probability at the QM description. However, swaps between the $[0^-]$ and $[0^+]$ ensembles are quite essential to speed up the ergodicity of the sampling. A way to keep the crossing probability exact at the QM description is to

**QuanTIS:** Combining MD efficiency with QM accuracy

FIG. 12. Artist impression of hypothetical reaction pathways on a free energy surface. The free energy landscape is metaphorically depicted as a Norwegian fjord landscape where the mountain needs to be crossed to reach the product state. Crossing the barrier implies the making and breaking of chemical bonds, and classical force fields are generally not accurate enough to treat these processes. The QuanTIS approach, therefore, treats all $[i^+]$ path ensembles at the QM level. However, if solute molecules drift away and towards each other in the solvent, there is no real chemistry taking place and classical force fields are the ideal tool to speed up the simulations. The figure shows a hypothetical reaction. At the right-hand side, the two solute molecules are close but do not have the right mutual orientation to react. The blue path is returning to the reactant state where it is *swapped* to the $[0^-]$ ensemble and continues with fast classical MD. Once the solutes meet again, their mutual orientation is much more favorable for the reaction to take place. A swap back to the $[0^+]$ ensemble even leads to a full reaction in this hypothetical case.

change the acceptance rule for this swap by the following equation:

$$P_{\text{acc}}(r_{\text{MM}}, r_{\text{QM}}) = \min \left[1, \exp \left(- \frac{1}{k_B T} \left[V_{\text{QM}}(r_{\text{MM}}) \right. \right. \right.$$
$$\left. \left. \left. - V_{\text{MM}}(r_{\text{MM}}) + V_{\text{MM}}(r_{\text{QM}}) - V_{\text{QM}}(r_{\text{QM}}) \right] \right) \right],$$
$$(23)$$

where V_{QM} is the potential at the QM level used in the $[0^+]$ ensemble (i.e., the potential energy is obtained from a QM calculation), and V_{MM} is the classical potential based on a force field used in the $[0^-]$ ensemble. r_{QM} and r_{MM} are the configuration parts of the phase points in the $[0^+]$ and $[0^-]$ ensembles, respectively, which are being swapped.

In the normal RETIS, $[0^-] \leftrightarrow [0^+]$ is always accepted. The acceptance rule, Eq. (23), ensures that $P_A(\lambda_B|\lambda_A)$ will converge to the same value as with RETIS using QM in all ensembles. Only the flux will be slightly different. However, for systems with many atoms, the acceptance of the swap will be very low since it depends on an exponent of absolute energy differences. Hence, this acceptance drops with system size. However, since systems used for *ab initio* MD are relatively small, this issue might be problematic but not insurmountable. For large systems, the approach could effectively be combined with spatial QMMM so that the potential energy difference is only due to a relatively small group of atoms.

Reference 76 discusses several tactics to keep the acceptance high in realistic systems. One solution is to mix the classical and QM potentials close to the λ_0 interface or to optimize the classical potential on-the-fly by adjusting the

force field parameters. Another practical approach is to break slightly the detailed balance and basically accept the move using a less stringent criterion. The results of Ref. 76 show that an *accept all* only affected slightly the initial crossing probabilities $P_A(\lambda_1|\lambda_0)$ and $P_A(\lambda_2|\lambda_1)$. Hence, the speed up in sampling and reduction in the statistical error outweighed the negative aspect of introducing a marginal systematic error.

D. Multiple state TIS

The original transition path sampling technique was designed to sample pathways connecting only two distinct stable states. In complex systems, multiple meta-stable minima can exist, even if the overall kinetics is two-state. For instance, a protein folding/unfolding transition is likely to encounter intermediate metastable states that are long lived on the molecular scale, while still short-lived with respect to the overall relaxation time. Direct path sampling between the unfolded and folded states can then become very inefficient as trial trajectories will get stuck in an intermediate state from which the escape itself is a rare event. While one can wait until such escapes from the intermediate state occur, even with reasonable acceptance using the stochastic one-way shooting moves[65] or the more recent *spring shooting*,[66] this might result in extremely long MD trajectories, thus lowering the overall path sampling efficiency dramatically. A simple solution to this problem is to conduct individual path sampling simulations for each sub-reaction.[77] While such a solution might work for a few well-defined intermediate states, the number of transitions scales as $N(N-1)$ for N states, and moreover the acceptance of pathways suffers due to the rejection of trajectories that do not connect the selected states. Multiple state TPS (MSTPS) and multiple state TIS (MSTIS) solve this problem by allowing the sampling of pathways that connect any two stable or intermediate states within one single path sampling simulation.[78] This approach therefore views the rare event process as a Markov State Model (MSM)[79–82] with transitions between each pair of states and a loss of memory in-between jumps. Different to PPTIS is that the intermediate states can be defined in a multidimensional collective variable space. On the other hand, PPTIS does not require well-defined intermediates as memory is included via the history dependence of the crossing probabilities (see Fig. 3). Both methods, strictly speaking, are not exact path sampling methods, unless the memory loss assumption is fully justified.

MSTPS/MSTIS (see Fig. 13) is a method to sample all transitions in such a MSM. The difference between MSTPS and MSTIS is the same as that between TPS and TIS: MSTPS is used to collect reactive pathways only, while MSTIS aims at performing an entire rate computation.

Analogous to TIS, the rate constant for a transition from a state $\mathcal{I}$ to a state $\mathcal{J}$ is

$$k_{\mathcal{IJ}} = \left[f_{\mathcal{I}} \prod_{s=0}^{m-1} P_{\mathcal{I}}(\lambda_{(s+1)\mathcal{I}}|\lambda_{s\mathcal{I}}) \right] P_{\mathcal{I}}(\lambda_{0\mathcal{J}}|\lambda_{m\mathcal{I}}) \qquad (24)$$

with $P_{\mathcal{I}}(\lambda_{0\mathcal{J}}|\lambda_{m\mathcal{I}})$ being the conditional probability that when the outermost interface $\lambda_{m\mathcal{I}}$ is crossed $\lambda_{0\mathcal{J}}$ will also be crossed,


FIG. 13. Cartoon of multiple state transition path (interface) sampling. Each stable state has a set of λ-interfaces with λ_m being the outermost interface (shown here as red dashed lines and the other interfaces are suppressed). Trajectories starting from one of the stable states, crossing the corresponding λ_m-interface, and ending in any of the other stable states contribute to the path ensemble (blue solid trajectories). Paths that return to their initial state without crossing the λ_m-interface will be rejected during the sampling (green dotted trajectory). Reprinted with permission from J. Rogal and P. G. Bolhuis, J. Chem. Phys. **129**, 224107 (2008). Copyright 2008 American Institute of Physics.

that is, state $\mathcal{J}$ is reached before the system returns to $\mathcal{I}$. The factor between square brackets denotes the flux through $\lambda_{m\mathcal{I}}$, which follows from a "regular" TIS simulation using the set of λ-interfaces for state $\mathcal{I}$ and needs to be evaluated only once for each stable state. The crossing probability $P_{\mathcal{I}}(\lambda_{0\mathcal{J}}|\lambda_{m\mathcal{I}})$ is computed in a single MSTIS run, from the number of paths, $n_{\mathcal{IJ}}$, starting from $\mathcal{I}$, crossing $\lambda_{m\mathcal{I}}$, and ending in $\mathcal{J}$, and dividing by the sum of all pathways starting from $\mathcal{I}$ and crossing $\lambda_{m\mathcal{I}}$,

$$P_{\mathcal{I}}(\lambda_{0\mathcal{J}}|\lambda_{m\mathcal{I}}) \approx \frac{n_{\mathcal{IJ}}}{\sum_{\mathcal{J}} n_{\mathcal{IJ}}}. \qquad (25)$$

As a transition $\mathcal{I} \to \mathcal{K} \to \mathcal{J}$ is described as two independent transitions $\mathcal{I} \to \mathcal{K}$ and $\mathcal{K} \to \mathcal{J}$, it is important that the defined states are really metastable so that the system spends sufficient time in these states to lose memory and guarantee Markovianity. As long as all stable states are long-lived with respect to the molecular time scale, this should not be a problem.

MSTPS/MSTIS requires only one path simulation to sample all possible transitions, instead of one path sampling simulation for each transition between each stable state pair. Moreover, acceptance ratios are higher because all trajectories connecting any two states contribute to the ensemble. In addition, decorrelation of subsequent pathways is faster due to switching between different types of pathways. This switching is essential. If the switching fails to occur, the original 2-state sampling is recovered. It follows that multiple state TPS should be employed in cases where the switching actually poses a problem to the 2-state sampling, i.e., systems with a low acceptance ratio due to trajectories that do not return to one of the two stable states.

MSTPS samples trajectories with their correct weights, i.e., transitions out of the same stable state; the ratio of pathways is equal to the ratio of the transition probabilities. When the transition probabilities are significantly different, e.g., if one transition is orders of magnitude more probable than

FIG. 14. Cartoon of the general principle of the single replica TIS for a simple 4-state system, where each state has two interfaces. Starting from the left top corner, the initial path crosses not only the current (black) interface but also the outermost interface, thus allowing an exchange between these interface ensembles. Subsequently, a shooting move happens to create a trajectory that reaches another stable state, enabling a state swap. After this swap, the outermost interface of the new state becomes the current interface. An interface exchange followed by a shooting move then leads to the red path depicted in the lower left corner cartoon. Continuing random combinations of such moves employing the Wang-Landau approach allows equal sampling of all states and interfaces. Reprinted with permission from W.-N. Du and P. G. Bolhuis, J. Chem. Phys. **139**, 044105 (2013). Copyright 2013 American Institute of Physics.

others, a biasing scheme based on the Wang-Landau sampling approach[73,83,84] can help enhance rare transitions. Alternatively, fast transitions can be excluded from the sampling.[85] When the TIS part of the MSTIS is suffering form sampling problems, MSTIS can be combined with RETIS, as was done in Ref. 39.

An additional challenge with the multiple state formalism is the choice of the order parameter λ. What is a good order parameter for one transition might not be a good order parameter for another transition, even for the same initial state. Within the MSTIS/RETIS approach, it is possible to combine multiple sets of interfaces based on different order parameters for each transition.[39]

E. Single replica TIS

While the combination of MSTIS and RETIS is very powerful, a practical implementation is not easy for large systems due to the multiple states and the large number of interfaces to be treated simultaneously. Moreover, the disparity in the length of the pathways makes a parallel implementation nontrivial. A software package to easily set up and conduct MSTIS with RETIS is currently under construction.[45] Another solution to this problem is to abandon the simultaneous simulation of all interfaces and only consider a single replica. Applying concepts from simulated tempering,[86,87] single replica transition interface sampling (SRTIS) can sample all interfaces in the MSTIS framework. SRTIS, instead of exchanging paths between replicas, exchanges the interfaces. Starting with an initial path in an initial interface, regular TIS shooting moves sample new paths. Occasionally, an interface swap move tries to switch to a neighboring interface and accepts this move when the path obeys the interface criteria (see Fig. 14). Straightforward application of such a scheme would strongly favor the interface near stable states and probably fail to sample barrier crossings. A Wang-Landau biasing scheme can create a flat sampling of interfaces via construction of the *density of paths* (DOP) at each interface. This approach allows an adaptive scheme in which stable states are added as the sampling of path space progresses. As the trajectory dynamics is unbiased, it can be combined via a reweighting scheme to yield the rate, free energy, and mechanism (see, e.g., Refs. 46, 73, and 74). Convergence of the Wang-Landau algorithm can be slow. It turns out that the optimal DOP is equal to the crossing probability. Indeed, a fixed bias based on the crossing probabilities tremendously speeds up convergence.[48]

The interface swap move only applies to an interface set belonging to a single initial state, where the paths all are required to start from the same state. A state swap move attempts to change the current initial state to a different state requiring a path-reversal of paths that connect two different states. After the state swap, the set of interfaces belonging to the new initial state is used. The acceptance probability for such a state swap again involves the DOPs of the two interfaces. While it is easiest to only allow state swaps between the outermost interfaces, it is also possible to swap interfaces with identical index or even allow for all-interface state swap. The latter is particularly advantageous when states are nested within interfaces (see Sec. VI F).

Similar to RETIS, it is useful to randomize within the stable states using the additional interface ensemble $[0^-]$,[27] by exchanging paths with the interface ensemble $[0^+]$ (minus move). In summary, SRTIS requires a set of five different path moves: shooting, interface swap, state swap, reversal, and minus moves (see Fig. 14). The method has been used to sample protein folding[50] and complex formation in patchy particle systems.[88,89]

F. Rate constant calculation for nested states

Straightforward application of Eq. (24) becomes invalid for complex systems with some states nested in-between interfaces of other states, as Eq. (24) assumes that transitions occur beyond the outermost interface λ_{mI}. One way of avoiding this

problem is via the path-type numbers introduced in Ref. 50. A path-type number $n_{\mathcal{I}\mathcal{J}}^i(\lambda_{k\mathcal{I}})$ is the number of paths in replica i joining states $\mathcal{I}$ and $\mathcal{J}$ that have crossed at maximum interface $\lambda_{k\mathcal{I}}$ (and thus by definition also all interfaces below k), where the superscript i indicates that the paths should obey the condition of replica i in the ensemble. Having set the maximum interface, we are allowed to reweight these numbers using the WHAM (weighted histogram analysis method) weights obtained from reweighting of the crossing probability and sum them over all interfaces k to yield the reweighted number of paths coming from state $\mathcal{I}$ and ending in state $\mathcal{J}$,

$$\tilde{n}_{\mathcal{I}\mathcal{J}} = \sum_{k=1}^{m} \bar{w}_{\mathcal{I}}^k \sum_{i=1}^{m} n_{\mathcal{I}\mathcal{J}}^i(\lambda_{k\mathcal{I}}), \tag{26}$$

where $\bar{w}_{\mathcal{I}}^k = (\sum_l^k 1/w_{\mathcal{I}}^l)^{-1}$, with $w_{\mathcal{I}}^l$ being the optimized WHAM weights for paths that have crossed interface $\lambda_{k\mathcal{I}}$ at maximum (note that these should be the same as the weights $w_{\mathcal{I}}^l$ obtained via the crossing probability).

Using the Wang-Landau scheme requires a reweighting of the states as well. While this is possible using correction factors based on the relative stability of the states, another approach makes use of the fact that in an unbiased ensemble each $\mathcal{I}\mathcal{J}$ path is as probable as the reversed $\mathcal{J}\mathcal{I}$ path. Therefore we split the path-type matrix, $\tilde{n}_{\mathcal{I}\mathcal{J}}$, into N separate matrices (one for each state) and symmetrize the $\mathcal{I}$th matrix: $\tilde{n}_{\mathcal{J}\mathcal{I}} = \tilde{n}_{\mathcal{I}\mathcal{J}}$. Setting all other entries of the $\mathcal{I}$th matrix to zero results in N matrices with only a nonzero $\mathcal{I}$th row and a nonzero $\mathcal{J}$th column. Joining these matrices with WHAM yields the individual weights for each state and to a $N \times N$ transition path type number matrix, $\tilde{n}_{\mathcal{I}\mathcal{J}}$. Normalization of this with the total number of paths going out of a state $\sum_{\mathcal{J}} \tilde{n}_{\mathcal{I}\mathcal{J}}$ gives the transition probability matrix $P(\lambda_{0\mathcal{J}}|\lambda_{1\mathcal{I}}) = \tilde{n}_{\mathcal{I}\mathcal{J}}/\sum_{\mathcal{J}} \tilde{n}_{\mathcal{I}\mathcal{J}}$, which can be directly used in Eq. (24), and takes into account all nested states.

VII. CONCLUSIONS

We outlined rare event sampling techniques with a focus on exact path sampling approaches and in particular the RETIS method and its developments. In the first part of the article, we described the basics of the theoretical foundations of RETIS and gave an overview on some of its recent applications. The second part is dedicated to the new developments that increase the applicability of the path sampling techniques to a broader range of problems. In addition, we discussed efficiency scaling issues and some recent algorithmic developments of the path sampling methods. This shows that exact path sampling, by means of RETIS, is getting more and more efficient and also the versatility of the approach has grown in the last few years. Thanks to the freely available open-source software packages like PyRETIS[44] and OPenPathSampling,[45] the RETIS approach will become accessible to non-developers. A further expansion of the popularity of the approach is, therefore, expected.

ACKNOWLEDGMENTS

[1]P. Gretener, AAPG Bull. **51**, 2197 (1967).
[2]F. Jeltsch *et al.*, Proc. R. Soc. B **264**, 495 (1997).
[3]B. Peters, *Reaction Rate Theory and Rare Events* (Elsevier, Amsterdam, The Netherlands, 2017).
[4]V. Tsoukalaa *et al.*, Oceanologia **58**, 71 (2016).
[5]P. Embrechts, C. Klüppelberg, and T. Mikosch, *Modelling Extremal Events: For Insurance and Finance* (Springer, Dordrecht, London, 1997), Vol. 33.
[6]A. Clauset and R. Woodard, Ann. Appl. Stat. **7**, 1838 (2013).
[7]D. Kashchiev, *Nucleation: Basic Theory With Applications* (Butterworth Heinemann, Burlington, MA, USA, 2003), pp. 3–16.
[8]D. Frenkel and B. Smit, *Understanding Molecular Simulations From Algorithms to Applications* (Academic Press, San Diego, California, USA, 2002).
[9]P. G. Bolhuis, C. Dellago, and D. Chandler, Faraday Discussions **110**, 421–436 (1998).
[10]P. G. Bolhuis and C. Dellago, *Reviews in Computational Chemistry* (John Wiley & Sons, Inc., Hoboken, NJ, USA, 2010), pp. 111–210.
[11]T. van Erp, Adv. Chem. Phys. **151**, 27 (2012).
[12]G. Torrie and J. Valleau, J. Comput. Phys. **23**, 187 (1977).
[13]E. Darve and A. Pohorille, J. Chem. Phys. **115**, 9169 (2001).
[14]G. D. Rodriguez, E. Darve, and A. Pohorille, J. Chem. Phys. **120**, 3563 (2004).
[15]L. Mones, N. Bernstein, and G. Csányi, J. Chem. Theory Comput. **12**, 5100 (2016).
[16]D. E. Shaw *et al.*, Science **330**, 341 (2010).
[17]R. G. Mullen, J.-E. Shea, and B. Peters, J. Chem. Phys. **140**, 041104 (2014).
[18]D. Moroni, P. Bolhuis, and T. van Erp, J. Chem. Phys. **120**, 4055 (2003).
[19]A. K. Faradjian and R. Elber, J. Chem. Phys. **120**, 10880 (2004).
[20]D. Moroni, T. S. van Erp, and P. G. Bolhuis, Phys. Rev. E **71**, 056709 (2005).
[21]P. Majek and R. Elber, J. Chem. Theory Comput. **6**, 1805 (2010).
[22]E. Vanden-Eijnden, M. Venturoli, G. Ciccotti, and R. Elber, J. Chem. Phys. **129**, 174102 (2008).
[23]C. Dellago, P. G. Bolhuis, F. S. Csajka, and D. Chandler, J. Comput. Phys. **108**, 1964 (1998).
[24]P. G. Bolhuis, C. Dellago, P. L. Geissler, and D. Chandler, J. Phys.-Conden. Mat. **12**(8A), A147–A152 (2000).
[25]C. Dellago, P. G. Bolhuis, and D. Chandler, J. Chem. Phys. **110**, 6617 (1999).
[26]T. van Erp, D. Moroni, and P. Bolhuis, J. Chem. Phys. **118**, 7762 (2003).
[27]T. van Erp, Phys. Rev. Lett. **98**, 268301 (2007).
[28]P. G. Bolhuis, J. Chem. Phys. **129**, 114108 (2008).
[29]T. E. Booth and J. S. Hendricks, Nucl. Technol.-Fusion **5**, 90 (1984).
[30]P. Melnik-Melnikov and E. Dekhtyaruk, Probab. Eng. Mech. **15**, 125 (2000).
[31]M. Villenaltamirano and J. Villenaltamirano, in *Queueing, Performance and Control in Atm North-Holland Studies in Telecommunication*, edited by J. W. Cohen and C. D. Pack (Elsevier Science Publisher B. V., Amsterdam, 1991), Vol. 15, pp. 71–76, 13th International Teletraffic Congress (ITC-13), Copenhagen, Denmark, 19-26 June 1991.
[32]H. Jonsson, Proc. Natl. Acad. Sci. U. S. A. **108**, 944 (2011).
[33]E. Wigner, Trans. Faraday Soc. **34**, 29 (1938).
[34]J. Keck, Discuss. Faraday Soc. **33**, 173 (1962).
[35]H. Eyring, J. Chem. Phys. **3**, 107 (1935).
[36]C. H. Bennett, ACS Symp. Ser. **46**, 63 (1977).
[37]D. Chandler, J. Chem. Phys. **68**, 2959 (1978).
[38]C. Dellago, P. Bolhuis, and D. Chandler, J. Chem. Phys. **108**, 9236 (1998).
[39]D. W. Swenson and P. G. Bolhuis, J. Chem. Phys. **141**, 044101 (2014).
[40]T. Dauxois, M. Peyrard, and A. R. Bishop, Phys. Rev. E **47**, R44 (1993).
[41]T. J. H. Vlugt, Ph.D. thesis, Universiteit van Amsterdam, 2000.
[42]G. Henkelman and H. Jonsson, J. Chem. Phys. **113**, 9978 (2000).
[43]A. Lervik, E. Riccardi, and T. S. van Erp, J. Comput. Chem. **38**(28), 2439–2451 (2017).
[44]PyRETIS: rare events in Python, www.pyretis.org, 2017.
[45]OpenPathSampling: A Python library to facilitate path sampling algorithms, openpathsampling.org, 2017.
[46]W. Lechner, C. Dellago, and P. G. Bolhuis, J. Chem. Phys. **135**, 154110 (2011).
[47]G. Menzl *et al.*, Proc. Natl. Acad. Sci. U. S. A. **113**, 13582 (2016).
[48]W.-N. Du and P. G. Bolhuis, J. Chem. Phys. **139**, 044105 (2013).
[49]W. Du and P. G. Bolhuis, Biophys. J. **108**, 368 (2015).
[50]W. Du and P. G. Bolhuis, J. Chem. Phys. **140**, 195102 (2014).
[51]S. Saroukhani *et al.*, J. Mech. Phys. Solids **90**, 203 (2016).
[52]T. S. van Erp, M. Moqadam, E. Riccardi, and A. Lervik, J. Chem. Theory Comput. **12**, 5398 (2016).
[53]M. Moqadam *et al.*, Phys. Chem. Chem. Phys. **19**, 13361 (2017).

[54] M. Moqadam *et al.*, J. Chem. Phys. **143**, 184113 (2015).
[55] T. S. van Erp, J. Chem. Phys. **125**, 174106 (2006).
[56] P. G. Bolhuis, D. Chandler, C. Dellago, and P. Geissler, Annu. Rev. Phys. Chem. **53**, 291 (2002).
[57] P. L. Geissler, C. Dellago, and D. Chandler, J. Phys. Chem. B **103**, 3706 (1999).
[58] R. G. Mullen, J.-E. Shea, and B. Peters, J. Chem. Theory Comput. **10**, 659 (2014).
[59] J. Juraszek, G. Saladino, T. S. van Erp, and F. L. Gervasio, Phys. Rev. Lett. **110**, 108106 (2013).
[60] K. V. Klenin, J. Chem. Phys. **141**, 074103 (2014).
[61] E. Riccardi, O. Dahlen, and T. S. van Erp, J. Phys. Chem. Lett. **8**, 4456 (2017).
[62] M. Gruenwald, C. Dellago, and P. L. Geissler, J. Chem. Phys. **129**, 194101 (2008).
[63] R. G. Mullen, J. E. Shea, and B. Peters, J. Chem. Theory Comput. **11**, 2421 (2015).
[64] C. N. Rowley and T. K. Woo, J. Chem. Phys. **131**, 234102 (2009).
[65] P. G. Bolhuis, J. Phys.: Condens. Matter **15**, S113 (2003).
[66] Z. F. Brotzakis and P. G. Bolhuis, J. Chem. Phys. **145**, 164112 (2016).
[67] C. Dellago, P. G. Bolhuis, and P. L. Geissler, Adv. Chem. Phys. **123**, 1 (2002).
[68] B. Peters and B. L. Trout, J. Chem. Phys. **125**, 054108 (2006).
[69] B. Peters, G. T. Beckham, and B. L. Trout, J. Chem. Phys. **127**, 034109 (2007).
[70] R. Best and G. Hummer, Proc. Natl. Acad. Sci. U. S. A. **102**, 6732 (2005).
[71] E. Weinan, W. Ren, and E. Vanden-Eijnden, Chem. Phys. Lett. **413**, 242 (2005).
[72] P. G. Bolhuis and W. Lechner, J. Stat. Phys. **145**, 841 (2011).
[73] J. Rogal and P. G. Bolhuis, J. Chem. Phys. **133**, 034101 (2010).
[74] W. Lechner *et al.*, J. Chem. Phys. **133**, 174110 (2010).
[75] A. Warshel and M. Levitt, J. Mol. Biol. **103**, 227 (1976).
[76] A. Lervik and T. S. van Erp, J. Chem. Theory Comput. **11**, 2440 (2015).
[77] J. Juraszek, J. Vreede, and P. G. Bolhuis, Chem. Phys. **396**, 30–44 (2012).
[78] J. Rogal and P. G. Bolhuis, J. Chem. Phys. **129**, 224107 (2008).
[79] S. Park and V. S. Pande, J. Chem. Phys. **124**, 054118 (2006).
[80] F. Noé, I. Horenko, C. Schütte, and J. C. Smith, J. Chem. Phys. **126**, 155102 (2007).
[81] J. D. Chodera *et al.*, J. Chem. Phys. **126**, 155101 (2007).
[82] J. H. Prinz *et al.*, J. Chem. Phys. **134**, 174105 (2011).
[83] F. Wang and D. P. Landau, Phys. Rev. Lett. **86**, 2050 (2001).
[84] F. Wang and D. P. Landau, Phys. Rev. E **64**, 056101 (2001).
[85] W. N. Du, K. A. Marino, and P. G. Bolhuis, J. Chem. Phys. **135**, 145102 (2011).
[86] E. Marinari and G. Parisi, Europhys. Lett. **19**, 451 (1992).
[87] M. Hagen *et al.*, J. Phys. Chem. B **111**, 1416 (2007).
[88] A. C. Newton, J. Groenewold, W. K. Kegel, and P. G. Bolhuis, Proc. Nat. Acad. Sci. U. S. A. **112**, 15308 (2015).
[89] A. C. Newton, J. Groenewold, W. K. Kegel, and P. G. Bolhuis, J. Chem. Phys. **146**(23), 234901 (2017).