EPJ manuscript No.
Scale-Bridging Techniques in Molecular Simulation

Practical and conceptual path sampling issues

P.G. Bolhuis[1,a] and C. Dellago[2]

[1]van 't Hoff Institute for Molecular Sciences and Amsterdam Center for Multiscale Modeling, Universiteit van Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands
[2]Faculty of Physics, University of Vienna, Boltzmanngasse 5, 1090 Wien, Austria

Abstract. In the past 15 years transition path sampling (TPS) has evolved from its basic algorithm to an entire collection of methods and a framework for investigating rare events in complex systems. The methodology is applicable to a wide variety of systems and processes, ranging from transitions in small clusters or molecules to chemical reactions, phase transitions, and conformational changes in biomolecules. The basic idea of TPS is to harvest dynamical unbiased trajectories that connect a reactant with a product, by a Markov Chain Monte Carlo procedure called shooting. This simple importance sampling yields the rate constants, the free energy surface, insight in the mechanism of the rare event of interest, and by using the concept of the committor, also access to the reaction coordinate. In the last decade extensions to TPS have been developed, notably the transition interface sampling (TIS) methods, and its generalization multiple state TIS. Combination with advanced sampling methods such as replica exchange and the Wang-Landau algorithm, among others, improves sampling efficiency. Notwithstanding the success of TPS, there are issues left to discuss, and, despite the method's apparent simplicity, many pitfalls to avoid. This paper discusses several of these issues and pitfalls: the choice of stable states and interface order parameters, the problem of positioning the TPS windows and TIS interfaces, the matter of convergence of the path ensemble, the matter of kinetic traps, and the question whether TPS is able to investigate and sample Markov state models. We also review the reweighting technique used to join path ensembles. Finally we discuss the use of the sampled path ensemble to obtain reaction coordinates.

1 Introduction

Using molecular simulation, a short hand for particle-based computer simulation, one can predict not only static properties of matter, such as equilibrium thermodynamics and mechanics, but also microscopic dynamical properties and kinetics[1]. The method of choice for the latter is molecular dynamics (MD), a microscopic time evolution of Hamilton's equation of motion for all particles in the system[2]. However, while force-field based all-atom molecular dynamics often demands short time steps, many processes in nature take place on much longer time scales. These timescales are often

[a] e-mail: p.g.bolhuis@uva.nl

caused by high free energy barriers between meta-stable states. Classical examples of this are nucleated phase transitions, chemical reactions, and biomolecular isomerisation. The long time-scales often make a direct approach of MD impractical. Therefore many computational techniques were developed to overcome the high barriers, for instance by biasing the system to move along a coordinate which drives the system over the barrier[1]. However, the choice of coordinate can severly alter the outcome of the simulation (see Ref. [3] for a discussion). To avoid this problem one can, instead of sampling configurations, focus on pathways between two states. By sampling these pathways one could get insight in the mechanism of the reaction. Action sampling[4], nudged elastic band[5,6], string methods[7,8] all belong to this category. In the late 1990's Chandler and coworkers invented the transition path sampling method as a means to investigate the mechanism of an activated event in a complex system[9–12, 3]. The method creates, via a random walk Markov Chain Monte Carlo process, an ensemble of unbiased trajectories between predefined initial and final states, which can be analyzed in terms of mechanism and kinetics. TPS achieves this by modifying an existing pathway by a procedure called shooting, and accepting or rejecting according to a Metropolis rule[9,10]. Since its inception the transition path sampling method has become a framework for investigating rare events using trajectories. Because its premise is rather straightforward, namely, enhanced sampling of paths, there have been many improvements and altered methods based on TPS, notably TIS[13–15], PPTIS[16], FFS[17,18], NEUS[19], RETIS[20,21], and STePS[22]. However, this review is not aimed at giving an in depth overview of the different methodologies for trajectory-based sampling. We refer the readers interested in such a review to Ref.[23]. Also, we do not aim to provide a technical overview of the methods. For that we refer to Refs.[24,25]. Rather, we discuss some issues that have turned out to be important over the years in implementing the method, despite its apparent simplicity. We also discus some misconceptions, and list some intrinsic problems in TPS of complex systems, and their solutions. The paper is organized as follows. We start with a short recollection of the advantages of path sampling. Then, in section 3, we discuss several issues, misconceptions, and problems arising in the sampling of paths. In section 4 we give a short (slightly technical) overview of the reweighting schemes used for TIS. Finally we discuss strategies for reaction coordinate analysis in section 5. We end with concluding remarks.

2 Advantage of TPS and related methods

2.1 No reaction coordinate required

The TPS method biggest selling point is that instead of imposing the reaction coordinate by biasing along a chosen collective variable (named the order parameter[1]), the method only requires a definition for the stable states and selects its own unbiased transition trajectories [3]. Randomly selecting a 'shooting point configuration from an initial trajectory TPS creates a new trial trajectory by integrating the equations of motions of the unbiased dynamics forward and backward in time. This unbiased dynamics can be deterministic, or stochastic, or any other dynamics, as long as it

[1] In this work we distinguish between the concepts of *order parameter*, *collective variable* and *reaction coordinate*. A collective variable is any function of the particle coordinates, whereas the term order parameter is here reserved for a collective variable that describes a high dimensional hyper-surface to define a stable state or interface, or to drive a biased free energy computation. A reaction coordinate is a (combination of) collective variable(s) that describes the reaction progress. See section 5 for a discussion.

obeys microscopic reversibility. The new trial trajectory can be accepted according to a Metropolis rule. In the simplest version, the Metropolis rule amounts to accepting a trial path when it connects the stable states, and reject. In this way, the technique leads to the same rare event trajectories, as a straightforward MD would have done, only exponentially faster, at a fraction of the computational cost. In this way, the technique leads to the same rare event trajectories as a straightforward MD would have done, only exponentially faster, at a fraction of the computational cost. In a post-processing analysis the reaction coordinate can be deduced from the resulting path ensemble, for instance using neural network technology, machine learning, or simple likelihood maximization (see section 5). Even simply plotting trajectories in different collective variable spaces already can lead to much qualitative insight, for instance about the order of events, or whether a certain transition is fast or slow. Because an ensemble creates many paths, one quickly looses track of the individual path. To show the most likely mechanism one can make us of a path density plot, which shows the density of paths as function of multiple collective variables[2][26,27].

2.2 TPS and TIS give correct rate (recrossings are counted)

Slowly constraining the path ensemble from being completely free toward a transition path ensemble connecting the initial and final state, a procedure akin to a reversible work integration, allows extraction of the rate constant for any activated process[9, 12]. In this integration the transmission coefficient correction to the transition state theory rate constant estimate is taken into account exactly. The TPS rate computation thus accounts for recrossings, trajectories that cross the barrier but still return to their state of origin. In contrast, a transition state theory based approach relying on computation of the free energy barrier, does not take such recrossings into account. The contributions of recrossings to the rate can be substantial when the order parameter along which the free energy computation is performed, is not the true reaction coordinate. This correcting factor can easily be of 3 orders of magnitude because when an order parameter is not sufficient to describe the so-called dividing surface (the hyper-surface that separates the products from the reactants) a simulation constrained at a particular order parameter value believed to be the TS, will still suffer from rare event dynamics and mostly will stay in either one of the basins. This situation is much more likely than one would think, certainly for complex systems, as there is simply no way that a single order parameter would precisely cut the highly dimensional phase space exactly along the dividing surface. However, path sampling corrects for the recrossings automatically[3,24].

The TIS algorithm provides an improvement in efficiency over the reversible integration scheme. Here the integration process is done by dividing the phase space using so-called interfaces and sampling the path ensembles for each interface. TIS defines a monotonically increasing set of interfaces $\{\lambda_0, \lambda_1, ...\lambda_n\}$ determined by an order parameter λ. The stable states A and B are bounded by the interfaces $\lambda_A = \lambda_0$ and $\lambda_B = \lambda_n$, respectively. Thus state A is defined as the set of configurations x for which $\lambda(x) < \lambda_A$, where the function $\lambda(x)$ maps x on the order parameter. For each interface TIS samples the paths that leave state A and cross the interface and return to a stable state, either A or B. The rate from a state A to B is expressed as

$$k_{AB} = \langle \phi_{A1} \rangle P_A(\lambda_B|\lambda_1) \tag{1}$$

[2] Note a common density plot, which plots the population of configurations, will often be overwhelmed by the stables states. In contrast, a path density plot counts paths only once for every 'bin' in the projection, which can highlight the barrier crossings. See for more information Ref. [26,27,62]

where ϕ_{A1} is the effective positive flux of crossing the first interface for trajectories leaving state A [13]. $P_A(\lambda_B|\lambda_A)$ is the crossing probability, the probability that a trajectory that leaves state A makes it all the way to state B. This probability is usually very low, and needs to be constructed from concatenating crossing probability histograms for each interface ensemble. This procedure also links the unbound path ensemble to the constrained one, and is hence equivalent to doing the TPS integration. An additional advantage in efficiency comes from the fact that the paths can be shortened to the absolute minimum, because a path can be halted if it enters a stable state. Particularly efficient is a combination of TIS with replica exchange methodology, which provides much better decorrelation of paths (see section 3.7). Furthermore, TIS is not very sensitive to the choice of the order parameter λ used to describe the interfaces[15]. While a variation in the order parameter might influence the efficiency of the calculation, it will not change the final outcome.

2.3 Access to the entire path space by reweighting

The path ensemble set resulting from TIS can be reweighted yielding the *reweighted path ensemble* (RPE)[28]. Attached to each path in this ensemble is the probability of observing it in a infinitely long MD trajectory. By reweighting each path observed in the TIS, one has thus access to an infinitely long trajectory (see Section 4 for a detailed discussion). Therefore all kinetic and thermodynamic projections immediately follow from the reweighted path ensemble. Besides rates and free energy, of special interest is the committor function. The committor for the final state B, $p_B(x)$, assigns to each configuration x a probability to reach the final state B rather than the initial state A when starting from that configuration with random velocities. The reweighted path ensemble allows projected committors in arbitrary collective variable spaces. This is useful for the reaction coordinate analysis.

2.4 Mechanistic insight through reaction coordinate analysis (committor analysis)

One of the main attractions of the TPS method is its ability to give unbiased mechanistic insights into the reaction mechanism. By direct inspection and statistical analysis of the pathways harvested with TPS, one can obtain information on the variables governing the course of the reaction. In this context, the committor has proven to be a particularly useful concept. Ideally such a analysis yields a reaction coordinate, a collective variable measuring the progress of the reaction. In Sec. 5 we will sketch several methods for the analysis of pathways and discuss issues arising in their application. Reaction coordinate analysis has been applied to several complex processes, for instance, ion pair dissociation [56], crystal nucleation [58] and protein conformational changes [26,27].

3 Issues and potential pitfalls in path sampling

3.1 Stable states need to be carefully defined: core sets

While the path sampling methodology does not require an order parameter that mimics the reaction coordinate, there still is a collective variable problem in the form of the stable state definitions (and the interface definition, for TIS). Fortunately, defining a stable state is in principle much easier than defining the mechanism of a

complex rare event without a priori knowledge. Using a regular simulating setup the stables states themselves are easily sampled, as in metastable states MD trajectories have a tendency to return to the same state. Analyzing these trajectories allows definition of the stable state regions, in the form a characteristic function $h_A(x) = 1$ if $x \in A$, 0 if $x \notin A$. Usually this characteristic function is defined using order parameters, e.g., $h_A(x) = 1$ if $\lambda_{A,min} < \lambda(x) < \lambda_{A,max}$. However, there are a few pitfalls that need to be avoided in such a definition of the states. The most important one is the possibility that paths are identified incorrectly as transition paths. This means that a path, identified as connecting A and B, recrosses the barrier and becomes an AA (or BB) path when extended in the forward (or backward) time direction. This false identification is caused by an incorrect or too loose definition of the characteristic functions for the stable state. As an example, suppose that in stable state A the value of the order parameter fluctuates between 0 and 0.4 and in stable state B it is between 0.6 and 1. However, because of recrossings, it is possible that a path coming from A reaches e.g. a value of 0.6 but has, in fact, still not made the transition. Such a path has not yet reached the so-called basin of attraction of B and will return quickly to A. When the states are identified by the fluctuations of the order parameters, as one would naively think would be appropriate, the TPS will allow such spurious paths, and get stuck, because once a paths lies entirely in the basin of attraction of a single state it is unlikely to escape this basin again. To avoid this problem, the stable state definitions, i.e. the characteristic functions, should be much stricter than one would naively expect. For instance, in the example, one could choose a definition of intervals [0,0.05] and [0.95,1] for A an B respectively. This sounds strange, because surely, these intervals are at the wings of the distribution of the order parameter. However, the point is that once the path is in a stable basin it quickly can explore the basins, and visit these values, on a timescale much smaller than the reaction time, while the chance to find this value in the other stable state is minimal. Hence, a possible criterion for defining the states could be the ratio of the population ρ inside the two

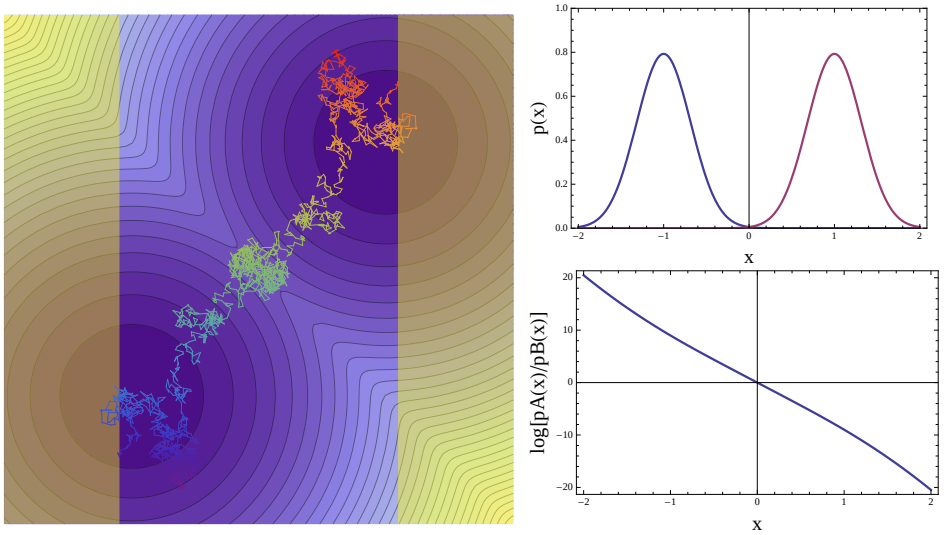

Fig. 1. Left: Cartoon of possible energy landscape with two minima, with a very strict definition in the x-axis. Top right: Equilibrium probability $\rho_A(x)$ (blue) and $\rho_B(x)$ (red) along the x-axis. Bottom right: relative logarithmic probability of ρ_A with respect to ρ_B.

basins of attractions as a function of the order parameter,

$$\log_{10} \rho_A(\lambda_{A,max})/\rho_B(\lambda_{A,max}) > 10. \quad (2)$$

Here, $\rho_A(\lambda)$ is proportional to the number of configurations x for which the committor $p_A(x) > 0.5$ and $\lambda(x) = \lambda$. So the criterion states that only one out of every 10^{10} configurations at the edge of the stable state definition of A (as characterized by $\lambda_{A,max}$) in fact belongs to B. If the order parameter is well chosen and the barrier is high enough, this is not difficult to realize (See Fig. 1 for an illustration). Of course, there is a chance that choosing only one order parameter is not sufficient to achieve this. However, a combination of order parameters can be used to make the state definition stricter. Such a combined characteristic function can be made by a (non)linear combination of order parameters $\lambda(x) = f(q_1(x), q_2(x)...)$, or a logical combination $h_A(x) = 1$ if $q_1(x) < q_{1,max} \wedge q_2(x) < q_{2,max}$.

We note that this way of defining states is related to the so-called core sets used in the MSM community[29].

3.2 Optimizing TIS simulations

As noted earlier, TIS is not as sensitive to the definition of interfaces between the stable states as other interface-based methods. The reason is that in TIS pathways are relaxed also in backward direction such that the pathways are not forced to start into the wrong configuration-space direction by unsuitably defined interfaces. The efficiency of a TIS calculation, however, can be improved considerably by a smart placement of the interfaces. In this section, we briefly discuss an adaptive procedure [30] designed to improve the interface definition and enhance the efficiency of the simulation. The basic idea of the method, based on an earlier optimization-procedure developed for FFS [31], is to place the interfaces in a way to minimize the statistical error of the calculated rate for a given total simulation time. As the accuracy of a TIS rate-calculation depends mainly on the accuracy of the probability to reach final state B once the boundary of the initial state A has been crossed, the optimization procedure concentrates on reducing the variance in the estimation of this crossing probability. As can be shown by considering the statistical properties of repeated shooting moves, the accuracy of the crossing probability is best for equal flux through all interfaces. To satisfy this constant flux condition, one repositions the interfaces using an interpolation formula obtained from a short TIS-simulation with unoptimised interfaces. This procedure places interfaces with higher density in the bottleneck regions of configuration space such that the overall numerical effort is concentrated on pathways passing through these region. Alternatively, one can also vary the number of paths sampled at each interface for fixed interface positions. Thus, in this approach one redistributes the computational effort on the interfaces without changing their position. In both cases the optimization can be carried out based on short preliminary calculations without creating a large computational overhead. TIS-calculations carried out for a simple two-dimensional model and for the dipole reorientation of ice structures inside carbon nanotubes indicate that this optimization procedure increases the efficiency of TIS-simulations of up to an order of magnitude [30] with respect to a simple uniform placement of interfaces.

3.3 TPS and TIS are less easy to implement than other path-based methods

One of the strengths of transition path sampling is the use of the backward time integration in the shooting move[9,24], which allows paths to relax and equilibrate to

correct unbiased path ensemble, and does make TIS less sensitive to the choice of the order parameter. However, the backward time integration step of the shooting move is not as easy to implement as the forward integration. It requires careful bookkeeping of the momenta at the shooting point (and also the rest of the trajectory). Therefore many path-based rare event studies have opted to use a forward shooting only method, such as the forward flux sampling (FFS) algorithm[17], which is conceptually simpler in spirit than TPS and TIS but only applies to systems with stochastic dynamics (whereas TPS/TIS can handle any microscopically reversible dynamics). While originally developed for non-equilibrium dynamics, because it does not require microscopic reversibility, the FFS scheme is also often used for studying equilibrium kinetics[18]. The basic FFS algorithm collects from a long dynamical trajectory, configurations that reach the first interface. From these points new stochastic trajectories are initiated and stopped when the paths reach either the next interface or return to the initial state (or are deemed to return). Points that reach the next interface are collected and the procedure is repeated, until paths from A to B are realized. The kinetic rate constant then follows from Eq 1. This procedure is indeed simpler than the TIS algorithm, and can be parallelized trivially, but also suffers from a (possibly strong) dependence of the initial stages of the FFS sampling. If (at least some of) the paths in the first stages of sampling are not representative of the entire transition, the final path ensemble will not be representative, and kinetic properties will be correct. We refer to a review by van Erp for a thorough analysis of this issue[15]. The usual remedy is to sample a very large number of points at each stage, so that this dependence will become minimal. But then of course the method becomes less efficient. In TIS, at each interface, the backward shooting allows relaxation of the pathways, and hence a better estimate of the kinetic mechanism and the related rate constants. As TIS requires not much more computational effort than FFS we recommend the use of TIS for studying equilibrium kinetics. Of course, when dealing with non-equilibrium dynamics one has no choice than to use only forward integration (e.g. FFS, NEUS), as the requirement for microscopic reversibility is no longer valid[17,18].

3.4 Path sampling is computationally expensive.

The TPS/TIS method is computationally expensive, more so than e.g. a biased method, such as meta-dynamics[32], but clearly not as much as straightforward MD, if the involved barrier is sufficiently high. As a rule of thumb, TPS should be more than two orders of magnitude more efficient than straightforward MD to be worthwhile. This will be already fulfilled for moderate barriers (e.g. of $5k_{\mathrm{B}}T$). If one is a priori very certain about the collective variable that describes the mechanism of the transition, e.g. a gas phase chemical reaction with a single atomic distance, the use of TPS is limited and one is better off doing enhanced configurational sampling, possibly in combination with the reactive flux approach. However, when the reaction coordinate describing the transition is far from clear, or very complex and convoluted, the TPS method might be useful to get insight in this reaction coordinate. A reaction coordinate analysis then leads to a improved description, which might immediately allow a reaction coordinate driven configurational sampling, or could be used in a TIS scheme to compute the rate constants.

3.5 How to track convergence using one way shooting

For a complex rough energy landscape, molecular dynamics trajectories connecting A with B over a high barrier can be long and tortuous. In that case, even for deterministic MD, the dynamics can be almost be viewed as stochastic, because the paths are

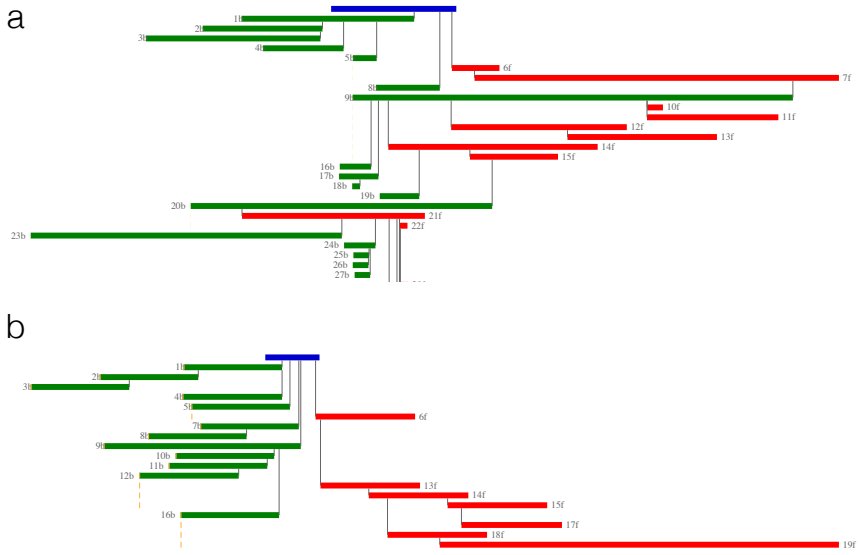

Fig. 2. Examples of sampling trees. Horizontal lines denote the trajectory from time slice 0, to the final slice L from left to right. Blue indicates the initial path. Green is a backward shot, red denotes a forward shot. Vertical lines denote the shooting point location. a) a reasonable tree, which after 20 accepted paths has completely changed the trajectories several times. b) an example of a incomplete sampling. The middle part of the trajectory has not changed at all. (Trees provided by Eva van Mastbergen.)

much longer than the Lyapunov time, which denotes the onset of molecular chaos[33, 24]. The randomly chosen shooting point on the trajectories is more often than not away from the dividing surface or the transition state. A path that is started from the shooting point will therefore likely be attracted to the same basin of the state it is in, and so does its backward shot. In that case, two-way shooting will fail more often than one-way shooting. To see this, consider that the chance that a two-way shot is accepted is roughly $p_A(x)p_B(x)$, where $p(x)$ is the committor of the configuration x, in this case the shooting point. The chance that a one-way (backward) shot is accepted is p_A. Hence, a two-way shot is accepted much less frequently than a one-way shot, certainly for points with high p_A (low $p_B = 1 - p_A$). This could be the situation for 90% of the trajectory. For such situations, the one-way shooting turns out to be more efficient[33]. However, one-way shooting does not produce a completely new path every shot, but only (roughly) half of it. The other half remains identical to the previous path. Therefore, the paths must decorrelate during the sampling. To check for decorrelation, it is convenient to draw a "shooting tree", that clearly shows when paths decorrelate: when a forward shot is accepted from a shooting point that was previously on a backward shot part of the path, or vice verse when a backward shot is accepted shot from a forward path. The alternation between forward and backward is thus crucial for the path sampling. A shooting tree will quickly show this. Figure 2 shows a clear example. Still, even when a tree looks correct, the path ensemble needs to be checked for sufficient decorrelation, e.g. using a correlation function[24] An alternative way to avoid low acceptance with two-way shooting of long paths, is to use precision shooting[34]. Here, the momenta at at the shooting point are modified so minimally, that the path retraces the old path to a substantial fraction of the trajec-

tory length. Of course, here also special attention needs to given to path decorrelation during the sampling.

3.6 Path sampling can get trapped in intermediate meta-stable states

Path sampling was originally developed for a system in which two states are separated by a single (rough) barrier. However, nothing stops systems from exhibiting intermediate states between A and B. In fact, for sufficiently complex transitions, it is almost impossible to avoid such intermediates. A transition path sampling simulation, conducted for such a process from A to B, will almost certainly visit these intermediate states. When the lifetime of an intermediate state is small, and on the same order as the duration of the direct molecular transition time from A to B this is not a problem. The average path length will become longer because of these traps, but not detrimentally so. However, problems arise when the intermediates cause the path to stay inside those states for timescales much longer than the transition time. The path length will grow exponentially, and computer power will be mostly wasted on just exploring the intermediate states, quickly rendering the sampling very inefficient. In case the number of intermediates is small, the simplest remedy, and often also the most effective, is to split the transition into multiple segments and treat each transition independently. This means that each intermediate is treated as a stable state, and defined in this way. Path sampling is then initiated in the usual way, and has to be repeated for each combination of states. For N states this boils down to $\frac{1}{2}N(N-1)$ independent transitions. This square dependence can be daunting, especially when N is running in the tens of states. Moreover, as TPS focuses on trajectories between specific pairs of states only, a small detour into any another state has to be rejected, making the sampling even more inefficient. A viable alternative to sampling all possible transitions independently is to make use of the multiple state TPS/TIS ensemble[35]. In this ensemble, all transition paths from all states to all states are allowed. In particular when all states are of similar stability, and barrier heights are not too different, this is an attractive approach. To obtain all the kinetic

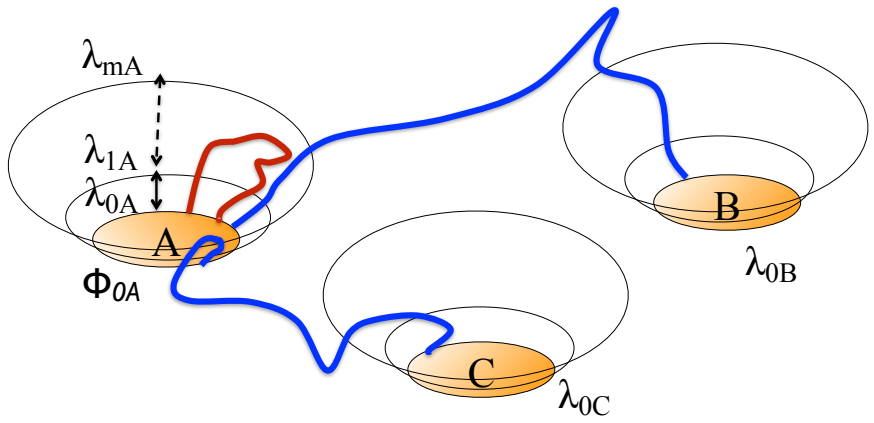

Fig. 3. Multiple state path sampling. The blue paths belong to all interface ensembles for A, including the outermost interface and multistate TPS ensembles. The red path does not reach the outermost interface and is hence only a valid path in the lower interface TIS ensembles.

rate constants, TIS computes for each state the crossing probability towards a predefined outermost interface (see Fig. 3). The final piece of the puzzle is then a multiple state TIS simulation in which all paths are allowed the connect two states, and cross the outermost interface of the initial state[35]. This reduces the computational effort from quadratic to linear scaling with the number of states N.

3.7 Multiple channels are not easily sampled: replica exchange path sampling

One of the major challenges in TPS /TIS is to sample complex processes occurring via multiple parallel routes or channels. An example could be a molecule passively transported through a membrane. This can clearly happen at multiple points in the membrane, that do not have to be identical. Another example is a chemical reaction that takes place along multiple distinct mechanisms. Finally, protein conformational changes often occur via a complex network. For such systems one would like to sample all possible parallel pathways with TPS or TIS. Since TPS is a Monte Carlo procedure, in principle, it should be able to sample all important regions. However, there might be high (free) energies barriers between channels that prevent it from doing so. This situation is similar to a regular Metropolis Monte Carlo simulation of a molecular system: e.g. a super-cooled liquid might take a long time to nucleate into a crystal. One way to overcome such barriers is to invoke a replica exchange approach [36], where multiple replicas sampled at different conditions are allowed to exchange in an extended ensemble. The mostly used 'vanilla' type replica variable is the temperature. While sampling pathways at the temperature of interest might suffer from large barriers between the distinct channels, at high temperature these barriers are easily overcome, and all channels can be sampled. Coupling the replicas then allows the replicas to flow from low to high temperatures and back again, and also correctly sample all possible channels at the temperature of interest [36].

For TIS there is also another possibility, namely the exchange between interface ensembles[20,21]. Neighboring interface ensembles can often exchange because there is a reasonable change that a path crossing interface i also crosses interface $i+1$. The exchange then allows paths to move between interfaces, from a transition path all the way to the first interface, and back again to the outermost (last) interface where the transition to B is very likely. When multiple channels exist, a pathways that starts in one channel, is able to first retract to the first interface and then extend again over the barrier toward the final state via a different channel. Decorrelation of pathways is furthermore hugely improved by the so-called minus move, in which a path in the first interface is exchanged with a new path coming from a straightforward run in the stable state. Decorrelation is further enhanced when paths from different initial states are exchanged, allowing a more efficient exploration of the path space.

An alternative way to overcome barriers in trajectory space consists in combining transition path sampling with the Wang-Landau flat histogram algorithm [37,38]. In this approach, the simulation is biased such that all values of a given variable (for instance, the energy) are sampled with uniform frequency within a certain range. The bias, which is adjusted on the fly, drives the simulation into regions that have not been visited before. Flat histogram calculations carried out in path space, with a bias applied to the energy or to the volume, have been shown to enhance the ergodicity of path sampling simulations considerably even if the transition occurs via several distinct pathways [37]. As a by-product, such simulations yield activation energies and activation volumes, which can be used to compute reaction rate constants based on thermodynamic integration starting from a reference state with known reaction rate [39].

3.8 RETIS is even more computationally expensive: Single Replica TIS

The application of replica exchange TIS (RETIS), which swaps interfaces among the replicas, solves major convergence problems in the path sampling framework and can be easily extended to the multiple state approach [40]. While solving convergence problems, RETIS has two major drawbacks: 1) it is computationally very expensive, because for a multiple state RETIS simulation one needs at least one set of interfaces per state, leading to a number of interfaces running in the hundreds[40]. 2) the method is not easily parallelized because paths in each interface ensemble likely have a different instantaneous path length, and hence each trajectory requires different run times to create. Of course one could design a clever asynchronous exchange scheme or use fancy job schedulers, but the fact remains that the effort to run such a scheme is very large to start with. Since computational effort is a major bottleneck in the study of any complex systems, it is worthwhile to try to reduce it. One option is to turn the replica exchange scheme into a single replica version similar to simulated tempering[41,42]. In this approach only one replica is being sampled at the time, and this replica moves through interface space by exchange moves[42]. The downside of the technique is that one needs to know the correct weights of the pathways at each interface a priori to achieve proper sampling. This is identical to the case of umbrella sampling, where one basically needs a good estimate for the free energy as a function of the order parameter. The single replica TIS approach does exactly this, by moving trough interface space using a biasing function, which turns out to be identical to the crossing probability. The crossing probability hence plays an analogous role to the free energy in umbrella sampling. This biasing function can be build up by a Wang-Landau algorithm which should slowly converge to the correct answer, or by an educatedly guessed fixed bias function, which can then be iteratively improved[42].

3.9 Can path sampling be merged with Markov state modeling?

For very complex systems such as protein folding or binding, even the task of defining stable state is not straightforward anymore. This is caused by the presence of a myriad of meta-stable states in such systems, each with its own intrinsic timescale. One of the most successful approaches of the last decade is the application of Markov state models to make sense of these complex systems[43,29]. In this approach one conducts many long MD simulations and analyzes these on the basis of a geometrical or kinetic criteria. That allows grouping of configurations into meta-stable sets. From the MD simulations one can then extract kinetics as well as thermodynamics by constructing kinetic transition matrices, and essentially solving the master equation numerically. While extremely successful, there seem to be two issues in MSM that would potentially be problematic: 1) there might be high barriers, related to rare but important events for the mechanism and kinetics which are not sampled in the relatively short MD trajectories. 2) It seems that one is doing too much work, because the overwhelming majority of the MD trajectories will sample metastable states but not the more relevant transition region. In this sense MSM should be seen as an analysis method rather than a sampling method. One approach is that of the adaptive seeding approach, in which new MD trajectories are restarted at places that were not encountered before [44]. However, it appears that some of the difficulties arising within the MSM approach can also be addressed by path sampling. Pande and coworkers pioneered these lines of thinking more then a decade ago[45], by using path sampling to improve MSMs. The recently developed single replica transition interface sampling (SRTIS) can provide a systematical approach. SRTIS is able to start with a single state, and when a new, previously unknown state is encountered,

signalled by a trapped pathway, this new trapping state can be analyzed and added to the allowed states. By adding states one by one to the database of states, and disfavoring paths that have been sampled over and over again, SRTIS is able to take the best of both worlds, and sample rare pathways while building up the MSM. This was effectively the approach that was taken in Ref. [46].

There are some important differences in the MSM and SRTIS approaches. While a typical MSM analysis can result in 10,000's of states, this will not be possible for SRTIS. In SRTIS the number of states should be of the order of ten, but not much more, to keep the problem tractable. However, of the 10,000 MSM states not all are stable on the order of a microsecond. On the contrary, most states are stable on the order of a nanosecond or less. This means that the MSM states are in fact not (yet) core sets in the sense of TPS stable states. However, it is perfectly possible to identify the core sets in the MSM, that are stable on timescales of tens of nanoseconds and more, and use those in the description of the stable states for path sampling. This approach is currently under investigation.

4 Data analysis of the path ensembles

4.1 Reweighting schemes

Once sampled sufficiently, the TPS and TIS path ensemble data needs to be post-processed in order to yield the rates, free energy, and the committor function for the reaction coordinate analysis. In the case of TIS, the path ensemble of the different interfaces need to be related to each other in a meaningful way. If one is just interested in a rough estimate of the kinetic rate constant, as defined in Eq. 1, computation of $P_A(\lambda_{I+1}|\lambda_i)$ as the fraction of trajectories in the path ensemble for interface i that make it to the next interface $i+1$ is sufficient[13]. However, the analysis can be made more rigorous using reweighting schemes such as the weighted histogram analysis method (WHAM) to compute the crossing probability $P_A(\lambda_B|\lambda_1)$, and hence the rate. Moreover, it turns out that the entire path ensemble can be reweighted, and projected in any desired way[28]. Since this reweighting scheme has not been reviewed before, we decided to give it a bit more attention here, as it might be of help to researchers that would like to reweight their path ensembles.

4.2 Reweighting the crossing probabilities

Provided the TIS ensembles contain a statistically sufficient number of decorrelated paths, the crossing probability $P_A(\lambda|\lambda_i)$ can be estimated by a histogram H_j^i constructed for each TIS λ_i ensemble. In practice the histogram is just a sum over the N_i sampled pathways for interface i

$$P_A(\lambda_j|\lambda_i) \approx H_j^i = \frac{1}{N_i}\sum_n^{N_i} \theta(\lambda_{max}[\mathbf{x}^{(n)}] - \lambda_j), \tag{3}$$

where every path $\mathbf{x}$ is a sequence of configurations $\mathbf{x} \equiv \{x_0, x_1, \ldots x_L\}$, the λ_{max} function returns the maximum value of λ along this path[3], and the θ function returns only unity for a non-negative argument, and zero otherwise. Note that index j refers to the histogram bins, with a specific λ_j value, whereas i refers to the interface index. In case j and i have the same resolution, i.e. the histogram has one bin per interface,

[3] This requires a monotonically increasing series of interfaces.

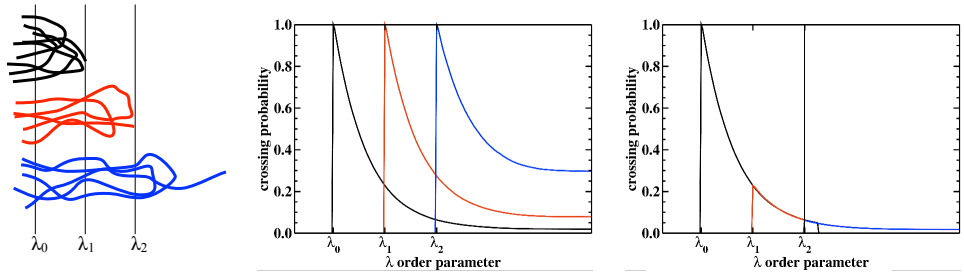

Fig. 4. Illustration of reweighting the crossing probability histograms. Left: a cartoon of the path ensembles for three interfaces. Middle: constructed normalized histograms. Right: reweighted histograms using WHAM.

the histogram H_j^j is identical to the path type number (see section 4.5). Note also that $H_j^j = 0$ for $\lambda_j < \lambda_i$, by definition. To avoid numerical issues, the histograms should be cut off, e.g. at 5% of the maximum value. The total crossing probability histogram can be constructed by e.g. the weighted histogram analysis method (WHAM), which computes the reweighted crossing probability histograms as

$$P_A(\lambda_j|\lambda_1) = \bar{w}^A_{k(\lambda_j)} \sum_{i=1}^{k(\lambda_j)} H_j^i, \tag{4}$$

where $k(\lambda) = \sum_{j=1}^{n} \theta(\lambda - \lambda_j)$ selects the correct window for the value of λ considered. The optimal WHAM weight is $\bar{w}^A_k = (\sum_{i=1}^{k} Z_A/Z_{A,i})^{-1}$, where estimates for the path ensemble partition sums $Z_{A,i}$ follow from solving the recursive relation

$$Z_{A,i} = \sum_{j}^{bins} \mathrm{sgn}(H_j^i) \frac{\sum_{k=0}^{i} H_j^k}{\sum_k^i \mathrm{sgn}(H_j^k) M_k / Z_{A,k}}, \tag{5}$$

where the sign function ensures that only positive entries of the histograms are taken into account and $M_k = \sum_j H_j^k$ is the total amount of entries in the histogram k, This equation can be solved iteratively. In the iteration the partition function of the first interface is always set to $Z_A = Z_{A,1} = 1$. Figure 4 shows an illustration of this process. An almost identical reweighting can be done using the Multistate Bennett Acceptance Ratio (MBAR) method[47].

4.3 Reweighting the paths themselves

As each TIS interface ensemble is a constrained subset of the total path ensemble, the reweighted path ensemble can construct the unbiased total path ensemble from the individual TIS ensembles, by associating each path with a probability or weight $\bar{w}^A_k$ depending on the kth interface window defined by $\lambda_k < \lambda < \lambda_{k+1}$. The weights turn out to be identical to the WHAM weights $\bar{w}^A_k$ for the crossing probability histogram, but now with $k = \sum_{i=1}^{n} \theta(\lambda_{max}[\mathbf{x}] - \lambda_i)$ the interface number that is maximally crossed.

A similar reweighting procedure can be done for state B (or any other state). The ensembles of each state should be multiplied by a factor adjusted such that the AB pathways are equally likely as BA paths. To complete the total path ensemble

one can add the paths from the stable state, e.g. from a minus interface sampling or from the stable state MD sampling.[20,28] We note that the path reweighting can in principle also be performed with the dynamical reweighting scheme of Minh and Chodera[48].

4.4 Projections of the reweighted path ensemble

Assigning the relevant weight to all paths, the reweighted path ensemble can be used to project out e.g. the free energy or committor[28,25]. For instance the probability density $\rho(\mathbf{q})$ as a function of collective variable space $\mathbf{q} = \{q_1, q_2....q_n\}$ is given by

$$\rho(\mathbf{q}) = C \sum_{n}^{N_p} W[\mathbf{x}^{(n)}] \sum_{l=0}^{L} \delta(\mathbf{q}(x_l^{(n)}) - \mathbf{q}), \tag{6}$$

where $W[\mathbf{x}]$ is the correct weight for path $\mathbf{x}$ (e.g. w_k^A or w_k^B) depending on where the path came from and how far it has traveled, N_p denotes all sampled paths in all interface ensembles, and C is a normalizing constant. The projected free energy up to a constant follows from

$$F(\mathbf{q}) = -k_B T \ln \rho(\mathbf{q}) + const, \tag{7}$$

where k_B is Boltzmann's constant.

The (projected) committor is given by

$$p_B(\mathbf{q}) = \frac{\sum_n^{N_p} W[\mathbf{x}^{(n)}] h_B(x_L^{(n)}) \sum_{l=0}^{L} \delta(\mathbf{q}(x_l^{(n)}) - \mathbf{q})}{\sum_n^{N_p} W[\mathbf{x}^{(n)}] \sum_{l=0}^{L} \delta(\mathbf{q}(x_l^{(n)}) - \mathbf{q})}, \tag{8}$$

where $h_B(\mathbf{x}) = 1$ if $x \in B$ and zero otherwise.

4.5 Path type reweighting

When applying multiple state TIS it is possible that intermediate states are overlapping with or even falling inside the outermost interface. Computing the crossing probability for that outermost interface and constructing the rate matrix might then lead to incorrect results. A possible remedy to this problem is to use *path type* reweighting[46]. For interface ensemble i the path type number $n_{IJ}^i(\lambda_{kI})$ is defined as the number of paths that start in state I, end in state J, and maximally reach interface λ_{kI}. Just as the crossing probability, the path type number can be reweighted to give the reweighted path type number $\tilde{n}_{IJ}(\lambda_{kI})$

$$\tilde{n}_{IJ}(\lambda_{kI}) = \bar{w}_k \sum_{i=1}^{m} n_{IJ}^i(\lambda_{kI}), \tag{9}$$

with the weights $\bar{w}_k$ determined from the WHAM procedure for the crossings probability (see Ref. [28]) The crossing probability $P_I(\lambda_{0J}|\lambda_{1I})$ for reaching state J (defined by λ_{0J}) from state I is now

$$P_I(\lambda_{0J}|\lambda_{1I}) = \frac{\sum_{k=1}^{m} \tilde{n}_{IJ}(\lambda_{kI})}{\sum_{J \in M} \sum_{k=1}^{m} \tilde{n}_{IJ}(\lambda_{kI})}. \tag{10}$$

Thus, the rate constant matrix is given by

$$k_{IJ} = \langle \phi_{1I} \rangle P_I(\lambda_{0J}|\lambda_{1I}), \tag{11}$$

where $\langle \phi_{1I} \rangle$ is the positive effective flux through the first interface λ_{1I}, which follows from a direct simulations or from the minus interfaces move[20].

5 Strategies for reaction coordinate analysis

5.1 The reaction coordinate and the committor

One of the central challenges in the computer simulation of complex systems is how to make sense of the vast amounts of data produced in such simulations and transform them into true understanding. For instance, molecular dynamics simulations produce long lists of numbers detailing the positions and velocities of all atoms in the system as a function of time. Now, which of these degrees of freedom need to be considered when building simplified models that capture the essence of the underlying molecular mechanisms and which can be replaced by random noise? In some cases, visual inspection of the atomistic trajectories might offer some clues on what is happening, but in many cases the relevant variables are collective and, as such, are difficult to identify based on watching a molecular movie. To overcome this problem, several machine learning algorithms have been developed recently to assist in the construction of low-dimensional models in order to rationalize the simulation results [49–51]. The issue of identifying important degrees of freedom is particularly relevant for rare event processes such as the nucleation of first order phase transitions or chemical reactions in solution, where knowledge of a valid reaction coordinate is crucial for understanding the reaction mechanism. However, this immediately begs the question, what actually is a reaction coordinate? Naturally a reaction coordinate should be a progress variable that describes the reaction. However, in our view the precise definition of such concept depends one the objective that one has in mind. In the following we will briefly discuss several possibilities.

1. The reaction coordinate is **the committor function** itself. It has been observed in the literature [52,53] that the committor function (also known as the *splitting probability* or *pfold*) is the perfect reaction coordinate as it predicts for a certain configuration the exact probability that one will reach the product state. In other words, the committor tells us how far the process has proceeded and what is likely to happen next. Moreover, the committor function (along with the equilibrium probability) permits to express all statistical properties of the reactive trajectories. While this seems a perfectly reasonable viewpoint, there are some comments to make. First, the statement is strictly true only in the case of stochastic dynamics. For deterministic dynamics a phase point has a binary value for the committor, either the trajectory goes to the product or it does not. Of course, an integration over the momenta restores the notion of the committor, but then it is not clear what the role of ballistic motion and velocity correlation is in this description. Second, and more importantly, even if the committor is a perfect description of the reaction coordinate, unfortunately it conveys not much information, and certainly not much physical insight. This is because the committor function, being a function, just yields a number. The real information in fact results from solving x in $p_B(x) = const$. Since x contains the entire system, this is a high dimensional $(3N-1)$ space itself, and not trivial (if not impossible) to analyze.
2. The next best approximation is an **optimal low dimensional model of the committor**. Indeed this seems a very good option, as such a model would give direct insight in mechanistic detail in terms of meaningful collective variables. This interpretation is thus much better than that in 1), as it gives the collective variables that are pertinent to the reaction, rather than a full high dimensional phase space. This approach allows to cast a mechanism in terms of an ordered sequence of events that occurs during the reaction.
3. When we are not really interested in what the reactive paths might do away from the dividing surface, the reaction coordinate could be approximated by a **low dimensional model of the dividing surface (separatrix)**. This is usually

the approach taken by variational TST [54] and gives insight in the transition state ensemble and direction of reactive flux. This type of reaction coordinate is also the aim of the likelihood maximization method of Peters and Trout[55], which uses the shooting point ensemble obtained with aimless shooting.

4. In many cases one is not interested in an accurate and precise reaction coordinate. A **reasonably good low dimensional model** of the transition already allows evaluation of the free energy along this reaction coordinate, rates, transmission coefficients. This definition does not really constitute a true reaction coordinate and is usually denoted *order parameter* or *collective variable*.
5. The coarsest description of a reaction coordinate is to find **reasonable order parameters or collective variables that distinguish reactants from products**. This allows a description in terms of (meta)stable states, and in e.g. allows performing TPS, but cannot be used to identify a transition state ensemble.

In our view the ultimate goal would be to analyze and model the committor function to obtain the reaction coordinate interpretation 2) in the above list, but we realize that in many case we should be happy already with achieving interpretation 3). In practice, the reaction coordinate as in 2) is available form a full analysis of the entire reweighted path ensemble[28, 57], while reaction coordinate definition 3) would be the result from analysis of a shooting point ensemble obtained with aimless shooting.

5.2 Committor analysis

While the committor carries the dynamical information required to characterize the transition, it is in general unclear how the committor can be expressed in terms of physically more transparent collective variables that permit to construct meaningful low-dimensional models from the wealth of information produced by molecular simulations. Nevertheless, the concept of the committor is the basis for several methods that have been suggested for the identification of the collective variables contributing to the reaction coordinate. Since the early days after the development of TPS it was already realized that the committor held the key for identifying whether a collective variable chosen to model the reaction coordinate was actually capable of describing the reaction correctly. Geissler et al introduced the so-called committor test or analysis [56]. To test whether the collective variable could act as a reaction coordinate, one only had to perform a simulation in which the sampling was constrained to the top of the barrier as predicted by that collective variable. Subsequently, committors could be computed for the configurations from this constrained ensemble. If these committors would have values around 0.5 the reaction coordinate would be well described by the candidate collective variable, but if the committor would instead have values very different from 0.5, e.g. only 0 or 1, the collective variable would be not be a good reaction coordinate. Of course, intermediate outcomes, such as a flat distribution of committor values, would also be possible.

The disadvantage of the above method was that for every candidate reaction coordinate this procedure had to be repeated. Therefore several groups developed analysis methods that only used the TPS data itself. The most developed one is the method of Peters and Trout [55], in which the parameters of a postulated model are adapted to maximize the likelihood to observe a particular sequence of accepted and rejected path sampling moves. This method has been generalized to include a non-linear dependence of the committor on the collective variables [55, 57–59] and has also been applied to optimize the definition of crystallinity in simulations of crystal nucleation [60]. In an alternative perspective, neural networks are used to search for combinations of physical variables that best approximate the committor [61]. All these methods, however, require the definition of a list of potential candidates

for important collective variables and, thus, rely on guesswork that is ineffective particularly for complex, heterogeneous systems. Since due to the rapid growth in raw computing power coupled with the advent of efficient sampling algorithms the amount of data generated by molecular simulations is quickly increasing, further progress in the development of machine learning approaches for the analysis of rare event simulations is sorely needed.

6 Concluding remarks

In this paper we have discussed several practical and conceptual issues in path sampling methodology: the problems that can arise with sampling the trajectories themselves, the reweighting of the path ensembles, and the strategies for reaction coordinate analysis. Naturally, we had to leave out many technical points, and moreover we realize that path sampling is not the magic solution to all rare event problems. There are also many other path based rare event techniques, which we did not mention. Nevertheless, we think that pointing out these problems and solutions will help researchers that are applying path sampling methodology to their own systems.

The path sampling framework continues to be developed, and even after 15 years there are still open issues. At the end of this paper we list a few of these. For instance, it would be convenient to have a better way to establish (de)correlation between paths. At this moment, we rely, besides the tree representations discussed in section 3 mostly on (time) correlation functions between collective variables along the paths. However, this is rather ad hoc, and does vary with a time shift. A more systematic approach for the quantification and control of correlation in path ensembles would be useful. Another possible development that was discussed in this paper would be the combination of path sampling with Markov state modeling. This would allow for systematic and automatic stable state recognition, and would be useful for very complex (bio)systems. Furthermore, better or automatic ways to optimize the path sampling would improve efficiency. This would include optimizing the used TIS order parameters as well as their placement. As pointed out in section 5 a machine learning algorithm to come up with good reaction coordinates based on committor or other path sampling data is going to be extremely helpful in understanding complex processes. Another area that could use improvement is the visualization of path ensembles. Of course trajectories can be plotted, and movies can be shown, but how to visualize an ensemble of trajectories of a complex process in space-time is not straightforward. Finally, it would be extremely useful for the community to have a path sampling software package that does all (or at least some) of the above.

Acknowledgments

This work is part of the research programme VICI 700.58.442, which is financed by the Netherlands Organization for Scientific Research (NWO). Financial support from the Austrian Science Fund FWF within the SFB Vicom (F41) and project P24681-N20 is gratefully acknowledged.

References

1. D. Frenkel and B. Smit. *Understanding Molecular Simulation, 2nd ed.*, Academic Press, San Diego, Calfifornia, 2002.

2. M. P. Allen, and D. J. Tildesley, *Computer Simulation of Liquids*, Oxford University Press, Oxford, 1987.
3. P. G. Bolhuis, D. Chandler, C. Dellago, and P. L. Geissler, Ann. Rev. Phys. Chem. **53**, 291 (2002).
4. R. Elber, A. Ghosh, A. Cárdenas, and H. Stern, Adv. Chem. Phys. **126**, 93 (2003).
5. G. Henkelman, G. Johannesson, and H. Jónsson, in Progress on Theoretical Chemistry and Physics, ed. S. D. Schwartz, Kluwer Academic Publishers (2000).
6. H. Jónsson, G. Mills, and K. W. Jacobsen, in *Classical and Quantum Dynamics in Condensed Phase Simulations*, B.J. Berne, G. Ciccotti, and D. Coker, Eds., World Scientific: Singapore, 1998.
7. W. E, W. Ren, and E. Vanden-Eijnden, Phys. Rev. B **66**, 052301 (2002).
8. L. Maragliano, A. Fischer, E. Vanden-Eijnden, and G. Ciccotti, J. Chem. Phys. **125**, 024106 (2006).
9. C. Dellago, P. G Bolhuis, F. S. Csajka, and D. Chandler, J. Chem. Phys. **108** 1964 (1998).
10. P. G. Bolhuis, C. Dellago, and D. Chandler, Faraday Discuss. **110**, 421-436 (1998).
11. C. Dellago, P. G. Bolhuis, and D. Chandler, J. Chem. Phys. **108**, 9236 (1998).
12. C. Dellago, P. G. Bolhuis, and D. Chandler, J. Chem. Phys. **110**, 6617 (1999).
13. T. S. van Erp, D. Moroni, and P. G. Bolhuis, J. Chem. Phys. **118** 7762 (2003).
14. T. S. van Erp and P. G. Bolhuis, J. Comp. Phys. **205** 157-181 (2005).
15. T.S. Van Erp, Advances in Chemical Physics, **151** 27 (2012)
16. D. Moroni, P. G. Bolhuis, and T. S. van Erp, J. Chem. Phys. **120**, 4055 (2004).
17. R. Allen, P. B. Warren, and P. R. ten Wolde, Phys. Rev. Lett. **94**, 018104 (2005).
18. R. J. Allen, D. Frenkel, and P. R. ten Wolde, J. Chem. Phys. **124**, 024102 (2006).
19. A. Dickson and A. R. Dinner, Annu. Rev. Phys. Chem. **61**, 44159 (2010)
20. T. S. van Erp, Phys. Rev. Lett. **98**, 268301 (2007).
21. P. G. Bolhuis, J. Chem. Phys. **129** 114108 (2008).
22. N. Guttenberg , A.R. Dinner, J. Weare, J. Chem. Phys. **136**, 234103 (2012)
23. C. Dellago and P.G. Bolhuis (2009) In C. Holm and K. Kremer (Eds.), Advanced computer simulation approaches for soft matter sciences III (Advances in polymer science, 221) (pp. 167-233). Berlin: Springer.
24. C. Dellago, P. G. Bolhuis, and P. L. Geissler, Adv. Chem. Phys. **123** 1 (2002).
25. P.G. Bolhuis and C. Dellago, Reviews in Comp. Chem.**27** 111210, (2010)
26. J. Juraszek and P.G. Bolhuis, Biophys. J. **95** 4246-4257 (2008).
27. J. Vreede, J. Juraszek and P.G. Bolhuis Proc. Nat. Acad. Sci. USA **107** , 2397-2402 (2010).
28. J. Rogal, W. Lechner, J. Juraszek, B. Ensing and P.G. Bolhuis J. Chem. Phys. 133 (17), 174109 (2010).
29. C. Schütte, F. Noé, J. Lu, M. Sarich, E. Vanden-Eijnden, J. Chem. Phys. **134**, 204105 (2011)
30. E. E. Borrero, M. Weinwurm and C. Dellago, J. Chem. Phys. **134**, 244118 (2011).
31. E. E. Borrero and F. A. Escobedo, J. Chem. Phys. **129**, 024115 (2008).
32. A. Laio and M. Parrinello, Proc. Natl. Acad. Sci. USA **99,** 12562 (2002).
33. P. G. Bolhuis, J. Phys. Cond. Matter **15**, S113 (2003).
34. M. Grünwald, P. L. Geissler, and C. Dellago, J. Chem. Phys. **129**, 19401 (2008).
35. J. Rogal and P. G. Bolhuis. J. Chem. Phys. 129 (22), 224107. (2008).
36. T. J. H. Vlugt and B. Smit, Phys. Chem. Comm. **2** 1 (2001).
37. E. E. Borrero and C. Dellago, J. Chem. Phys. **133**, 134112 (2010).
38. F. Wang and D. P. Landau, Phys. Rev. Lett. **86**, 2050 (2001).
39. C. Dellago and P. G. Bolhuis, Mol. Sim. **30**, 795 (2004).
40. D.W.H. Swenson and P.G. Bolhuis J. Chem. Phys., **141** 044101 (2014)
41. E. Marinari and G. Parisi, Europhys. Lett. 19, 451 (1992).
42. W. Du and P. G. Bolhuis, J. Chem. Phys. **139**, 044105 (2013).
43. J.H. Prinz, Jan-Hendrik, H. Wu, M. Sarich, B. Keller, M. Senne, M. Held, J.D. Chodera, C. Schütte, and F. Noé, J. Chem. Phys. **134** 174105 (2011)

44. X. Huang, G. R. Bowman, S. Bacalladoc and V. S. Pande, Proc. Nat. Acad. Sci. USA **106** , 19765-19769 (2009)
45. N Singhal, CD Snow, VS Pande J. Chem. Phys. **121**, 415-425 (2004)
46. W. Du and P. G. Bolhuis, J. Chem. Phys. **140**, 195102 (2014).
47. M. R. Shirts and J. D. Chodera, J. Chem. Phys. **129**, 129105 (2008)
48. D. Minh and J. D. Chodera, J. Chem. Phys. **131**, 134110 (2009).
49. A.L. Ferguson, A.Z. Panagiotopoulos, P.G. Debenedetti, I.G. Kevrekidis, Proc. Natl. Acad. Sci. USA **107**, 13597 (2010).
50. M. A. Rohrdanz, W. Zheng, M. Maggioni, and C. Clementi, J. Chem. Phys. **134**, 124116 (2011).
51. G. A. Tribello, M. Ceriotti, and M. Parrinello, Proc. Nat. Acad. Sci. USA **109**, 5196 (2012).
52. G. Hummer, J. Chem. Phys. **120**, 516 (2004).
53. W. E, W. Ren, E. Vanden-Eijnden, Chem. Phys. Lett. **413**, 242 (2005).
54. D. G. Truhlar, B. C. Garrett, and S. J. Klippenstein, J. Phys. Chem. **100**, 31 (1996).
55. B. Peters and B. L. Trout, J. Chem. Phys. **125**, 054108 (2006).
56. P. L. Geissler, C. Dellago, and D. Chandler, J. Phys. Chem. B **103**, 3706 (1999).
57. W. Lechner, J. Rogal, J. Juraszek, B. Ensing and P.G. Bolhuis, J. Chem. Phys. 133 (17), 174110 (2010).
58. W. Lechner, C. Dellago and P. G. Bolhuis, Phys. Rev. Lett. **106**, 085701 (2011).
59. C. Leitold and C. Dellago, J. Chem. Phys. **141**, 134901 (2014).
60. S. Jungblut, A. Singraber, and C. Dellago, Mol. Phys. **111**, 3527 (2013).
61. A. Ma and A. R. Dinner, J. Phys. Chem. B **109**, 6769 (2005).
62. P.G. Bolhuis and W. Lechner J. Stat. Phys. **145** , 841-859 (2011)