

TRANSITION PATH SAMPLING

CHRISTOPH DELLAGO

*Department of Chemistry, University of Rochester,
Rochester, NY 14627*

PETER G. BOLHUIS

*Department of Chemical Engineering, University of Amsterdam,
1018 WV Amsterdam, The Netherlands*

PHILLIP L. GEISLER

*Department of Chemistry, Massachusetts Institute of Technology,
Cambridge, MA 02139*

CONTENTS

- I. Introduction
- II. Defining the Transition Path Ensemble
 - A. Dynamical Path Probability
 - B. Reactive Path Probability
 - C. Deterministic Dynamics
 - D. Stochastic Dynamics
 - E. Monte Carlo Dynamics
 - F. Defining the Initial and Final Region
- III. Sampling the Transition Path Ensemble
 - A. Shooting Moves
 - 1. Shooting Moves for Deterministic Dynamics
 - 2. Selecting Phase-Space Displacements
 - 3. Momentum Rescaling
 - 4. Efficiency of Deterministic Shooting Moves
 - 5. Shooting Moves for Stochastic Dynamics
 - B. Shifting Moves
 - 1. Shifting Moves for Deterministic Dynamics
 - 2. Shifting Moves for Stochastic Dynamics

Advances in Chemical Physics, Volume 123, Edited by I. Prigogine and Stuart A. Rice.
ISBN 0-471-21453-1 © 2002 John Wiley & Sons, Inc.

- C. Memory Requirements
 - D. Stochastic Trajectories as Sequences of Random Numbers
 - E. Other Algorithms
 - 1. Local Algorithm for Stochastic Pathways
 - 2. Dynamical Algorithm
 - 3. Configurational Bias Monte Carlo
 - F. Parallel Tempering
 - G. Generating an Initial Path
 - H. Transition Path Sampling with an Existing Molecular Dynamics Program
 - IV. Computing Rates and Reversible Work
 - A. Population Fluctuations
 - B. Reversible Work
 - C. Umbrella Sampling
 - D. A Convenient Factorization
 - E. Correspondence with Reactive Flux Theory
 - F. How Long Should Pathways Be?
 - V. Analyzing Transition Pathways
 - A. Reaction Coordinates and Order Parameters
 - B. The Separatrix and the Transition State Ensemble
 - C. Computing Committors
 - D. Committor Distributions
 - E. Path Quenching
 - VI. Outlook
- Acknowledgments
References

1. INTRODUCTION

In this chapter, we review in detail the theory and methodology of transition path sampling. This computational technique is an importance sampling of reactive trajectories, the rare but important dynamical pathways that bridge stable states. We discuss the statistical view of dynamics underlying the method. Within this perspective, ensembles of trajectories can be sampled and manipulated in close analogy to standard techniques of statistical mechanics. Because transition path sampling does not require foreknowledge of reaction mechanisms, it is a natural tool for studying complex dynamical structures of high-dimensional systems encountered at the frontiers of physics, chemistry, and biology.

The dynamics of many such systems involve rare but important transitions between long-lived stable states. These stable states could be, for example, distinct inherent structures of a supercooled liquid, reactants and products of a chemical reaction, or native and denatured states of a protein. In each case, the system spends the bulk of its time fluctuating within stable states, so that transitions occur only rarely. In order to understand such processes in detail, it is necessary to distinguish reaction coordinates, whose

fluctuations drive transitions between stable states, from orthogonal variables, whose fluctuations may be viewed as random noise. In principle, computer simulations can provide such insight. Because the times separating successive transitions are long, however, conventional simulations most often fail to exhibit the important dynamics of interest.

A straightforward approach to such problems is to follow the time evolution of the system with molecular dynamics simulations until a reasonable number of events has been observed. The computational requirements of such a procedure are, however, impractically excessive for most interesting systems. For example, a specific water molecule in liquid water has a lifetime of about 10 h, before it dissociates to form hydronium and hydroxide ions. Thus, only a few ionization events occur every hour in a system of, say, 100 water molecules. Since the simulation of molecular motions proceeds in time steps of roughly 1 fs, approximately 10^{18} steps would be required to observe just one such event. Such a calculation is beyond the capabilities of the fastest computers available today and in the foreseeable future.

A different strategy, often used to study chemical reactions, is to search for the dynamical bottlenecks the system passes through during a transition. For a simple system, with an energy landscape as depicted in Figure 1.1(a), this can often be accomplished by enumerating stationary points on the potential energy surface [1,2]. Neglecting the effects of entropy, local minima exemplify stable (or metastable) states. Saddle points exemplify transition states, activated states from which the system may access different

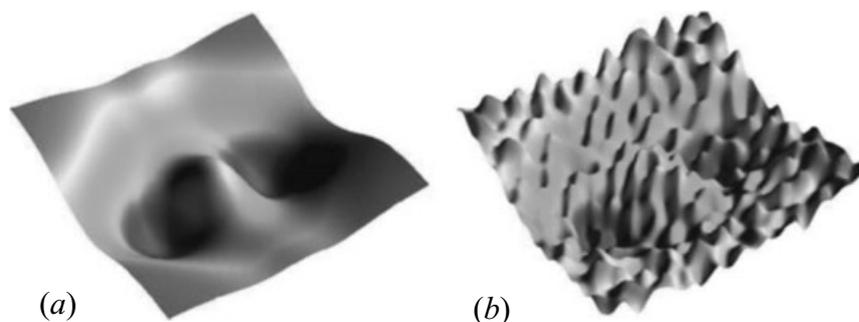


Figure 1.1. Prototypical potential energy surface of a simple system (a) and of a complex system (b). In a simple, low-dimensional system, dynamical bottlenecks for transitions between long-lived stable states most often coincide with saddle points on the potential energy surface. Locating these stationary points reveals the reaction mechanism. In a typical complex system, the potential energy surface is rugged and has countless local minima and saddle points. Nevertheless, there can be well-defined long-lived stable states and rare transitions between them. Such transitions can occur via a multitude of different transition pathways.

stable states via small fluctuations. One can often infer the mechanism of a reaction by comparing stable and transition states. Transition rates can subsequently be calculated by computing the reversible work to reach the transition state from a stable state, and then initiating many fleeting trajectories from the transition state [3].

The situation is dramatically different for complex systems, classified by Kadanoff [4] as having “many chaotically varying degrees of freedom interacting with one another.” Figure 1.1(b) shows how one might envision the energy landscape of such a system. As in the simple system, long-lived stable states are separated by an energetic barrier. But the stationary points exemplifying this barrier comprise only a small fraction of the total set of saddle points, as is generally the case for complex systems. An incomplete enumeration of stationary points is thus insufficient to locate transition states of interest. One might hope instead to guide the search for transition states using physical intuition, in effect postulating the nature of reaction coordinates. But these variables can be highly collective, and therefore difficult to anticipate. In the case of electron transfer, for example, the relevant coordinate is an energy gap that depends on many atomic coordinates. A specific value of the energy gap can be realized in many different ways. Similarly, reaction coordinates for protein folding are expected to depend on many protein and solvent degrees of freedom.

In order to overcome these problems inherent to the study of rare events in complex systems, we have developed a computer simulation technique based on a statistical mechanics of trajectories [5]. In formulating this method, we have recognized that transitions in complex systems may be poorly characterized by a single sequence of configurations, such as a minimum energy pathway. Indeed, a large set of markedly different pathways may be relevant. We term the properly weighted set of reactive trajectories the *transition path ensemble*. Defining this ensemble does not require prior knowledge of a reaction coordinate. Rather, it is sufficient to specify the reactants and products of a transition. This is a crucial feature of this method, since knowledge of a reaction coordinate is usually unavailable for complex systems.

To sample the transition path ensemble, we have developed efficient Monte Carlo procedures [6,7] that generate random walks in the space of trajectories. As a result of this “transition path sampling”, one obtains a set of reactive trajectories, from which the reaction mechanism (or mechanisms) can be inferred. Since trajectories generated in the transition path sampling method are true dynamical trajectories, free of any bias, the ensemble of harvested paths can also be used to calculate reaction rates. The high efficiency of these algorithms has significantly widened the range of processes amenable to computer simulation. As was demonstrated in applications of the methodology [5–26] the spectrum of tractable problems now

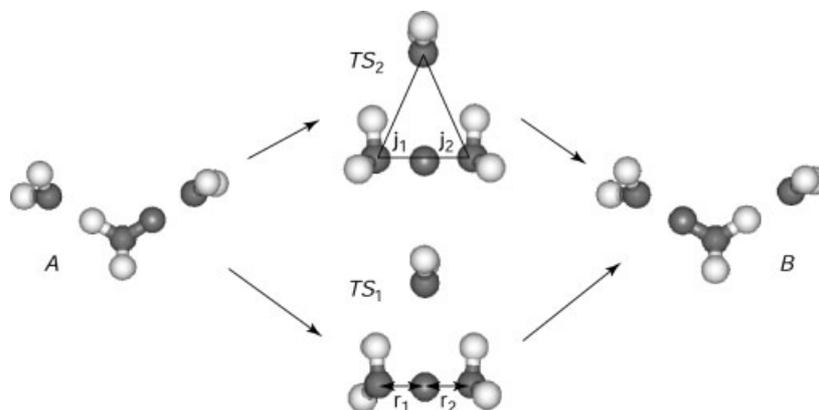


Figure 1.2. The protonated water trimer consists of three water molecules and one excess proton. In the equilibrium configuration of this cluster, the excess proton (dark gray) is bound to a central water molecule forming a well-defined hydronium ion (configuration *A*) [14]. Proton transfer from this central water molecule to one of the other molecules can occur through two different transition state regions denoted by TS_1 and TS_2 in the figure [10]. In the final state (configuration *B*), a different water molecule holds the excess proton. The transfer of the proton, shown in dark gray, requires rearrangement of the cluster's hydrogen-bonding structure. The angles φ_1 and φ_2 indicated in the upper transition state configuration, TS_2 , are used to define an order parameter characterizing the stable states *A* and *B*. The distance r_1 and r_2 indicated in the lower transition state configuration are the distances between the transferring proton and the donating and accepting water molecule, respectively.

includes chemical reactions in solution, conformational transitions in biopolymers, and transport phenomena in condensed matter systems.

In this chapter, we present the foundations and methodology of transition path sampling comprehensively, including details important for its implementation. Readers interested in a broad overview of the perspective exploited by the method, and several of its applications, are encouraged to consult a recent review [18]. In the following sections, we first discuss the theoretical basis of transition path sampling, namely, a statistical mechanics of trajectories. We then describe how reactive trajectories may be efficiently sampled, and subsequently analyzed. The practical simplicity of the method is emphasized by outlining essential algorithms in separate schemes. Computer code exemplifying the application of these algorithms can be downloaded from the website <http://www.pathsampling.org>.

Applications of the method are not discussed separately in this chapter, but are instead used to exemplify important aspects of the method. In particular, proton transfer in the protonated water trimer as shown in Figure 1.2 serves as an illustration for many of the techniques discussed in this chapter. Details of this process are discussed in the caption of Figure 1.2.

II. DEFINING THE TRANSITION PATH ENSEMBLE

In an ergodic system, every possible trajectory of a particular duration occurs with a unique probability. This fact may be used to define a distribution functional for dynamical paths, upon which the statistical mechanics of trajectories is based. For example, with this functional one can construct partition functions for ensembles of trajectories satisfying specific constraints, and compute the reversible work to convert between these ensembles. In later sections, we will show that such manipulations may be used to compute transition rate constants. In this section, we derive the appropriate path distribution functionals for several types of microscopic dynamics, focusing on the constraint that paths are reactive, that is, that they begin in a particular stable state, A , and end in a different stable state, B .

A. Dynamical Path Probability

Let us denote a trajectory of length \mathcal{T} by $x(\mathcal{T})$. While in principle the time evolution of the system is continuous, it is convenient to discretize time and represent a trajectory by an ordered sequence of states, $x(\mathcal{T}) \equiv \{x_0, x_{\Delta t}, x_{2\Delta t}, \dots, x_{\mathcal{T}}\}$. Consecutive states, or *time slices*, are separated by a small time increment, Δt . Accordingly, such a representation consists of $L = \mathcal{T}/\Delta t + 1$ states. Each of the states x along the trajectory contains a complete set of variables describing the system. For a system that evolves according to Newtonian dynamics, for example, the state $x \equiv \{r, p\}$ consists of the coordinates, r , and momenta, p , of all particles. For a system that evolves according to the rules of Brownian dynamics, x denotes only the configuration of the system. A trajectory connecting stable states A and B is schematically depicted in Figure 1.3.

The statistical weight, $\mathcal{P}[x(\mathcal{T})]$, of a particular trajectory $x(\mathcal{T})$, depends on the distribution of initial conditions and on the specific propagation rules describing the time evolution of the system. Consider, for example, a Markovian process in which state x_t evolves into state $x_{t+\Delta t}$ over a time Δt with probability $p(x_t \rightarrow x_{t+\Delta t})$. In this case, the dynamical path probability can be expressed as a product of short-time transition probabilities,

$$\mathcal{P}[x(\mathcal{T})] = \rho(x_0) \prod_{i=0}^{\mathcal{T}/\Delta t - 1} p(x_{i\Delta t} \rightarrow x_{(i+1)\Delta t}) \quad (1.1)$$

Here, $\rho(x_0)$ denotes the distribution of states x_0 serving as starting points for trajectories $x(\mathcal{T})$. For example, these initial conditions might be distributed according to the canonical ensemble, $\rho(x_0) \propto \exp\{-\beta H(x_0)\}$, where

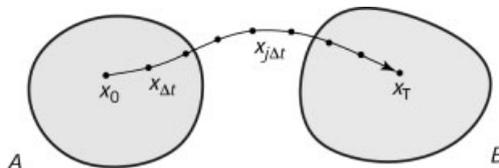


Figure 1.3. A transition pathway $x(\mathcal{T}) \equiv \{x_0, x_{\Delta t}, \dots, x_{j\Delta t}, \dots, x_{\mathcal{T}}\}$ connecting stable regions A and B .

$H(x)$ is the Hamiltonian of the system. Both $\rho(x_0)$ and $p(x_{i\Delta t} \rightarrow x_{(i+1)\Delta t})$ are assumed to be normalized.

B. Reactive Path Probability

Since we are interested only in reactive trajectories connecting A and B , we now restrict the path ensemble to trajectories beginning in region A at time zero and ending in region B at time \mathcal{T}

$$\mathcal{P}_{AB}[x(\mathcal{T})] \equiv h_A(x_0)\mathcal{P}[x(\mathcal{T})]h_B(x_{\mathcal{T}})/Z_{AB}(\mathcal{T}) \quad (1.2)$$

Here, $h_A(x)$ and $h_B(x)$ are the population functions, or characteristic functions, of regions A and B , respectively. The function $h_A(x) = 1$ if its argument, x , lies in A , and vanishes otherwise. The characteristic function $h_B(x)$ is similarly defined. In most cases, $h_A(x)$ and $h_B(x)$ depend only on the configuration part r of state x , but situations may arise in which it is advantageous to introduce characteristic functions that depend on momenta as well. In practical applications, the stable regions A and B must be characterized carefully. We will return to this issue in Section II.F.

In Eq. (1.2), $Z_{AB}(\mathcal{T})$ is a factor normalizing the distribution of trajectories,

$$Z_{AB}(\mathcal{T}) \equiv \int \mathcal{D}x(\mathcal{T}) h_A(x_0)\mathcal{P}[x(\mathcal{T})]h_B(x_{\mathcal{T}}) \quad (1.3)$$

The fact that $Z_{AB}(\mathcal{T})$ has the form of a partition function is important for the calculation of rate constants, as we will discuss later. The notation $\int \mathcal{D}x(\mathcal{T})$, borrowed from the theory of path integrals, indicates a summation over all pathways $x(\mathcal{T})$. For a discretized path, this summation corresponds to an integration over states at each time slice of the path.

The probability functional from Eq. (1.2) is a statistical description of all pathways of length \mathcal{T} connecting reactants with products. We call this set

of appropriately weighted paths the *transition path ensemble*. Pathways $x(\mathcal{T})$, which do not begin in A or do not end in B , have zero probability in this ensemble. Reactive trajectories, on the other hand, may have a non-vanishing probability, depending on the dynamical path weight $\mathcal{P}[x(\mathcal{T})]$.

The perspective exploited by transition path sampling, namely, a statistical description of pathways with endpoints located in certain phase-space regions, was first introduced by Pratt [27], who described stochastic pathways as chains of states, linked by appropriate transition probabilities. Others have explored similar ideas and have constructed ensembles of pathways using ad hoc probability functionals [28–35]. Pathways found by these methods are reactive, but they are not consistent with the true dynamics of the system, so that their utility for studying transition dynamics is limited. Trajectories in the transition path ensemble from Eq. (1.2), on the other hand, are true dynamical trajectories, free of any bias by unphysical forces or constraints. Indeed, transition path sampling *selects* reactive trajectories from the set of all trajectories produced by the system’s intrinsic dynamics, rather than *generating* them according to an artificial bias. This important feature of the method allows the calculation of dynamical properties such as rate constants.

C. Deterministic Dynamics

Consider a system whose time evolution is described by a set of ordinary differential equations,

$$\dot{x} = \Gamma(x) \quad (1.4)$$

where \dot{x} indicates the time derivative of x and $\Gamma(x)$ is a function of x only. The time evolution of such a dynamical system is deterministic in the sense that initial conditions x_0 completely determine the trajectory for all times. Newton’s equations of motion, for example, have this form.

$$\dot{r} = \frac{\partial H(r, p)}{\partial p}, \quad \dot{p} = -\frac{\partial H(r, p)}{\partial r} \quad (1.5)$$

Other examples for this type of dynamics include the equations of motion based on the extended Lagrangian of Car and Parrinello [36] and equations of motion for various thermostatted systems [37]. The equations of motion for hydrodynamic flow can also be cast in this form [38].

Solving the equations of motion (1.4) yields the propagator ϕ_t , which maps the initial state of the system to that at time t :

$$x_t = \phi_t(x_0) \quad (1.6)$$

Because this mapping takes the state x_t into exactly one state $x_{t+\Delta t}$ at time Δt later, the short time transition probability is represented by a Dirac delta function:

$$p(x_t \rightarrow x_{t+\Delta t}) = \delta[x_{t+\Delta t} - \phi_{\Delta t}(x_t)] \quad (1.7)$$

Note that here the argument of the delta function is a high-dimensional vector. Accordingly, the delta function of Eq. (1.7) is actually a product of delta functions, one for each coordinate. The reactive path probability for a deterministic trajectory is therefore

$$\mathcal{P}_{AB}[x(\mathcal{T})] = \rho(x_0)h_A(x_0) \left\{ \prod_{i=0}^{\mathcal{T}/\Delta t - 1} \delta[x_{(i+1)\Delta t} - \phi_{\Delta t}(x_{i\Delta t})] \right\} h_B(x_{\mathcal{T}})/Z_{AB}(\mathcal{T}) \quad (1.8)$$

The normalization factor $Z_{AB}(\mathcal{T})$ is given by

$$Z_{AB}(\mathcal{T}) = \int dx_0 \rho(x_0)h_A(x_0)h_B(x_{\mathcal{T}}) \quad (1.9)$$

where integrations over the states along the path have been carried out at all times except time zero.

D. Stochastic Dynamics

Often, the analysis of molecular systems can be simplified by replacing certain degrees of freedom by random noise. With this replacement, the remaining degrees of freedom x evolve according to a Langevin equation. In the simplest case, the random noise is uncorrelated in time, giving

$$\begin{aligned} \dot{r} &= \frac{p}{m} \\ \dot{p} &= \mathcal{F}(r) - \gamma p + \mathcal{R} \end{aligned} \quad (1.10)$$

Here, $\mathcal{F}(r)$ is the force derived from the potential energy $V(r)$. The friction constant γ and the random force \mathcal{R} are related through the fluctuation dissipation theorem: $\langle \mathcal{R}(t)\mathcal{R}(0) \rangle = 2m\gamma k_B T \delta(t)$, where T is the temperature and k_B is Boltzmann's constant. The random thermal noise \mathcal{R} compensates for the energy dissipated by the frictional term $-\gamma p$. Because we focus on finite segments of trajectories, the treatment of noise that is correlated in time is awkward within the specific methodology presented in this chapter.

But even in this non-Markovian case, every finite trajectory of the primary variables has a well-defined probability, and transition path sampling can in principle be carried out with sufficient generalizations.

Various integration algorithms have been derived to solve the equation of motion (1.10) over small time increments Δt [39]. Applying these operations repeatedly yields stochastic trajectories of arbitrary length. Typically, these integration algorithms have the form

$$x_{t+\delta t} = x_t + \delta x_S + \delta x_R \quad (1.11)$$

While the systematic part δx_S is fully determined by x_t , the random part δx_R is drawn from a distribution $w(\delta x_R)$. For Langevin dynamics, $w(\delta x_R)$ is a multivariate Gaussian distribution as derived by Chandrasekhar [40]. The random component of the small time step propagator “smears out” the time evolution of the system such that many different states are accessible starting from the same initial state. The single-step transition probability is given by

$$p(x_t \rightarrow x_{t+\Delta t}) = w(\delta x_R) \left| \frac{\partial \delta x_R}{\partial x_{t+\Delta t}} \right| \quad (1.12)$$

where the Jacobian on the right-hand side arises from the variable transformation from random displacement δx_R to phase space point $x_{t+\Delta t}$. For the most widely used integration algorithm [39] this Jacobian is unity, simplifying the path weight considerably [6]. Furthermore, the transition probability derived from the Langevin equation is normalized, and it conserves the canonical distribution as required by the fluctuation–dissipation theorem [6]. By concatenating transition probabilities of Eq. (1.12) as in Eq. (1.1) and imposing boundary conditions as in Eq. (1.2), we obtain the reactive path probability

$$\mathcal{P}_{AB}[x(\mathcal{T})] = \rho(x_0) h_A(x_0) \left\{ \prod_{i=0}^{\mathcal{T}/\Delta t - 1} w[x_{(i+1)\Delta t} - x_{i\Delta t} - \delta x_S] \right\} h_B(x_{\mathcal{T}}) / Z_{AB}(\mathcal{T}) \quad (1.13)$$

In contrast to deterministic pathways, stochastic trajectories are not completely determined by the initial state x_0 . Accordingly, the probability functional $\mathcal{P}_{AB}[x(\mathcal{T})]$ explicitly depends on the entire path $x(\mathcal{T})$.

E. Monte Carlo Dynamics

The dynamics of some complex systems, such as spin systems, lattice gases, and certain models of proteins, are most naturally studied using Monte

Carlo simulations [8, 41–43]. Such simulations proceed stepwise in a biased random walk through the space of possible states. In this walk, a trial state is generated from the current state of the system, and is accepted with a probability that depends on the relative weights of the trial and original state in the ensemble of interest. If the trial state is not accepted, then the original state is retained. The acceptance probability is constructed so that every state x is visited with a frequency proportional to its equilibrium weight $\rho(x)$ [44]. This procedure is called importance sampling, because the most important states, those with the largest weight, are visited most often. States with negligible weight are rarely observed.

Various acceptance rules have been devised for importance sampling, the simplest of which is the so-called Metropolis algorithm [45]. In this case,

$$p(x_t \rightarrow x_{t+\Delta t}) = \omega(x_t \rightarrow x_{t+\Delta t}) + \delta(x_t - x_{t+\Delta t})Q(x_t) \quad (1.14)$$

where

$$\omega(x_t \rightarrow x_{t+\Delta t}) = \eta(x_t \rightarrow x_{t+\Delta t}) \min \left[1, \frac{\rho(x_{t+\Delta t})}{\rho(x_t)} \right] \quad (1.15)$$

is the probability for an accepted trial move from x_t to $x_{t+\Delta t}$ and

$$Q(x) = 1 - \int dx' \omega(x \rightarrow x') \quad (1.16)$$

is the total rejection probability for a move starting from x . Trial states are generated by selecting a random displacement from the distribution $\eta(x_t \rightarrow x_{t+\Delta t})$, which is assumed to be symmetric, that is, $\eta(x \rightarrow x') = \eta(x' \rightarrow x)$. In simple Monte Carlo simulations $\eta(x \rightarrow x')$ is a rapidly decaying function of the magnitude of the displacement, but more complicated choices are sometimes useful [44]. The min-function in Eq. (1.15) returns the smaller of its arguments. Thus, a trial state is always accepted if it has a larger weight than the original state, but is rejected with a finite probability if its weight is lower. The transition probability from Eq. (1.14) is normalized and conserves the equilibrium distribution $\rho(x)$.

By concatenating Monte Carlo transition probabilities according to Eq. (1.1), one obtains the probability of a particular stochastic path $x(\mathcal{T})$ generated in a Metropolis Monte Carlo simulation. The time variable t describing the progress of this stochastic process is artificial. This *Monte Carlo time* can be approximately mapped to a physical timescale by comparing known dynamical properties such as transport coefficients [8,41].

F. Defining the Initial and Final Region

Transition path sampling does not require knowledge of a *reaction coordinate* describing the progress of a transition through the dynamical bottleneck region. It is only necessary that the initial and final states of the transition are characterized carefully. While this requirement is considerably less stringent, its satisfaction is crucial. Typically, regions A and B are defined by distinct ranges of a low dimensional *order parameter* q . Identifying an order parameter that truly discriminates between A and B can be quite challenging.

A successful order parameter must satisfy several criteria. First, regions A and B must be large enough to accommodate typical equilibrium fluctuations in their corresponding basins of attraction. The basin of attraction of a specific stable state consists of all configuration from which trajectories nearly always relax into that stable state. If this first criterion is not met, important transition pathways will likely be overlooked. Second, and more importantly, the regions spanned by $h_A(x)$ and $h_B(x)$ should be located entirely within the corresponding basins of attraction. In other words, region A should not overlap with the basin of attraction of B , and vice versa. If this second criterion is not met, transition path sampling may harvest nonreactive trajectories. This situation is illustrated in Figure 1.4, which shows two pathways on a free energy surface $w(q, q')$.

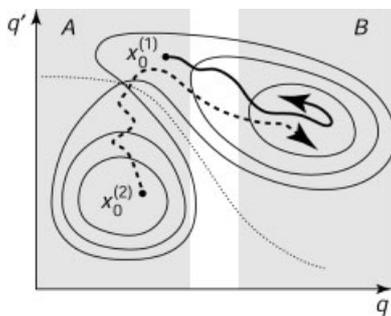


Figure 1.4. Contour lines of a free energy landscape $w(q, q')$ in which the coordinate q fails to unambiguously separate the basins of attraction of the stable states. Regions A and B (gray) accommodate most equilibrium fluctuations in the respective stable states, but the overlap of A with the basin of attraction of the final state leads to pathways that are not truly reactive. Although the trajectory depicted as a solid line with initial state $x_0^{(1)}$ starts in A and ends in B it does not cross the transition state surface (dotted line) separating the basins of attraction. The trajectory depicted as a dashed line with initial condition $x_0^{(2)}$, on the other hand, starts in the basin of attraction of A and ends in the basin of attraction of B . With the definition of the initial region A and the final region B depicted in the figure the transition path sampling algorithm is most likely to sample nonreactive trajectories rather than reactive ones.

pathways begin in A and end in B , only one of them actually crosses the transition state surface (dashed line). Although the other path (solid line) also begins with $h_A = 1$, its initial point $x_0^{(1)}$ is in fact located in the basin of attraction of state B . Thus, this pathway does not exhibit a true transition.

Suitable definitions of regions A and B may require considerable trial and error. Fortunately, it is straightforward to diagnose an unsuccessful order parameter. For example, most short trajectories initiated from the state $x_0^{(1)}$ will quickly visit states with values of q characteristic of state B . In other words, the probability to relax into B is close to one. (This relaxation probability plays an important role in the analysis of transition pathways, as will be discussed in detail in Section V.) In contrast, the probability to relax into B from $x_0^{(2)}$ is negligible. When relaxation probabilities indicate that definitions of A and B do not exclude nonreactive trajectories, the nature and/or ranges of the order parameter must be refined.

III. SAMPLING THE TRANSITION PATH ENSEMBLE

Transition path sampling is an importance sampling of trajectories, akin to the importance sampling of configurations described in Section II.E. Specifically, it is a biased random walk in the space of trajectories, in which each pathway is visited in proportion to its weight in the transition path ensemble. Because trajectories that do not exhibit the transition of interest have zero weight in this ensemble, they are never visited. In this way, attention is focused entirely on the rare but important trajectories, those that are reactive.

We accomplish the random walk through trajectory space as follows: Beginning with a trajectory $x^{(o)}(\mathcal{T})$ [here, the superscript (o) stands for old] whose weight $\mathcal{P}_{AB}[x^{(o)}(\mathcal{T})]$ in the transition path ensemble is nonzero, we generate a new trajectory $x^{(n)}(\mathcal{T})$. Using the terminology of conventional Monte Carlo techniques we call this procedure to create a new trajectory a “trial move”. Efficient methods for such generation of new trajectories by modifying existing ones will be discussed in the following sections. Next, we accept the newly generated path with a certain probability. There are many ways to construct an appropriate acceptance probability. The simplest is based on *detailed balance* of moves in trajectory space. This criterion requires that the frequency of accepted moves from $x^{(o)}(\mathcal{T})$ to $x^{(n)}(\mathcal{T})$ is exactly balanced by the frequency of reverse moves:

$$\mathcal{P}_{AB}[x^{(o)}(\mathcal{T})]\pi[x^{(o)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})] = \mathcal{P}_{AB}[x^{(n)}(\mathcal{T})]\pi[x^{(n)}(\mathcal{T}) \rightarrow x^{(o)}(\mathcal{T})] \quad (1.17)$$

where, $\pi[x(\mathcal{T}) \rightarrow x'(\mathcal{T})]$ is the conditional probability to make a move from $x(\mathcal{T})$ to $x'(\mathcal{T})$ given an initial path $x(\mathcal{T})$. In our case, $\pi[x(\mathcal{T}) \rightarrow x'(\mathcal{T})]$ is a product of the probability to generate $x'(\mathcal{T})$ from $x(\mathcal{T})$, $P_{\text{gen}}[x(\mathcal{T}) \rightarrow x'(\mathcal{T})]$, and the probability $P_{\text{acc}}[x(\mathcal{T}) \rightarrow x'(\mathcal{T})]$ to accept the trial path $x'(\mathcal{T})$:

$$\pi[x(\mathcal{T}) \rightarrow x'(\mathcal{T})] = P_{\text{gen}}[x(\mathcal{T}) \rightarrow x'(\mathcal{T})] \times P_{\text{acc}}[x(\mathcal{T}) \rightarrow x'(\mathcal{T})] \quad (1.18)$$

From the detailed balance condition one obtains a condition for the acceptance probability:

$$\frac{P_{\text{acc}}[x^{(0)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})]}{P_{\text{acc}}[x^{(n)}(\mathcal{T}) \rightarrow x^{(0)}(\mathcal{T})]} = \frac{\mathcal{P}_{AB}[x^{(n)}(\mathcal{T})]P_{\text{gen}}[x^{(n)}(\mathcal{T}) \rightarrow x^{(0)}(\mathcal{T})]}{\mathcal{P}_{AB}[x^{(0)}(\mathcal{T})]P_{\text{gen}}[x^{(0)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})]} \quad (1.19)$$

This condition can be satisfied conveniently using the Metropolis rule [45]:

$$P_{\text{acc}}[x^{(0)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})] = \min \left\{ 1, \frac{\mathcal{P}_{AB}[x^{(n)}(\mathcal{T})]P_{\text{gen}}[x^{(n)}(\mathcal{T}) \rightarrow x^{(0)}(\mathcal{T})]}{\mathcal{P}_{AB}[x^{(0)}(\mathcal{T})]P_{\text{gen}}[x^{(0)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})]} \right\} \quad (1.20)$$

Since the old trajectory $x^{(0)}(\mathcal{T})$ is reactive, that is, $h_A[x_0^{(0)}] = 1$ and $h_B[x_{\mathcal{T}}^{(0)}] = 1$, the acceptance probability can be written as

$$P_{\text{acc}}[x^{(0)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})] = h_A[x_0^{(n)}]h_B[x_{\mathcal{T}}^{(n)}] \times \min \left\{ 1, \frac{\mathcal{P}[x^{(n)}(\mathcal{T})]P_{\text{gen}}[x^{(n)}(\mathcal{T}) \rightarrow x^{(0)}(\mathcal{T})]}{\mathcal{P}[x^{(0)}(\mathcal{T})]P_{\text{gen}}[x^{(0)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})]} \right\} \quad (1.21)$$

where $x_0^{(n)}$ and $x_{\mathcal{T}}^{(n)}$ are the first and last state of the new trajectory, respectively. An acceptance probability of the form $\min[1, \alpha]$ may be realized by accepting a new pathway unconditionally when $\alpha > 1$, and with probability α when $\alpha < 1$. A simple way to implement this rule is to accept a trial move whenever a random number ζ drawn from a uniform distribution in the interval $[0, 1]$ is smaller than α . According to Eq. (1.21), only new pathways connecting regions A and B have a nonzero acceptance probability. If a new pathway does not begin in A or end in B it is rejected. A summary of the algorithm is given in Scheme 1.1.

An important feature of the acceptance probability in Eq. (1.21) is that a new pathway with lower statistical weight than the old one is accepted with finite probability. As a result, ‘‘barriers’’ in path space can be surmounted, facilitating ‘‘relaxation’’ toward the most important regions in path space. Thus, it is not essential that the first reactive trajectory have high statistical weight.

1. Generate a new pathway $x^{(n)}(\mathcal{T})$ from the current one, $x^{(0)}(\mathcal{T})$, with generation probability $P_{\text{gen}}[x^{(0)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})]$.
2. Accept or reject the new pathway according to a Metropolis acceptance criterion obeying detailed balance with respect to the transition path ensemble $\mathcal{P}_{AB}[x(\mathcal{T})]$.
3. If the new trajectory is accepted, it becomes the current one. Otherwise the old trajectory is retained as the current trajectory.
4. Repeat starting from 1.

Scheme 1.1. Metropolis Monte Carlo sampling algorithm for transition pathways.

In order to implement the algorithm described above, one must be able to generate trial paths from existing ones. While in principle there are many ways to do this, the efficiency of the algorithm depends crucially on the nature of trial moves. Specifically, it is important that the average acceptance probability for generated paths not be too small. For transition path sampling, satisfying this condition requires that trial paths have a reasonable weight $\mathcal{P}[x(\mathcal{T})]$ in the dynamical path ensemble, as well as a reasonable chance of connecting states A and B .

Desirable trial trajectories for transition path sampling have significant dynamical weights, $\mathcal{P}[x(\mathcal{T})]$, and thus resemble natural trajectories. The simplest way to obtain such paths is to apply the very same propagation rules that define the natural dynamics of the system. We have employed this strategy to construct two basic types of trial moves, which we call “shooting” and “shifting”. In both cases, new trajectories are obtained by applying dynamical propagation rules to a phase space point that is taken (and possibly modified) from an existing transition pathway. The resulting paths have significant dynamical weights by construction. In addition, trial paths generated by shooting and shifting have a good chance of exhibiting successful transitions from A to B , since they are grown from phase-space points on or near reactive trajectories. In Sections III.A. and III.B., we describe shooting and shifting moves in detail, and indicate how they may be successfully implemented for various classes of dynamics.

A. Shooting Moves

In a shooting move, a phase-space point $x_r^{(0)}$ is selected at random from the chain of states comprising an old pathway, $x^{(0)}(\mathcal{T})$. This state may be modified in some way, for example by displacing the atomic momenta by a small amount. Trajectory segments are then “shot off” forward and backward in time from the modified state $x_r^{(n)}$, applying the appropriate dynamical rules until the new path extends from time zero to time \mathcal{T} . This procedure is schematically depicted in Figure 1.5.

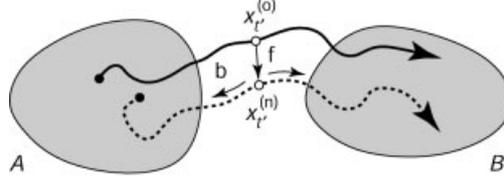


Figure 1.5. In a shooting move, a new pathway (dashed line) is generated from an old one (solid line) by first selecting a time slice $x_{t'}^{(o)}$ of the old path and modifying it to obtain $x_{t'}^{(n)}$. Then, a new pathway is constructed by generating forward (f) and backward (b) trajectory segments initiated at the modified time slice $x_{t'}^{(n)}$. If the new trajectory is reactive, that is, if it starts in A and ends in B , the path is accepted with a finite probability and rejected otherwise.

Because trajectories obtained by shooting reflect the system's underlying dynamics, the corresponding generation and acceptance probabilities are relatively simple. Specifically, the generation probability for the forward trajectory segment, beginning at time t' and ending at time \mathcal{T} , is

$$P_{\text{gen}}^f[\text{o} \rightarrow \text{n}] = \prod_{i=t'/\Delta t}^{\mathcal{T}/\Delta t - 1} p[x_{i\Delta t}^{(n)} \rightarrow x_{(i+1)\Delta t}^{(n)}] \quad (1.22)$$

This generation probability is identical to the dynamical path weight for the forward trajectory. Similarly, the generation probability for the backward trajectory segment, beginning at time t' and ending at time zero, is

$$P_{\text{gen}}^b[\text{o} \rightarrow \text{n}] = \prod_{i=1}^{t'/\Delta t} \bar{p}[x_{i\Delta t}^{(n)} \rightarrow x_{(i-1)\Delta t}^{(n)}] \quad (1.23)$$

Here, the appropriate small time step probability, $\bar{p}(x \rightarrow x')$, describes the evolution of the system backward in time.

By combining the generation probability for the modified time slice $x_{t'}^{(n)}$ with the generation probabilities of the forward and backward segments of the new trajectory, one obtains the complete generation probability for the new trajectory,

$$P_{\text{gen}}[x^{(o)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})] = p_{\text{gen}}[x_{t'}^{(o)} \rightarrow x_{t'}^{(n)}] \prod_{i=t'/\Delta t}^{\mathcal{T}/\Delta t - 1} p[x_{i\Delta t}^{(n)} \rightarrow x_{(i+1)\Delta t}^{(n)}] \\ \times \prod_{i=1}^{t'/\Delta t} \bar{p}[x_{i\Delta t}^{(n)} \rightarrow x_{(i-1)\Delta t}^{(n)}] \quad (1.24)$$

Here, $p_{\text{gen}}[x_{t'}^{(o)} \rightarrow x_{t'}^{(n)}]$ denotes the probability to obtain state $x_{t'}^{(n)}$, which we also call the *shooting point*, by modification of state $x_{t'}^{(o)}$.

The generation probability from Eq. (1.24) can now be used to determine the acceptance probability of a shooting move with the help of Eq. (1.21). For this purpose, the ratio appearing as the second argument of the min-function in Eq. (1.21) must be determined. By using the dynamical path probability from Eq. (1.1) for the old and the new paths and the generation probability from Eq. (1.24), we obtain

$$\frac{\mathcal{P}[x^{(n)}(\mathcal{T})]P_{\text{gen}}[x^{(n)}(\mathcal{T}) \rightarrow x^{(o)}(\mathcal{T})]}{\mathcal{P}[x^{(o)}(\mathcal{T})]P_{\text{gen}}[x^{(o)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})]} = \frac{\rho[x_0^{(n)}]p_{\text{gen}}[x_{t'}^{(n)} \rightarrow x_{t'}^{(o)}]}{\rho[x_0^{(o)}]p_{\text{gen}}[x_{t'}^{(o)} \rightarrow x_{t'}^{(n)}]} \prod_{i=0}^{t'/\Delta t - 1} \frac{p[x_{i\Delta t}^{(n)} \rightarrow x_{(i+1)\Delta t}^{(n)}]}{\bar{p}[x_{(i+1)\Delta t}^{(n)} \rightarrow x_{i\Delta t}^{(n)}]} \times \frac{\bar{p}[x_{(i+1)\Delta t}^{(o)} \rightarrow x_{i\Delta t}^{(o)}]}{p[x_{i\Delta t}^{(o)} \rightarrow x_{(i+1)\Delta t}^{(o)}]} \quad (1.25)$$

where factors have canceled because the trial trajectory $x^{(n)}(\mathcal{T})$ was generated using the propagation rule of the underlying dynamics.

This general result simplifies considerably if the dynamics conserve a stationary distribution, $\rho_{\text{st}}(x)$. This condition is very general, applying to systems at equilibrium, nonequilibrium systems in a steady state, and nonequilibrium systems relaxing to equilibrium with time-translationally invariant dynamics. In this case, p and \bar{p} are related in a simple way by microscopic reversibility,

$$\frac{p(x \rightarrow y)}{\bar{p}(y \rightarrow x)} = \frac{\rho_{\text{st}}(y)}{\rho_{\text{st}}(x)} \quad (1.26)$$

By substituting this relation into Eq. (1.25) and by using Eq. (1.21), the acceptance probability for a shooting move is written

$$P_{\text{acc}}[x^{(o)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})] = h_A[x_0^{(n)}]h_B[x_{\mathcal{T}}^{(n)}] \min \left\{ 1, \frac{\rho[x_0^{(n)}]}{\rho[x_0^{(o)}]} \frac{\rho_{\text{st}}[x_0^{(o)}]}{\rho_{\text{st}}[x_0^{(n)}]} \frac{\rho_{\text{st}}[x_{t'}^{(n)}]}{\rho_{\text{st}}[x_{t'}^{(o)}]} \frac{p_{\text{gen}}[x_{t'}^{(n)} \rightarrow x_{t'}^{(o)}]}{p_{\text{gen}}[x_{t'}^{(o)} \rightarrow x_{t'}^{(n)}]} \right\} \quad (1.27)$$

The above acceptance probability depends only on phase-space points at two times, at 0 and t' .

Often, the distribution of initial conditions is an equilibrium or steady-state distribution, and is therefore identical to the conserved distribution, $\rho(x) = \rho_{\text{st}}(x)$. In this case, the acceptance probability simplifies still further,

$$P_{\text{acc}}[x^{(o)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})] = h_A[x_0^{(n)}]h_B[x_{\mathcal{T}}^{(n)}] \times \min \left\{ 1, \frac{\rho[x_{t'}^{(n)}]}{\rho[x_{t'}^{(o)}]} \times \frac{p_{\text{gen}}[x_{t'}^{(n)} \rightarrow x_{t'}^{(o)}]}{p_{\text{gen}}[x_{t'}^{(o)} \rightarrow x_{t'}^{(n)}]} \right\} \quad (1.28)$$

For symmetric generation probabilities of the shooting point, this equation becomes

$$P_{\text{acc}}[x^{(0)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})] = h_A[x_0^{(n)}]h_B[x_{\mathcal{T}}^{(n)}] \min \left\{ 1, \frac{\rho[x_{t'}^{(n)}]}{\rho[x_{t'}^{(0)}]} \right\}. \quad (1.29)$$

Next, we will assume that generation probabilities for phase-space modifications are symmetric. If they are not, the asymmetry must be taken into account as prescribed by Eq. (1.28). The simplicity of the acceptance probability in Eq. (1.29) entails an algorithmic simplicity for shooting moves. In order to evaluate P_{acc} , one needs only to compute the relative weights of old and new phase-space points at time t' and determine whether the new path begins in region A and ends in region B .

Equation (1.29) also suggests an efficient implementation of shooting moves. A shooting move is initiated by selecting a time slice along the existing path. After generating the shooting point $x_{t'}^{(n)}$ from the old time slice $x_{t'}^{(0)}$ a first acceptance–rejection decision can be made. The shooting point is accepted if $\rho[x_{t'}^{(n)}]/\rho[x_{t'}^{(0)}] > 1$ or if a random number ζ drawn from a uniform distribution in the interval $[0, 1]$ is smaller than $\rho[x_{t'}^{(n)}]/\rho[x_{t'}^{(0)}]$. Otherwise the trial shooting move is aborted. If, on the other hand, the shooting point is accepted, one proceeds with growing either one of the forward or backward trajectory segments. If the appropriate boundary condition is satisfied (e.g., if the forward trajectory segment ends in region B), the other segment is grown. The path is finally accepted if the boundary condition for this latter trajectory segment is satisfied as well. A rejection at any stage of this procedure allows subsequent steps to be skipped. In this case, the trial move is rejected and the old path is retained. Using this sequential algorithm reduces the cost of rejected moves, saving considerable computational effort (typically 10–30%).

Naively, one might expect that a gain in computational efficiency might be obtained by growing the shorter of the two trajectory segments first. In case of rejection, the longer and more expensive of the two segments need not be grown. In most cases, however, the shorter trajectory segment has a much higher probability to reach the appropriate stable state than the longer one, which is due to the fact that the shorter trajectory segment has less time to diverge from the old path than the longer segment. Also, the short trajectory segment is less likely to have to pass the dynamical bottleneck separating the stable states. Consequently, early rejection is not likely to occur when growing the short segment first and the long trajectory segment must be determined in most cases. This effect can compensate the potential advantage of occasional early rejections. Indeed, for some systems

it might even be advantageous to grow the longer trajectory segment first and proceed with growing the shorter one only in case of acceptance. Since in general the associated efficiency increase is only of the order of a few percent, it is good practice to carry out integration of, say, the backward segment first and after that the forward segment regardless of which one of the two is shorter.

Shooting moves derive their efficiency from the tendency of trajectories to diverge in phase space, so that subsequent paths may be quite different. But it is important that trial paths are not too different, so that they have reasonable chance of connecting A and B . In deterministic systems, the degree of divergence between old and new trajectories in shooting moves depends on the magnitude of modification of $x_t^{(o)}$. Arbitrarily small modification gives new paths that are arbitrarily similar to old paths. Large displacements give very different paths. It is thus possible to tune the magnitude of modifications to give a desired acceptance probability on average, as will be discussed in Section III.A.4. For stochastic systems, small path displacements may not be possible, since even an unmodified state can yield a different path.

1. Shooting Moves for Deterministic Dynamics

For deterministic dynamics, applying dynamical propagation rules to a phase-space point x_t^o of an existing trajectory simply regenerates that trajectory. Shooting moves must therefore include modification of the shooting point, for example, by adding a random perturbation δx to $x_t^{(o)}$,

$$x_t^{(n)} = x_t^{(o)} + \delta x \quad (1.30)$$

This can be done in a symmetric way by drawing δx from a distribution $w(\delta x)$ for which $w(\delta x) = w(-\delta x)$. In this procedure, the generation probability of the new shooting point $x_t^{(n)}$ from $x_t^{(o)}$ is identical to that for the reverse move, that is, $p_{\text{gen}}[x_t^{(o)} \rightarrow x_t^{(n)}] = p_{\text{gen}}[x_t^{(n)} \rightarrow x_t^{(o)}]$. Selecting symmetric perturbations can be complicated when allowed values of x are restricted by internal constraints of the system. Methods for generating such displacements are discussed in Section III.A.2.

It is usually sufficient to modify only the momentum part of the selected time slice $x_t^{(o)}$, leaving the configurational part unchanged. In this case, $p_t^{(n)} = p_t^{(o)} + \delta p$ and $r_t^{(n)} = r_t^{(o)}$ (see Fig. 1.6). For some applications, however, it is advantageous to change both the configuration and momentum parts of $x_t^{(o)}$ [13].

For deterministic dynamics, a trial trajectory is obtained by integrating the equations of motion from the shooting point. The forward segment of

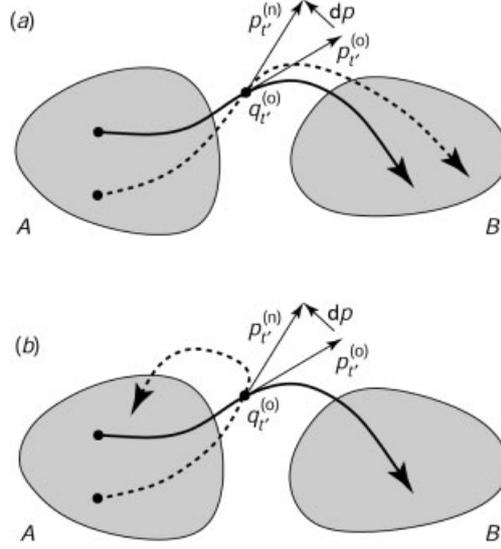


Figure 1.6. In a shooting move for deterministic trajectories, a time slice $x_{r'}^{(o)} = \{q_{r'}^{(o)}, p_{r'}^{(o)}\}$ on the old path (solid line) is selected at random and the corresponding momenta $p_{r'}^{(o)}$ are changed by a small random amount δp . Integration of the equations of motion backward to time 0 and forward to time t starting from the modified state $x_{r'}^{(n)} = \{q_{r'}^{(n)}, p_{r'}^{(n)}\}$ yields the new trajectory $x^{(n)}(\mathcal{T})$ (dashed line). If this trajectory is connecting A and B , it will be accepted with a nonvanishing probability (panel a). Otherwise, it will be rejected (panel b).

the trajectory is generated simply by integrating for the appropriate number of small time steps. The backward segment, on the other hand, must be generated with the direction of time inverted, that is, with a negative time step. In the case of time-reversible dynamics, this is accomplished by first inverting all momenta, and then integrating forward in time. Formally, $\bar{p}(x \rightarrow y) = p(\bar{x} \rightarrow \bar{y})$. Here, $\bar{x} = \{r, -p\}$ is obtained from $x = \{r, p\}$ by inverting all momenta. In the resulting chain of states, momenta are then inverted so that the backward trajectory segment evolves properly in time. The complete protocol for performing shooting moves with deterministic dynamics is summarized in Scheme 1.2.

We verify the soundness of this procedure by demonstrating that the reversibility condition in Eq. (1.26) is satisfied. A small step in the backward trajectory segment described above has probability

$$\bar{p}(y \rightarrow x) = \delta[x - \phi_{-\Delta t}(y)] = \delta[x - \phi_{\Delta t}^{-1}(y)] \quad (1.31)$$

1. Randomly select a time slice $x_t^{(o)}$ on an existing trajectory $x^{(o)}(\mathcal{T})$.
2. Modify the selected time slice by adding a random displacement: $x_t^{(n)} = x_t^{(o)} + \delta x$. The random displacement must be consistent with the ensemble of initial conditions and should be symmetric with respect to the reverse move.
3. Accept the new shooting point with probability $\min[1, \rho(x_t^{(n)})/\rho(x_t^{(o)})]$. Abort the trial move if the shooting point is rejected.
4. If the shooting point is accepted, integrate the equations of motion forward to time \mathcal{T} starting from $x_t^{(n)}$.
5. Abort the trial move if the final point of the path segment, $x_{\mathcal{T}}^{(n)}$, is not in B and continue otherwise.
6. Integrate the equations of motion backward to time 0 starting from $x_t^{(n)}$.
7. Accept the new trajectory if its initial point $x_0^{(n)}$ is in A and reject it otherwise.
8. In case of a rejection, the old trajectory is counted again in the calculation of path averages. Otherwise the new trajectory is used as the current one.

Scheme 1.2. Shooting algorithm for deterministic trajectories.

For the forward trajectory,

$$p(x \rightarrow y) = \delta[y - \phi_{\Delta t}(x)] = \delta[\phi_{\Delta t}^{-1}(y) - x] |\partial \phi_{\Delta t}(x) / \partial x|^{-1} \quad (1.32)$$

where $|\partial \phi_{\Delta t}(x) / \partial x|$ is the Jacobian associated with time evolution of duration Δt . This Jacobian describes the contraction or expansion of an infinitesimal comoving volume element in phase space. Combining these small time step probabilities,

$$\frac{p(x \rightarrow y)}{\bar{p}(y \rightarrow x)} = \left| \frac{\partial \phi_{\Delta t}(x)}{\partial x} \right|^{-1} = \frac{\rho_{\text{st}}(y)}{\rho_{\text{st}}(x)} \quad (1.33)$$

The final equality results from the fact that the phase-space flow generated by time evolution conserves probability. Because the reversibility criterion is satisfied, the simple acceptance probabilities from Eqs. (1.27) and (1.29) can be used, provided a stationary distribution ρ_{st} exists.

For some applications, the distribution of initial conditions is distinct from the stationary distribution preserved by a system's dynamics. For example, one might be interested in the relaxation of a system that has been driven away from equilibrium. One might even be interested in an unstable dynamics that does not preserve a stationary distribution. In both cases, it is possible to sample trajectories using the same type of shooting move described above, provided the distribution of initial conditions $\rho(x)$ is well defined. To derive an appropriate acceptance probability for this case, we return to Eq. (1.25). By using Eq. (1.32) and the chain rule for the Jacobian,

we obtain

$$P_{\text{acc}}[x^{(0)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})] = h_A[x_0^{(n)}]h_B[x_{\mathcal{T}}^{(n)}] \min \left\{ 1, \frac{\rho[x_0^{(n)}]}{\rho[x_0^{(0)}]} \frac{|\partial x_{t'}^{(0)}/\partial x_0^{(0)}|}{|\partial x_{t'}^{(n)}/\partial x_0^{(n)}|} \right\} \quad (1.34)$$

where we have again assumed a symmetric generation probability, that is, $p_{\text{gen}}[x_{t'}^{(0)} \rightarrow x_{t'}^{(n)}] = p_{\text{gen}}[x_{t'}^{(n)} \rightarrow x_{t'}^{(0)}]$.

The Jacobian $J(t) \equiv |\partial x_{t'}/\partial x_0|$ appearing in Eq. (1.34) can in principle be calculated by integrating

$$\frac{dJ(t)}{dt} = \Lambda(x)J(t) \quad (1.35)$$

along the trajectory of interest with $J(0) = 1$. In Eq. (1.35), $\Lambda(x) \equiv \text{Tr}(\partial \dot{x}/\partial x)$ is the so-called *phase-space compressibility* of the dynamical system at x . In general, $\Lambda(x)$ is nonzero, and the Jacobian consequently differs from unity. But for Newtonian dynamics, Liouville's theorem guarantees that the volume of comoving phase-space elements is conserved [46]. The Jacobian is thus unity, and P_{acc} has a particularly simple form:

$$P_{\text{acc}}[x^{(0)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})] = h_A[x_0^{(n)}]h_B[x_{\mathcal{T}}^{(n)}] \min \left\{ 1, \frac{\rho[x_0^{(n)}]}{\rho[x_0^{(0)}]} \right\} \quad (1.36)$$

Remarkably, even when a stationary distribution does not exist, the acceptance probability for Newtonian shooting moves depends only on the relative weights of the initial conditions, and on the reactivity of the trial trajectory.

2. Selecting Phase-Space Displacements

For the shooting algorithm we have described, acceptance probabilities are particularly simple if phase-space modifications have a symmetric generation probability. If an asymmetry is present, that is, $p_{\text{gen}}[x_{t'}^{(0)} \rightarrow x_{t'}^{(n)}] \neq p_{\text{gen}}[x_{t'}^{(n)} \rightarrow x_{t'}^{(0)}]$, and is not accounted for in the acceptance probability, the trajectories harvested by path sampling will not be correctly weighted. As indicated above, modifying states in a symmetric way can be challenging when constraints are present. A general formal procedure for taking into account linear constraints on momenta was presented in the Appendix of [10]. Here, we describe straightforward procedures for displacing momenta symmetrically, while satisfying common constraints such as fixed total linear or angular momentum, or rigid intramolecular bonds. These simple pro-

cedures are equivalent to the more general methods presented in [10]. For clarity, we will discuss this issue using proton transfer in the protonated water trimer as discussed in Section I as an illustrative example.

Proton transfer in the protonated water trimer has been studied extensively with transition path sampling using empirical and ab initio models [10,15]. In these studies, shooting moves were implemented by using momentum displacements δp only. Since in the classical limit an isolated cluster evolves at constant energy E according to Newton's equation of motion, the simulations were carried out in the microcanonical ensemble, that is, $\rho(x) \propto \delta[E - H(x)]$. Furthermore, the dynamics conserves the total linear momentum \mathbf{P} and the total angular momentum \mathbf{L} . Thus the complete distribution of initial conditions is

$$\rho(x) \propto \delta[E - H(x)]\delta[\mathbf{P}(x)]\delta[\mathbf{L}(x)] \quad (1.37)$$

where both the total linear momentum and the total angular momentum are assumed to vanish. The delta functions in expression (1.37) take into account the reduced dimensionality of the accessible phase space caused by conservation of these quantities.

To obtain a nonvanishing acceptance probability for shooting moves the perturbation δp used to construct a new shooting point $x_r^{(n)}$ must be consistent with the ensemble of initial conditions. For the cluster at constant energy, total linear momentum, and total angular momentum, the displacement δx must be chosen to conserve these quantities. Furthermore, as mentioned above, it must produce a symmetric path generation probability, which can be accomplished as follows.

First, one selects a momentum displacement δp from a Gaussian distribution that is added to the old momentum vector $p_r^{(o)}$. The total linear momentum of the new momentum p' is set to zero by subtracting $\sum_i \mathbf{p}'_i/N$ from all single particle momenta. Next, the total angular momentum $\mathbf{L}' = \sum \mathbf{r}_i \times \mathbf{p}'_i$ is set to zero. This can be accomplished with a procedure proposed by Laria et al. [25] and used in simulations of water clusters. For this purpose, one first calculates the angular velocity $\boldsymbol{\omega} = I^{-1}\mathbf{L}'$, where $I \equiv \sum_i m_i(\mathbf{r}_i^2 - \mathbf{r}_i \mathbf{r}_i)$ is the inertia tensor. Note that calculation of the angular velocity $\boldsymbol{\omega}$ requires inversion of the inertia tensor. Then, one calculates new momenta

$$\mathbf{p}''_i = \mathbf{p}'_i - m_i \boldsymbol{\omega} \times \mathbf{r}_i \quad (1.38)$$

Finally, the new momenta are rescaled to obtain the appropriate total kinetic energy. It can be shown that provided the center of mass of the system is located in the origin, the new momenta obey all constraints

dictated by the conserved total energy, linear, and angular momentum of the system.

This procedure guarantees that the generation probability for new trajectories is equal to the probability of the reverse move. It therefore explicitly satisfies detailed balance. Similar procedures can be used to remove components of δp violating the constraint of constant bond lengths. Molecules with fixed bond lengths are often used in molecular simulations to eliminate the fast oscillatory motion induced by stiff intramolecular bonds. If the distribution of initial conditions is microcanonical and the dynamics is Newtonian, the acceptance probability for a shooting move generated according to this procedure is

$$P_{\text{acc}}[x^{(0)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})] = h_A[x_0^{(n)}]h_B[x_{\mathcal{T}}^{(n)}] \quad (1.39)$$

This acceptance probability implies that trajectories are always accepted if they connect regions A and B .

Phase-space displacements obeying the linear constraints of vanishing total linear momentum and total angular momentum, and of fixed bond lengths can be also generated using an iterative procedure such as RATTLE [47].

3. Momentum Rescaling

If the distribution of initial conditions $\rho(x_0)$, such as the canonical distribution, allows for variations of the energy, shooting points with different energies must be created. This can be accomplished by adding a momentum displacement δp chosen from an appropriate distribution to a given momentum $p_i^{(0)}$ without rescaling the momenta to a fixed total energy. Large momentum changes, however, most likely produce large changes in the total energy of the system and therefore lead to a low acceptance probability. This problem can be solved by alternating constant energy shooting moves with moves in which the energy is changed by rescaling the momenta. This approach allows one to control the change in momentum and the change in energy independently.

Such an approach was necessary in recent work by Geissler and Chandler [16] in which they studied the nonlinear response of a polar solvent to electronic transitions of a solute with the transition path sampling method. In this study, nonequilibrium trajectories relaxing from states obtained by inversion of the solute dipole were harvested with a variant of the shooting algorithm. Since such solvation dynamics occurs rapidly, the trajectories of interest are only tens of femtoseconds in duration. Accordingly, large momentum displacements at the shooting point are required to generate new pathways sufficiently different from the old ones.

The energy changing move alternated with constant energy shooting moves described in Section III.A can be carried out as follows. First, a time slice $x_{t'}^{(o)}$ along the old path $x^{(o)}(\mathcal{T})$ is randomly chosen. At time slice $x_{t'}^{(o)}$, the system has a kinetic energy of $K_{t'}^{(o)}$. Then, a new kinetic energy $K_{t'}^{(n)}$ at time t' is selected from a distribution $\psi(K)$. Although any form of $\psi(K)$ can be selected, a smart choice of $\psi(K)$ can increase the efficiency of the simulation considerably. In general, the distribution $\psi(K)$ should be similar to the actual distribution of kinetic energies observed in the system under study. In the study carried out by Geissler and Chandler [16], for example, $\psi(K)$ was chosen to be the distribution of kinetic energies in an equilibrium system in contact with a heat bath. Next, a new time slice $x_{t'}^{(n)}$ is generated by rescaling the old momenta $p_{t'}^{(o)}$ with a factor of $\gamma = (K_{t'}^{(n)}/K_{t'}^{(o)})^{1/2}$ to give the kinetic energy $K_{t'}^{(n)}$. Note that rescaling of the momenta does not affect the total linear momentum and the total angular momentum provided that they vanish.

In the above procedure, shooting points are generated with relative probability,

$$\frac{p_{\text{gen}}[x_{t'}^{(n)} \rightarrow x_{t'}^{(o)}]}{p_{\text{gen}}[x_{t'}^{(o)} \rightarrow x_{t'}^{(n)}]} = \frac{\psi[K_{t'}^{(o)}]dK_{t'}^{(o)}}{\psi[K_{t'}^{(n)}]dK_{t'}^{(n)}} \times \frac{\delta \left[p_{t'}^{(o)} - \sqrt{\frac{K_{t'}^{(o)}}{K_{t'}^{(n)}}} p_{t'}^{(n)} \right]}{\delta \left[p_{t'}^{(n)} - \sqrt{\frac{K_{t'}^{(n)}}{K_{t'}^{(o)}}} p_{t'}^{(o)} \right]} \quad (1.40)$$

Due to the presence of the delta function the differential $dK_{t'}^{(o)}$ of the old kinetic energy can be expressed as function of the differential $dK_{t'}^{(n)}$ of the new kinetic energy,

$$dK_{t'}^{(o)} = \sum_{\alpha} \frac{1}{m_{\alpha}} p_{\alpha,t'}^{(o)} dp_{\alpha,t'}^{(o)} = \left[\frac{K_{t'}^{(o)}}{K_{t'}^{(n)}} \right] \sum_{\alpha} \frac{1}{m_{\alpha}} p_{\alpha,t'}^{(n)} dp_{\alpha,t'}^{(n)} = \left[\frac{K_{t'}^{(o)}}{K_{t'}^{(n)}} \right] dK_{t'}^{(n)} \quad (1.41)$$

where $p_{\alpha,t'}$ is component α of momentum $p_{t'}$ and m_{α} is the mass associated with degree of freedom α . By using this result and by applying Eq. (1.32) to simplify the ratio of delta functions in Eq. (1.40), we obtain

$$\frac{p_{\text{gen}}[x_{t'}^{(n)} \rightarrow x_{t'}^{(o)}]}{p_{\text{gen}}[x_{t'}^{(o)} \rightarrow x_{t'}^{(n)}]} = \frac{\psi[K_{t'}^{(o)}]}{\psi[K_{t'}^{(n)}]} \times \left[\frac{K_{t'}^{(o)}}{K_{t'}^{(n)}} \right]^{f/2-1}. \quad (1.42)$$

where f is the number of independent degrees of freedom. In determining f , the number of constraints acting on the system must be properly taken into account.

Assuming the dynamics conserves phase-space volume and correctly taking into account the asymmetric generation probability from the above equation yields an acceptance probability of

$$P_{\text{acc}}[x^{(0)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})] = h_A[x_0^{(n)}]h_B[x_{\mathcal{T}}^{(n)}] \times \min \left\{ 1, \frac{\rho[x_0^{(n)}]}{\rho[x_0^{(0)}]} \times \frac{\psi[K_{t'}^{(0)}]}{\psi[K_{t'}^{(n)}]} \times \left[\frac{K_{t'}^{(0)}}{K_{t'}^{(n)}} \right]^{f/2-1} \right\} \quad (1.43)$$

Thus, new pathways connecting A and B will be accepted depending on the change in kinetic energy at time t' and on the change of the probability of the initial conditions at time 0.

4. Efficiency of Deterministic Shooting Moves

The efficiency of shooting moves can be controlled by tuning the magnitude of phase-space displacements δx . An optimal magnitude minimizes the correlations between harvested paths. If the magnitude of δx is very small the new trajectory essentially retraces the old one. Although the average acceptance probability is near unity in this case, the random walk in trajectory space will progress slowly because of strong correlations between subsequent trajectories. If, on the other hand, the magnitude of δx is large, the new trajectory will drastically differ from the old one due to the chaotic nature of the underlying dynamics. In this case, the new trajectory most likely will not connect A with B and the resulting acceptance probability is near zero. Pathways are repeated many times before a new path is accepted. Again, correlations between subsequent pathways decay slowly. The optimum magnitude of the perturbation δx lies between these two extremes, which is completely analogous to the situation in conventional Monte Carlo simulations. In that case, the optimum acceptance probability is often near 50% [44].

The effect of the magnitude of δx on the efficiency of transition path sampling can be systematically analyzed by calculating correlation functions of various quantities as a function of the number of simulations cycles. Ideally, such correlation functions decay quickly, indicating that path space is sampled with high efficiency. In [11], we carried out such an efficiency analysis for transition path sampling of isomerizations of a model dimer immersed in a soft sphere liquid. In that study, we calculated correlation functions

$$c(n) \equiv \frac{\langle \delta G(0) \delta G(n) \rangle_{AB}}{\langle \delta G^2 \rangle_{AB}} \quad (1.44)$$

where n is the number of simulation cycles. The fluctuation δG is defined as $\delta G(n) \equiv G(n) - \langle G \rangle_{AB}$ and $G[x(\mathcal{T})]$ is a quantity depending on the path coordinates $x(\mathcal{T})$. The notation $\langle G \rangle_{AB}$ indicates a weighted average of the path functional $G[x(\mathcal{T})]$ over all pathways in the transition path ensemble,

$$\langle G \rangle_{AB} \equiv \int \mathcal{D}x(\mathcal{T}) \mathcal{P}_{AB}[x(\mathcal{T})] G[x(\mathcal{T})] \quad (1.45)$$

Since the sequence of pathways $\{x_i(\mathcal{T})\}$ generated in a transition path sampling procedure visits pathways according to their weight in the transition path ensemble, path averages $\langle G \rangle_{AB}$ can be calculated as averages over the sequence of pathways $\{x_i(\mathcal{T})\}$,

$$\langle G \rangle_{AB} = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_i^N G[x_i(\mathcal{T})] \quad (1.46)$$

where N is the total number of simulations cycles. Accordingly, the correlation function $c(n)$ can be expressed as

$$c(n) = \frac{\lim_{N \rightarrow \infty} (1/N) \sum_i^N \delta G[x_i(\mathcal{T})] \delta G[x_{i+n}(\mathcal{T})]}{\lim_{N \rightarrow \infty} (1/N) \sum_i^N \delta G^2[x_i(\mathcal{T})]} \quad (1.47)$$

Correlation functions $c(n)$ are shown in Figure 1.7(a) for three different quantities G : the final region characteristic function $h_B(x_{\mathcal{T}|2})$ at the path midpoint $x_{\mathcal{T}|2}$, the potential energy V_{tr} of the system when the dimer surpasses the potential energy barrier separating the stable states, and the transition time τ . The transition time τ is the time required to reach final region B after the system has left the initial region A . These correlation functions were obtained in a simulation for which the average acceptance probability for shooting moves was $\sim 40\%$.

From the correlation functions $c(n)$ one can compute the number of correlated cycles, n_c by determining at which cycle n the correlation function falls below a certain threshold, say, 0.5. The number n_c of correlated cycles is a measure of the efficiency of the simulation. For the isomerization of the model dimer, we have studied how n_c depends on the magnitude of the momentum displacement. To obtain a more generally usable criterion for the choice of δp , we have plotted n_c as a function of the average acceptance probability P_{acc} of shooting moves. This acceptance probability is monotonically related to the magnitude of the momentum displacement δp . A small δp yields a high acceptance probability and a large δp leads to a small one.

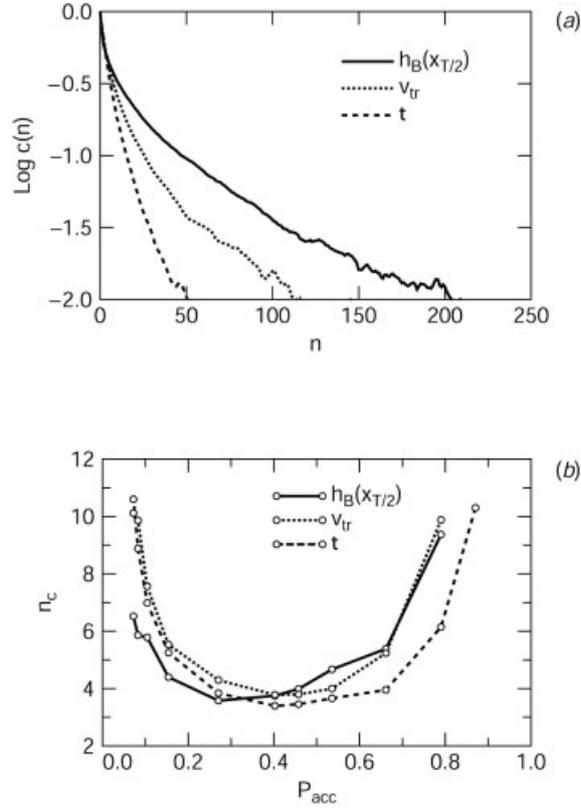


Figure 1.7. (a) Typical correlation functions for the three quantities $h_B(x_{T/2})$ (solid line), V_{tr} (dotted line), and τ (dashed line) as a function of the number of simulation cycles calculated for isomerizations in a model dimer [11]. (b) Number of correlated cycles as a function of the acceptance probability P_{acc} . The different line types refer to results obtained from analyzing correlation functions of $h_B(x_{T/2})$ (solid line), V_{tr} (dotted line), and τ (dashed line).

The results of this analysis are depicted in Figure 1.7(b), which shows the number of correlated cycles n_c as a function of the average acceptance probability P_{acc} . The three different curves represent results obtained by analyzing the correlation functions of $h_B(x_{T/2})$, V_{tr} , and τ . In all three cases, the number of correlated cycles is high for low and high acceptance probabilities and has a minimum for intermediate acceptance probabilities. These results suggest that, as a rule of thumb, the magnitude of the random displacement δx should be chosen to obtain an acceptance probability of $\sim 40\%$. Since rejected moves are computationally less expensive on the

average than accepted moves, acceptance probabilities even lower than 40% might be optimal in some cases.

The curves depicted in Figure 1.7 indicate that in transition path sampling a statistically independent trajectory is obtained after a small number (say three to five) of accepted trial moves. Thus, the computational cost of harvesting N independent trajectory of length \mathcal{T} is comparable to the effort required to compute a single long trajectory of length $N\mathcal{T}$. This linear scaling in trajectory length \mathcal{T} and number of trajectories N allows application of the transition path sampling method to the simulation of complex, high-dimensional systems, such as chemical reactions in solution or isomerizations of biomolecules.

In some cases, the applicability of current shooting algorithms can be limited by the Lyapunov instability of the underlying dynamics. This mechanical instability leading to chaotic behavior, is best quantified with the so called Lyapunov exponents λ describing the divergence in phase-space of neighboring trajectories. On the average, the separation δ between two initially close points in phase-space grows in time like $\delta \propto \exp(\lambda t)$. Due to this fast exponential divergence, trajectories generated by “shooting” from only slightly different phase-space points become completely different after a relatively short time. In typical simple liquids, Lyapunov exponents are of the order of $\lambda \sim 1\text{ps}^{-1}$ [48]. Accordingly, a perturbation of the order of the computer precision, $\sim 10^{-15}$ for double precision arithmetic (i.e., the smallest possible perturbation on a digital computer), needs of the order of 10 ps to grow to the scale of the liquid structure. Doubling the number of bits in the representation of numbers would increase this time only by a factor of 2. The average shooting acceptance probability of trajectories longer than this characteristic time is therefore low. Accordingly, processes in which the crucial fluctuations take much longer to occur than this critical time are difficult to study with current shooting algorithms [17].

5. Shooting Moves for Stochastic Dynamics

Shooting moves are efficient for stochastic dynamics as well. In most aspects, these moves are identical to the shooting moves employed for deterministic dynamics. Some important details, however, are different. These differences will be discussed in this section. The shooting procedure is schematically illustrated in Figure 1.8 and summarized in Scheme 1.3.

Imagine a stochastic path $x^{(0)}(\mathcal{T})$ of length \mathcal{T} starting in region A and ending in region B . One can randomly select a point $x_r^{(0)}$ along this path, initiate a stochastic trajectory starting from $x_r^{(0)}$ and integrate it forward to time \mathcal{T} using the propagation rule corresponding to the underlying dynamics, for example Eq. (1.11) or (1.14). Then the backward segment of the

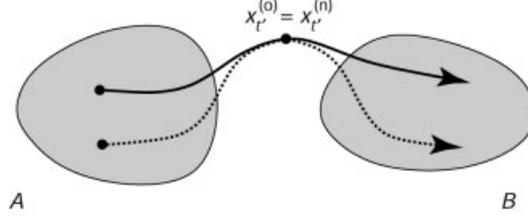


Figure 1.8. In a shooting move for stochastic dynamics forward and backward trajectory segments can be initiated from an unchanged shooting point $x_{t'}^{(n)} = x_{t'}^{(o)}$.

trajectory is generated using the time inverted dynamics satisfying Eq. (1.26). Alternatively, the backward segment can be grown before the forward segment. In contrast to the case of deterministic dynamics, it is not necessary to modify $x_{t'}^{(o)}$ before shooting, because the random nature of the dynamical propagation will alone cause the new path to diverge from the old one. In this case, the generation probability $p_{\text{gen}}[x_{t'}^{(o)} \rightarrow x_{t'}^{(n)}]$ is trivially symmetric and the acceptance probability of Eq. (1.28) simplifies to

$$P_{\text{acc}}[x^{(o)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})] = h_A[x_0^{(n)}]h_B[x_{\mathcal{T}}^{(n)}] \quad (1.48)$$

Thus, any new pathway connecting regions A and B can be accepted.

For Brownian dynamics described in Section II.D, the backward trajectory segment can be generated by integration of the equations of motion with an appropriate small time step algorithm [6]. Before initiating a trajectory, the momenta belonging to the shooting point are inverted. Then, a new trajectory segment of length t' is initiated at the

1. Select a time slice $x_{t'}^{(o)}$ at random from the time slices of an existing path $x^{(o)}(\mathcal{T})$.
2. Compute a new stochastic trajectory segment from t' to \mathcal{T} starting from $x_{t'}^{(o)}$ using the propagation rule of the underlying dynamics.
3. The new trajectory segment is rejected if its last state does not lie in region B , that is, if $h_B[x_{\mathcal{T}}^{(n)}] = 0$. Otherwise one proceeds with the next step.
4. If necessary, invert momenta at the selected time slice $x_{t'}^{(n)}$.
5. Compute a stochastic trajectory back to time 0 starting from the configuration with inverted momenta using the propagation rule of the underlying dynamics.
6. Invert the momenta along the newly generated trajectory segment so that the whole path evolves in forward direction.
7. The new trajectory is accepted if its initial state lies in region A , that is, if $h_A[x_0^{(n)}] = 1$.

Scheme 1.3. Shooting algorithm for stochastic trajectories.

shooting point with inverted momenta and is integrated back to time 0. After that, the momenta in this newly generated path segment are reversed such that the whole path evolves in forward direction. Accordingly, the small time step generation probability associated with this type of backward shot is

$$\bar{p}(x \rightarrow y) = p(\bar{x} \rightarrow \bar{y}) \quad (1.49)$$

where, again, $\bar{x} = \{r, -p\}$ is obtained from $x = \{r, p\}$ by inverting all momenta. It can be easily shown that this generation probability satisfies condition (1.26) with the canonical distribution as the stationary distribution [6]. Hence, the acceptance probability from Eq. (1.48) is valid for stochastic trajectories generated by numerical solution of Langevin's equation of motion.

In the high friction limit, the inertial term appearing in the Langevin equation can be neglected leading to simplified equations of motion in which the state of the system is completely described by the particle positions alone. In this case, the above procedure can be carried out using the unmodified forward integration algorithm to generate the backward trajectory segment.

As for Langevin dynamics in the high friction limit time does not appear in the Metropolis Monte Carlo transition probability from Eq. (1.14). Forward and backward propagation are therefore indistinguishable and the backward trajectory is generated with the forward propagation rule. Accordingly, the backward generation probability is identical to the forward generation probability, that is, $\bar{p}(x \rightarrow y) = p(x \rightarrow y)$. Condition (1.26) therefore becomes identical to the detailed balance condition which the Monte Carlo algorithm satisfies by construction,

$$\frac{p(x \rightarrow y)}{p(y \rightarrow x)} = \exp\{-\beta[H(y) - H(x)]\} \quad (1.50)$$

Hence, Eq. (1.48) is valid also for Monte Carlo dynamics and a new trajectory can be accepted if it connects region A with region B .

Another important characteristic of shooting moves for stochastic trajectories is that it is not necessary to shoot forward and backward simultaneously, that is, replacing only the forward or backward segment of an old path gives a trial path with significant dynamical weight. In contrast, growing only one of the two segments for deterministic dynamics yields trajectories with zero dynamical weight.

B. Shifting Moves

A shifting move translates an existing pathway forward or backward in time. As depicted in Figure 1.9, this move is graphically similar to the “reptation” motion of a polymer confined to a microscopic tube [49]. In a shifting move, a trial trajectory is obtained by first deleting a segment of length δt from the beginning (or end) of an existing trajectory. A new trajectory segment of length δt is then grown from the opposite end of the old path, by applying the dynamical propagation rules. These operations effectively shift the pathway forward (or backward) in time.

As in the case of shooting moves, the generation and acceptance rules for shifting moves are quite simple. The generation probability for a forward shifting move is just the dynamical weight of the newly generated trajectory segment appended to the end of the path:

$$P_{\text{gen}}^f[x^{(o)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})] = \prod_{i=(\mathcal{T}-\delta t)/\Delta t}^{\mathcal{T}/\Delta t-1} p[x_{i\Delta t}^{(n)} \rightarrow x_{(i+1)\Delta t}^{(n)}] \quad (1.51)$$

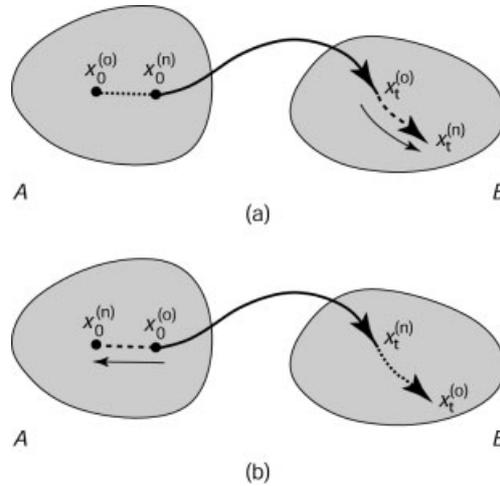


Figure 1.9. In a forward shifting move as shown in panel (a), a new trajectory is generated by removing $\delta t/\Delta t$ time slices (dotted line) from the beginning of the old path $x^{(o)}(\mathcal{T})$ and regrowing the same number of time slices (dashed line) at the end of the old path. If the time interval δt is small, most part of the new trajectory coincides with the old one and only a small segment of the new trajectory must be integrated. Panel (b) shows a backward shifting move in which $\delta t/\Delta t$ time slices (dotted line) are removed from the end of the old path and a new path segments if grown backward starting from the beginning of the old pathway (dashed line).

Here, it is understood that for times $t' = 0$ through $t' = t - \delta t$, the new path is identical to a portion of the old path,

$$x_{i\Delta t}^{(n)} = x_{i\Delta t + \delta t}^{(o)} \quad \text{for } i = 0, \dots, (\mathcal{T} - \delta t)/\Delta t \quad (1.52)$$

In a backward shifting move, a new trajectory segment affixed to the beginning of the path is obtained by time-reversed propagation. Correspondingly, the dynamical weight of this segment is composed of time-reversed transition probabilities $\bar{p}(x \rightarrow y)$. The generation probability for backward shifting is thus

$$P_{\text{gen}}^b[x^{(o)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})] = \prod_{i=1}^{\delta t/\Delta t} \bar{p}[x_{i\Delta t}^{(n)} \rightarrow x_{(i-1)\Delta t}^{(n)}] \quad (1.53)$$

where it is understood that

$$x_{i\Delta t + \delta t}^{(n)} = x_{i\Delta t}^{(o)} \quad \text{for } i = 0, \dots, (\mathcal{T} - \delta t)/\Delta t \quad (1.54)$$

An appropriate acceptance probability for shifting moves is most easily obtained by requiring that forward and backward moves are performed with equal frequency. We also require that the shifting length, δt , which is required to be a multiple of Δt , is drawn from the same distribution for forward and backward moves. With these restrictions, it is sufficient that detailed balance is satisfied for each pair of forward and backward moves

$$\begin{aligned} \mathcal{P}_{AB}[x^{(o)}(\mathcal{T})] P_{\text{acc}}^f[x^{(o)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})] P_{\text{gen}}^f[x^{(o)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})] \\ = \mathcal{P}_{AB}[x^{(n)}(\mathcal{T})] P_{\text{acc}}^b[x^{(n)}(\mathcal{T}) \rightarrow x^{(o)}(\mathcal{T})] P_{\text{gen}}^b[x^{(n)}(\mathcal{T}) \rightarrow x^{(o)}(\mathcal{T})] \end{aligned} \quad (1.55)$$

By substituting Eqs. (1.51) and (1.53) into Eq. (1.55), we obtain

$$\begin{aligned} \frac{P_{\text{acc}}^f[x^{(o)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})]}{P_{\text{acc}}^b[x^{(n)}(\mathcal{T}) \rightarrow x^{(o)}(\mathcal{T})]} \\ = h_A[x_0^{(n)}] h_B[x_{\mathcal{T}}^{(n)}] \times \frac{\rho[x_0^{(n)}]}{\rho[x_0^{(o)}]} \times \prod_{i=0}^{\delta t/\Delta t - 1} \frac{\bar{p}[x_{(i+1)\Delta t}^{(o)} \rightarrow x_{i\Delta t}^{(o)}]}{p[x_{i\Delta t}^{(o)} \rightarrow x_{(i+1)\Delta t}^{(o)}]} \end{aligned} \quad (1.56)$$

where many terms have canceled since $(\mathcal{T} - \delta t)/\Delta t$ time slices of the new and the old path are identical (but have different indices). Further cancellation is due to the identity of path probability of the new path and the generation probability associated with the newly grown part of the new path.

As is the case of shooting moves, this ratio simplifies considerably if $p(x \rightarrow y)$ and $\bar{p}(y \rightarrow x)$ are related by the microscopic reversibility criterion from Eq. (1.26). In this case, satisfactory acceptance probabilities for forward and backward shifting moves are

$$P_{\text{acc}}^f[x^{(o)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})] = h_A[x_0^{(n)}]h_B[x_{\mathcal{T}}^{(n)}] \min \left\{ 1, \frac{\rho[x_0^{(n)}]}{\rho[x_0^{(o)}]} \times \frac{\rho_{\text{st}}[x_0^{(o)}]}{\rho_{\text{st}}[x_0^{(n)}]} \right\}, \quad (1.57)$$

and

$$P_{\text{acc}}^b[x^{(n)}(\mathcal{T}) \rightarrow x^{(o)}(\mathcal{T})] = h_A[x_0^{(o)}]h_B[x_{\mathcal{T}}^{(o)}] \min \left\{ 1, \frac{\rho[x_0^{(o)}]}{\rho[x_0^{(n)}]} \times \frac{\rho_{\text{st}}[x_0^{(n)}]}{\rho_{\text{st}}[x_0^{(o)}]} \right\}, \quad (1.58)$$

where we have used the fact that $x_{\delta t}^{(o)} = x_0^{(n)}$. If, further, the distribution of initial conditions $\rho(x)$ is identical to the stationary distribution $\rho_{\text{st}}(x)$, Eqs. (1.57) and (1.58) become

$$P_{\text{acc}}^f[x^{(o)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})] = h_A[x_0^{(n)}]h_B[x_{\mathcal{T}}^{(n)}], \quad (1.59)$$

and

$$P_{\text{acc}}^b[x^{(n)}(\mathcal{T}) \rightarrow x^{(o)}(\mathcal{T})] = h_A[x_0^{(o)}]h_B[x_{\mathcal{T}}^{(o)}] \quad (1.60)$$

From these results, it follows that shifting should be accepted provided the trial path $x^{(n)}(\mathcal{T})$ connects states A and B . The average acceptance probability of a shifting move can be easily controlled by changing the magnitude of the time interval δt .

Short shifting moves are computationally inexpensive and facilitate convergence of quantities averaged over the transition path ensemble. However, shifting moves are essentially orthogonal to shooting moves; alone, small shifting moves do not change the portions of trajectories near the transition state region. Therefore shifting moves do not sample trajectory space ergodically and must be used in conjunction with other trial moves, such as shooting moves.

Shifting moves are slightly different for stochastic and deterministic dynamics. First, we discuss the application of the shifting algorithm to systems evolving deterministically and then discuss shifting moves for stochastic dynamics.

1. Shifting Moves for Deterministic Dynamics

In a forward shifting move of length δt for deterministic dynamics, the new trajectory segment is generated by integrating the equations of motion of the system for $\delta t/\Delta t$ time steps starting from the endpoint $x_{\mathcal{T}}^{(0)}$ of the old path. Similarly, the backward shifting move is carried out by integrating the equations of motion for $\delta t/\Delta t$ time steps backward in time starting from $x_0^{(0)}$. Accordingly, the single step generation probabilities for the forward and the backward move are

$$p(x \rightarrow y) = \delta[y - \phi_{\Delta t}(x)] \quad (1.61)$$

and

$$\bar{p}(y \rightarrow x) = \delta[x - \phi_{-\Delta t}(y)] \quad (1.62)$$

respectively. As shown in Section III.A.1, the generation probabilities $p(x \rightarrow y)$ and $\bar{p}(y \rightarrow x)$ associated with this type of procedure satisfy the reversibility condition expressed in Eq. (1.26). In equilibrium, the acceptance probability for a shifting move for deterministic trajectories is therefore simply given by Eqs. (1.59) and (1.60), that is, the new pathway obtained by shifting the old one in time can be accepted whenever it connects A with B . The shifting algorithm is summarized in Scheme 1.4.

Shifting moves for deterministic trajectories are particularly inexpensive when carried out in long uninterrupted sequences. In this case, the pathway is constrained to a one-dimensional (1D) manifold defined by the equations

Forward Shifting

1. Randomly select a time interval δt from a distribution $w(\delta t)$.
2. Copy the $(\mathcal{T} - \delta t)/\Delta t$ last time slices of the old path to the first $(\mathcal{T} - \delta t)/\Delta t$ time slices of the new path, that is, $x_{i\Delta t + \delta t}^{(n)} = x_{i\Delta t}^{(0)}$ for $i = 0, \dots, (\mathcal{T} - \delta t)/\Delta t$.
3. Integrate the equations of motion forward for $\delta t/\Delta t$ time steps starting from $x_{\mathcal{T} - \delta t}^{(n)}$.
4. Accept the new path $x^{(n)}(\mathcal{T})$ if it is reactive and reject it otherwise.

Backward Shifting

1. Randomly select a time interval δt from a distribution $w(\delta t)$.
2. Copy the $(\mathcal{T} - \delta t)/\Delta t$ first time slices of the old path to the last $(\mathcal{T} - \delta t)/\Delta t$ time slices of the new path, that is, $x_{i\Delta t}^{(n)} = x_{i\Delta t + \delta t}^{(0)}$ for $i = 0, \dots, (\mathcal{T} - \delta t)/\Delta t$.
3. Integrate the equations of motion backward for $\delta t/\Delta t$ time steps starting from $x_{\delta t}^{(n)}$.
4. Accept the new path $x^{(n)}(\mathcal{T})$ if it is reactive and reject it otherwise.

Scheme 1.4. Shifting algorithm for deterministic trajectories.

of motion. By storing states along this manifold as they are obtained, shifting moves can eventually be performed with no cost, since the time evolution of new trajectory segments is already known.

If no stationary distribution $\rho_{\text{st}}(x)$ exists or if it is unknown, a different approach must be taken. For this purpose, we return to Eq. (1.56) and insert the forward and backward generation probabilities from Eqs. (1.61) and (1.62). Using Eq. (1.32) and applying the chain rule for products of Jacobians, we finally obtain the following expression for the acceptance probability of a forward shifting move,

$$P_{\text{acc}}^{\text{f}}[x^{(0)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})] = h_A[x_0^{(n)}]h_B[x_{\mathcal{T}}^{(n)}] \min \left\{ 1, \frac{\rho[x_0^{(n)}]}{\rho[x_0^{(0)}]} \times \left| \frac{\partial x_{\delta t}^{(n)}}{\partial x_0^{(0)}} \right| \right\} \quad (1.63)$$

For a new trajectory $x^{(n)}(\mathcal{T})$ generated with a backward shifting move from the old trajectory $x^{(0)}(\mathcal{T})$, the acceptance probability is slightly different. By using analogous manipulations one obtains

$$P_{\text{acc}}^{\text{b}}[x^{(0)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})] = h_A[x_0^{(n)}]h_B[x_{\mathcal{T}}^{(n)}] \min \left\{ 1, \frac{\rho[x_0^{(n)}]}{\rho[x_0^{(0)}]} \times \left| \frac{\partial x_{\delta t}^{(n)}}{\partial x_0^{(0)}} \right|^{-1} \right\} \quad (1.64)$$

If the dynamics is phase-space volume conserving, that is, $|\partial x_t / \partial x_0| = 1$, such as for Newtonian dynamics, the acceptance probability for both forward and backward shifting moves reduces to

$$P_{\text{acc}}[x^{(0)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})] = h_A[x_0^{(n)}]h_B[x_{\mathcal{T}}^{(n)}] \min \left\{ 1, \frac{\rho[x_0^{(n)}]}{\rho[x_0^{(0)}]} \right\}. \quad (1.65)$$

In this case, a new trajectory obtained with a shifting move is accepted with a probability of $\min\{1, \rho[x_0^{(n)}]/\rho[x_0^{(0)}]\}$ provided the new trajectory connects A with B .

2. Shifting Moves for Stochastic Dynamics

Analogous shifting moves can be carried out also for stochastic trajectories. In a forward reptation move for stochastic trajectories time slices removed from one end of the path are regrown on the other using the integration rules of the underlying dynamics. In a backward reptation move, on the other hand, a new path segment is grown backward using the same procedures used to generate the backward trajectory of a shooting move. As was shown in Section III.A, generation probabilities associated with these

Forward Shifting

1. Randomly select a time interval δt from a distribution $w(\delta t)$.
2. Copy the $(\mathcal{T} - \delta t)/\Delta t$ last time slices of the old path to the first $(\mathcal{T} - \delta t)/\Delta t$ time slices of the new path, that is, $x_{i\Delta t + \delta t}^{(n)} = x_{i\Delta t}^{(o)}$ for $i = 0, \dots, (\mathcal{T} - \delta t)/\Delta t$.
3. Generate $\delta t/\Delta t$ new time steps starting from $x_{\mathcal{T} - \delta t}^{(n)}$ using the propagation rule of the underlying stochastic dynamics.
4. Accept the new path $x^{(n)}(\mathcal{T})$ if it is reactive and reject it otherwise.

Backward Shifting

1. Randomly select a time interval δt from a distribution $w(\delta t)$.
2. Copy the $(\mathcal{T} - \delta t)/\Delta t$ first time slices of the old path to the last $(\mathcal{T} - \delta t)/\Delta t$ time slices of the new path, that is, $x_{i\Delta t}^{(n)} = x_{i\Delta t + \delta t}^{(o)}$ for $i = 0, \dots, (\mathcal{T} - \delta t)/\Delta t$.
3. Invert the momenta belonging to $x_{\delta t}^{(n)}$.
4. Generate $\delta t/\Delta t$ time steps backward starting from $\bar{x}_{\delta t}^{(n)}$ by applying the propagation rule corresponding to the underlying stochastic dynamics.
5. Invert the momenta in the newly generated path segment.
6. Accept the new path $x^{(n)}(\mathcal{T})$ if it is reactive and reject it otherwise.

Scheme 1.5. Shifting algorithm for stochastic trajectories.

procedures satisfy the reversibility condition from Eq. (1.26). Therefore, the acceptance probability from Eqs. (1.59) and (1.60) can be used for stochastic reptation moves based on Langevin and Monte Carlo dynamics. According to this acceptance probability a shifting move for stochastic trajectories is always accepted if both the starting point and the end point of the new path lie in their respective stable states. The reptation algorithm for stochastic trajectories is summarized in Scheme 1.5.

C. Memory Requirements

Molecular dynamics simulations proceed in discrete time steps that are comparable to the shortest characteristic time of atomic motions, often ~ 1 fs. A 1-ps pathway is thus represented by $\sim 10^3$ microstates. Storing such pathways in computer memory quickly becomes costly for large systems. Fortunately, it is not necessary for the purposes of transition path sampling to store every microstate belonging to a given pathway. The algorithms we have described are exact even when a path is represented by just a few states. Clearly, though, a shooting move can only be initiated from a microstate that has been stored. For efficient sampling, it is thus advantageous to store states with a time resolution that captures typical fluctuations of the reaction coordinate. In our experience with molecular liquids, storage of states is necessary only about every 10 fs leading to considerable memory savings. Note that such a reduction of memory requirements is not

possible for other algorithms such as the local algorithm or the dynamical algorithm discussed in Section III.F. In these cases, it is necessary to maintain full trajectories in memory at all times.

For some systems, storing even just tens of microstates can be burdensome. For example, in the density functional theory based method of Car and Parrinello [36] the occupied single-particle Kohn–Sham orbitals $\{\psi_i\}$ are propagated in time together with the nuclear positions using a set of fictitious equations of motion. Shooting moves to harvest Car–Parrinello trajectories can be carried out by selecting a time slice along an existing path and changing only the nuclear momenta before computing a new trajectory. Then, the new path is accepted according to the criteria described in the preceding paragraphs. Due to the large amount of memory necessary to store the Kohn–Sham orbitals and their time derivatives, it is, however, not possible to keep a complete pathway consisting of typically 100 copies of the system in memory with current computational resources. Therefore, a few time slices for shooting are selected randomly before calculating the trajectory and a complete set of data including the Kohn–Sham orbitals is stored only at these specific time slices [15]. Shooting moves can then be initiated from the preselected states. In the unlikely case that shooting moves are rejected from each of these states, one regenerates the current path, storing a complete description of the system at a new set of randomly selected times along the path. Such a technique has been used to harvest Car–Parrinello trajectories for proton transfer in the protonated water trimer [15] and for autoionization in liquid water [24].

D. Stochastic Trajectories as Sequences of Random Numbers

In order to use shooting moves with the simple acceptance probability in Eq. (1.29), one must be able to generate forward and backward trajectories that satisfy microscopic reversibility. So far, we have considered several dynamical systems for which this task is straightforward. But there are many systems for which we do not know how to construct time-reversed trajectory segments that satisfy Eq. (1.26). For example, a stochastic system driven by a nonconservative force may converge to a steady state, but the corresponding stationary distribution generally cannot be written down. This class of dynamical phenomena, which may exhibit rare but important fluctuations [50, 51], are in principle amenable to transition path sampling. But a different type of trial move is needed to implement the method. Crooks and Chandler [26] devised such a move.

The basic idea of the Crooks–Chandler algorithm described in [26] is to represent a pathway by the sequence of random numbers, the noise history, used to generate it. In this representation, local changes of the noise history

can be employed to generate nonlocal changes of trajectories. To see how this can be achieved, consider a system evolving according to Langevin's equation of motion in the high-friction limit:

$$\dot{r} = \mathcal{F}(r, t) + \mathcal{R}(t) \quad (1.66)$$

where r represents the configuration of the system and \mathcal{R} is a stochastic force with correlations $\langle \mathcal{R}(0)\mathcal{R}(t) \rangle = \varepsilon\delta(t)$. The system is driven out of equilibrium by a systematic force $\mathcal{F}(r, t)$ that is either explicitly time-dependent or nonconservative. The above equations of motion can be integrated over short times by representing \mathcal{R} with a set of random numbers ξ [38]. The appropriate distribution for ξ is Gaussian, with zero mean and a variance ε . Because a random number must be drawn for each degree of freedom, ξ is a vector with the dimensionality of configuration space, f . The sequence of random numbers along a trajectory specifies the noise history $\mathcal{R}(t)$. Accordingly, the probability to observe a certain stochastic trajectory $x(\mathcal{T})$ is proportional to the probabilities $p(\xi)$ of all vectors of random numbers ξ used to generate that trajectory,

$$\mathcal{P}[x_0; \xi(\mathcal{T})] = \rho(x_0) \prod_{i=0}^{\mathcal{T}/\Delta t - 1} p(\xi_{i\Delta t}) = \rho(x_0) \prod_{i=0}^{\mathcal{T}/\Delta t - 1} \frac{1}{(2\pi\varepsilon)^{f/2}} \exp\{-|\xi_{i\Delta t}|^2/2\varepsilon\} \quad (1.67)$$

Here, $|\xi_t|$ is the magnitude of the vector ξ drawn for time t .

This representation of the dynamical path probability emphasizes specific realizations of the random force \mathcal{R} . It enables an interesting class of trial moves in trajectory space: A new path $x^{(n)}(\mathcal{T})$ is obtained from an old path $x^{(0)}(\mathcal{T})$ by replacing the noise $\xi_t^{(0)}$ at one randomly chosen time t' by a new set of random numbers $\xi_t^{(n)}$ drawn from the distribution $p(\xi)$. The noise histories of the new and the old path are identical except at time t' , and the new path coincides with the old one from time 0 to time t' . From time $t' + \Delta t$ to time \mathcal{T} , the new path is obtained by integrating the equation of motion with the modified noise history. All new time slices after time t' will differ from those of the old path. A local modification in noise space thus generates a small but global move in trajectory space.

Due to the dissipative nature of these dynamics, trajectories initiated from two nearby points converge rapidly if they possess the same noise history. As a consequence, a new trajectory generated from a reactive old trajectory has a high probability to be reactive as well. Such moves in noise space therefore have a high acceptance probability. This algorithm is summarized in Scheme 1.6.

The generation probability for trial paths obtained from this type of move is

$$P_{\text{gen}}[x^{(0)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})] = p(\xi_{t'}^{(n)}) \quad (1.68)$$

1. Randomly select a time slice $x_{t'}^{(0)}$ on the existing trajectory $x^{(0)}(\mathcal{T})$.
2. Replace the set of random numbers $\xi_{t'}^{(0)}$ belonging to time slice $x_{t'}^{(0)}$ by a new set of random numbers $\xi_{t'}^{(n)}$ drawn from a Gaussian distribution with the appropriate mean and variance.
3. Determine the new path $x^{(n)}(\mathcal{T})$ from $t' + \Delta t$ to \mathcal{T} using the modified set of random numbers to integrate the equations of motion.
4. Accept the new path if the final point $x_{\mathcal{T}}^{(n)}$ is in region B .

Scheme 1.6. Crooks–Chandler noise history algorithm for stochastic trajectories.

where $\xi_{t'}^{(n)}$ is the new random number at time t' . By inserting this generation probability and the path probability from Eq. (1.67) into Eq. (1.21), we obtain a corresponding acceptance probability

$$P_{\text{acc}}[x^{(0)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})] = h_A[x_0^{(n)}]h_B[x_{\mathcal{T}}^{(n)}] \times \min \left\{ 1, \frac{\rho[x_0^{(n)}] \prod_{i=0}^{\mathcal{T}/\Delta t - 1} p[\xi_{i\Delta t}^{(n)}]}{\rho[x_0^{(0)}] \prod_{i=0}^{\mathcal{T}/\Delta t - 1} p[\xi_{i\Delta t}^{(0)}]} \times \frac{p[\xi_{t'}^{(0)}]}{p[\xi_{t'}^{(n)}]} \right\} \quad (1.69)$$

Most factors in the second argument of the min-function cancel because the noise histories of the new and the old trajectories differ only at time t' . Because such a move does not modify the initial time slice of the path [$x_0^{(n)} = x_0^{(0)}$], Eq. (1.69) simplifies to

$$P_{\text{acc}}[x^{(0)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})] = h_B[x_{\mathcal{T}}^{(n)}] \quad (1.70)$$

Thus, a new path obtained by applying local changes in noise space is accepted if it ends in region B .

As described, this algorithm harvests an ensemble of trajectories with identical initial states x_0 , and is therefore incomplete. But sampling the distribution of initial conditions is in fact very simple for the system considered, since $\rho(x) = \rho_{\text{st}}(x)$. One need only apply dynamical propagation rules [with a randomly selected realization of $\mathcal{R}(t)$] to an existing initial state. States along the resulting trajectory are correctly weighted by $\rho(x)$. In order to make the algorithm complete, these new initial states are occasionally used to grow trial trajectories with an existing noise history. Again, the dissipative nature of these dynamics ensures that such a trial trajectory has a reasonable chance of exhibiting the rare fluctuation of interest. If the path does satisfy the desired boundary conditions, it is accepted with unit probability. With this complete algorithm, Crooks and Chandler [26] have successfully sampled large fluctuations in nonequilibrium systems that were previously inaccessible.

E. Other Algorithms

Preceding sections have described the most efficient algorithms we have devised to sample transition paths in complex systems. Other algorithms, while sound in principle, are far less efficient than shooting and shifting. For completeness, we discuss these less efficient algorithms briefly.

1. Local Algorithm for Stochastic Pathways

It is possible to sample stochastic transition pathways by making only local displacements of trajectories. For example, a randomly chosen time slice $x_{t'}^{(0)}$ of an existing pathway may be modified by adding a small displacement δx to positions and momenta, $x_{t'}^{(n)} = x_{t'}^{(0)} + \delta x$. All other time slices remain unchanged. This modification, which is local in time, gives a different but finite path probability $\mathcal{P}_{AB}[x(\mathcal{T})]$. If the displacement δx is chosen from a symmetric distribution, $w(\delta x) = w(-\delta x)$, then the generation probability for the new path is identical to that for the reverse move. According to Eq. (1.21), the new path should be accepted with probability

$$P_{\text{acc}}[x^{(0)}(\mathcal{T}) \rightarrow x^{(n)}(\mathcal{T})] = \min \left\{ 1, \frac{p[x_{t'-\Delta t}^{(0)} \rightarrow x_{t'}^{(n)}]p[x_{t'}^{(n)} \rightarrow x_{t'+\Delta t}^{(0)}]}{p[x_{t'-\Delta t}^{(0)} \rightarrow x_{t'}^{(0)}]p[x_{t'}^{(0)} \rightarrow x_{t'+\Delta t}^{(0)}} \right\} \quad (1.71)$$

where we have assumed that an intermediate time slice is modified, $0 < t' < \mathcal{T}$. Acceptance probabilities for displacing initial and final time slices are somewhat different [5].

The local algorithm suffers from several shortcomings. The most serious of these is its inefficiency. In contrast to shooting and shifting moves, local moves sample phase space with a rate proportional to L^3 , where L is the number of time slices in a path [52]. This unfavorable scaling makes the local algorithm impractical for all but the simplest problems. Furthermore, all time steps of a path must be stored in computer memory, a burdensome task for high-dimensional systems. Finally, the local algorithm cannot be applied to deterministic trajectories. To date, we have only been able to sample deterministic trajectories with shooting and shifting moves.

A local algorithm for Metropolis Monte Carlo trajectories must be constructed carefully. Due to the finite probability of exactly repeated states in these paths, the corresponding transition probability includes a singular term [see Eq. (1.14)]. The generation algorithm for local path moves must take this singularity into account properly. Appropriate acceptance probabilities are given in [5]. (H. C. Andersen has drawn our attention to an omission in [5]. In Metropolis Monte Carlo trajectories sequences of multiple rejections can occur. Attempts to modify time slices in the interior

of such a sequence always lead to rejections. More specifically, Eqs. (17) and (18) of [5] must be modified to include sequences of the form $r_{\tau-1} = r_{\tau} = r_{\tau+1}$ (in the notation of [5]). In that case, the acceptance probability for a local move from r_{τ} to r'_{τ} vanishes, that is, $P_{\text{acc}}[r_{\tau} \rightarrow r'_{\tau}] = 0$.

2. Dynamical Algorithm

For stochastic pathways, such as those generated according to Langevin's equation of motion, the path probability functional $\mathcal{P}_{AB}[x(\mathcal{T})]$ can be written as

$$\mathcal{P}_{AB}[x(\mathcal{T})] \equiv \exp(-S_{AB}[x(\mathcal{T})]), \quad (1.72)$$

where the functional $S_{AB}[x(\mathcal{T})]$ depending on all path coordinates $x(\mathcal{T})$ is called the path ‘‘action’’. The form of Eq. (1.72), similar to the canonical distribution function $\exp[-\beta V(r)]$, suggests that we sample the path distribution $\mathcal{P}_{AB}[x(\mathcal{T})]$ with dynamical methods based on generalizations of Newton's equation of motion. For this purpose, path space is extended to include also ‘‘momenta’’ $y(\mathcal{T}) \equiv M\dot{x}(\mathcal{T})$ of the path coordinates $x(\mathcal{T})$. Here, M is an artificial mass associated with the path coordinates and the dot indicates a derivative with respect to an artificial time θ . This new space $\{x(\mathcal{T}), y(\mathcal{T})\}$ has twice the dimensionality of the original path space. One then defines a path Hamiltonian functional $H_P[x(\mathcal{T}), y(\mathcal{T})]$ by adding a ‘‘kinetic energy’’ to the path action,

$$H_P[x(\mathcal{T}), y(\mathcal{T})] \equiv \sum_{\alpha} \frac{y_{\alpha}^2}{2M} + S[x(\mathcal{T})] \quad (1.73)$$

From this Hamiltonian functional one can derive a set of equations of motion capable of moving the path $x(\mathcal{T})$ through path space in artificial time θ ,

$$\begin{aligned} \frac{dx_t}{d\theta} &= \frac{\partial H_P}{\partial y_t} = \frac{y_t}{M} \\ \frac{dy_t}{d\theta} &= -\frac{\partial H_P}{\partial x_t} = -\frac{\partial S_{AB}[x(\mathcal{T})]}{\partial x_t} \end{aligned} \quad (1.74)$$

If these equations are equipped with an appropriate thermostat, for example, a stochastic Andersen thermostat [53] or a deterministic Nosé–Hoover thermostat [37], the resulting distribution in path space is consistent with $\exp(-S_{AB}[x(\mathcal{T})]) = \mathcal{P}_{AB}[x(\mathcal{T})]$. Alternatively, a stochastic Langevin equation may be employed to drive the path through path space.

The forces necessary to integrate the artificial equations of motion are obtained by differentiating $S_{AB}[x(\mathcal{T})] = -\ln \mathcal{P}_{AB}[x(\mathcal{T})]$ with respect to the path coordinates x_t . Since $\mathcal{P}_{AB}[x(\mathcal{T})]$ is a product of transition probabilities, the distribution of initial conditions and the boundary conditions constraining the path to start in A and end in B , the path action $S_{AB}[x(\mathcal{T})]$ is a sum of terms originating from the different contributions to the path probability. The boundary conditions $h_A(x_0)$ and $h_B(x_{\mathcal{T}})$ act as hard walls confining the endpoints of the path to regions A and B , respectively. These hard walls require special attention when integrating the equations of motion (1.74).

Although correct in principle, dynamical path sampling algorithms are far inferior to shifting and shooting algorithms in terms of efficiency. One reason for this deficiency is that due to the strong coupling between subsequent time slices especially for low friction very small time steps must be used in order to reproduce the associated high-frequency oscillations in the path dynamics correctly. Thus, path motion in trajectory space proceeds slowly. The sampling rate is further decreased by the necessity of computing second derivatives of the potential in order to determine forces acting on the path variables. Such calculations can be onerous especially in the case of first principles molecular dynamics simulations. For these reasons dynamical algorithms have not been applied to processes occurring in complex systems so far.

3. *Configurational Bias Monte Carlo*

In its original form, the configurational bias Monte Carlo algorithm is a method to sample polymers in the melt in an efficient way [54]. The basic idea of the algorithm is to regrow entire polymers in a biased fashion in order to avoid unfavorable overlaps. Due to the formal similarity between polymers and stochastic pathways a configurational bias Monte Carlo scheme can be constructed to sample transition pathways [5]. To date, this algorithm has been only used to study the kinetics of hydrogen bonds in liquid water [8].

Also, algorithms familiar from path integral simulations [52], such as the staging algorithm [55], can be used to sample the transition path ensemble for stochastic trajectories. These algorithms are, however, less efficient than the shooting and shifting algorithms, which are the most efficient transition path sampling algorithms found so far.

F. **Parallel Tempering**

The transition path ensemble includes all pathways that make successful transitions on the time scale of the path length \mathcal{T} . For complicated energy

surfaces, however, it may be difficult to sample all statistically relevant pathways. For example, proton transfer occurs in the protonated water trimer via two different classes of pathways [15]. Switching from one class of transition pathways to the other is hindered by a high energetic barrier. One might expect similar problems in sampling protein folding pathways for physiologically relevant temperatures. Such ergodicity problems are akin to those encountered in Monte Carlo simulations of supercooled liquids and glasses, where locally stable states are separated by high energy barriers. Various techniques have been developed to overcome these problems in the context of conventional Monte Carlo simulations, including J-walking [56], multicanonical sampling [57] and parallel tempering [58–60]. In principle, all these methods can be generalized to improve sampling of the transition path ensemble. Vlugt and Smit [23] recently showed that parallel tempering is especially suitable for transition path sampling.

In a conventional parallel tempering calculation, several simulations are performed in parallel at different temperatures. One allows exchange of configurations between different temperature levels with a probability satisfying detailed balance. At low temperatures, the system is confined to the vicinity of a local energy minimum, but at high temperatures the system can surmount energetic barriers readily. With parallel tempering, the system is able to move efficiently through configuration space, while sampling the statistically most important, low energy regions with correct probability.

Parallel tempering can be used to sample transition pathways analogously. Pathways are harvested in parallel at N different temperatures T_1, T_2, \dots, T_N . Shooting and shifting moves are complemented with exchange of pathways between different temperature levels. Appropriate acceptance rules satisfying detailed balance can be easily derived [23]. Parallel tempering not only improves sampling efficiency at each temperature level, but also provides temperature-dependent results at basically no extra cost. The structure of this algorithm is ideal for implementation on massively parallel computers.

To test parallel tempering in conjunction with transition path sampling, we have applied it to a simple toy model. In this two dimensional (2D) system (see Fig. 1.10), a “molecule” is immersed in a fluid of purely repulsive disks. The molecule consists of three atoms bonded with harmonic springs that repel each other at short distance. All atoms have the same mass and size and the system evolves according to Newton’s equation of motion. Initial conditions are weighted by a canonical distribution. The triatomic molecule can reside in two stable states: state A , in which atoms 1, 2, and 3 are arranged in a clockwise manner (as in Fig. 1.4), and state B , in which the arrangement is counterclockwise. The two stable states are distinguished

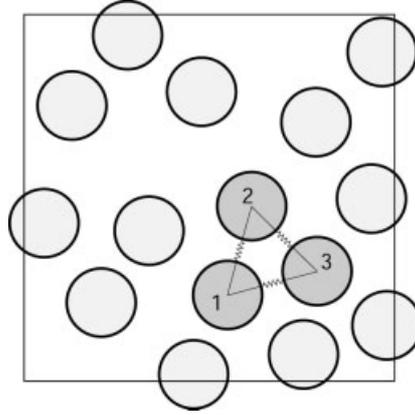


Figure 1.10. A molecule (dark gray) of three 2D atoms with short-range repulsion and connected with harmonic springs is immersed in a fluid of purely repulsive disks. The molecule can reside in two long lived enantiomers.

with the order parameter $q \equiv (y_{12}x_{13} - x_{12}y_{13})/r_{12}$, where $\mathbf{r}_{13} \equiv (x_3 - x_1, y_3 - y_1)$ and $\mathbf{r}_{12} \equiv (x_2 - x_1, y_2 - y_1)$ are the vectors from atom 1 to atom 2 and from atom 1 to atom 3, respectively. The order parameter q measures the position of particle 3 with respect to particles 1 and 2.

Fluid atoms act as a heat bath for the triatomic molecule, providing activation energy for transitions between states A and B and dissipation to stabilize the molecule following a transition. Three different pathways connect states A and B . In each of them, a different atom of the molecule squeezes its way through the gap formed by the other two atoms. In each of the transition states, the three atoms are collinear with a different atom at the center.

We have studied transitions between states A and B with transition path sampling in the canonical ensemble. For the parameters used in the simulation, the energetic barrier separating states A and B has a height of $\sim 22\epsilon$ in Lennard–Jones units. We used a total of 16 particles (including the atoms of the molecule) with periodic boundary conditions. If the sampling is ergodic, that is, if all relevant regions in the space of all transitions paths are visited, one expects that the three classes of pathways occur with the same frequency. However, such ergodic sampling is observed only at temperatures larger than $k_B T \simeq 1.0\epsilon$. For lower temperatures, the sampling is effectively confined to one of the three transition states, even for simulations of $> 10^6$ cycles. Results of these simulations are shown in Figure 1.11.

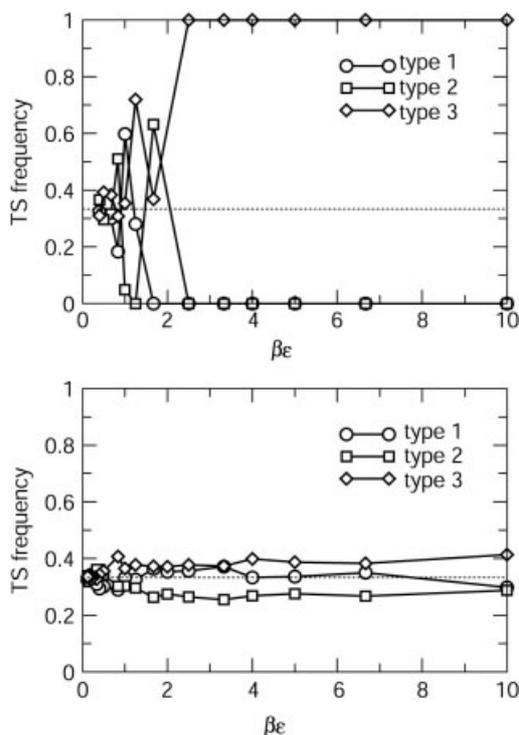


Figure 1.11. Fraction of transition pathways of type 1, 2, and 3 found in the simulations as a function of $\beta\epsilon = \epsilon/k_B T$. The panel on the top shows the results obtained without parallel tempering. The bottom panel depicts the results of parallel tempering simulations with 20 temperature levels.

These results demonstrate that parallel tempering can enhance sampling considerably. Without parallel tempering severe sampling problems occur for $\beta \equiv 1/k_B T > 1/\epsilon$. With parallel tempering, all path space regions are visited with the correct probability.

G. Generating an Initial Path

The transition path sampling techniques we have described assume that an initial reactive pathway is available. Generating such a pathway is therefore an important step in applying the method. In the simplest cases, a trajectory connecting A and B can be obtained by running a long molecular dynamics (or stochastic dynamics) simulation. For most applications, however, the

rarity of transitions rules out this straightforward approach, and an initial trajectory must be procured by other means.

The path used to initiate sampling need not be representative of the transition path ensemble. This fact is a particular strength of our methodology. In fact, the initial path need not even be a true dynamical pathway. Any chain of states connecting A and B is in principle sufficient to begin transition path sampling. Subsequent importance sampling will naturally guide pathways to the most important regions of trajectory space. This “equilibration” is directly analogous to the initial relaxation of conventional Monte Carlo simulations. Suppose, for example, one wants to harvest configurations of an atomic liquid. If a typical state of the liquid is unavailable, one may begin by placing atoms on the sites of a periodic lattice. Such a regular arrangement is clearly atypical of the liquid phase. But importance sampling will quickly carry the fluid from this unlikely configuration to more representative states. In a similar way, transition path sampling efficiently carries reactive pathways from unlikely regions of trajectory space to those representative of natural transitions.

An artificial chain of states connecting basins of attraction can be constructed in various ways. In fact, the many methods that focus on nondynamical transition paths may be useful for this task [28, 30, 31]. The more closely an initial pathway resembles natural reactive trajectories, the more quickly transition path sampling will reach the important regions of trajectory space. We describe here a few useful techniques for generating reactive chains of states.

One systematic way to generate an initial pathway is to gradually transform an ensemble of all trajectories starting in A to one including only those starting in A and ending in B . As discussed in Section IV, such a conversion of path ensembles is computationally demanding. But if the transformation is performed quasistatically, the resulting pathway will be truly representative of the transition path ensemble. In this case, no equilibration of the initial path is needed.

High-temperature pathways can also be used to initiate transition path sampling. Consider, for example, the folding of a protein. At physiological temperatures, folding and unfolding occur very rarely on a molecular timescale. But at higher temperature, unfolding is sufficiently facile that it can be simulated with straightforward molecular dynamics [61]. Such a high-temperature trajectory can be used to begin sampling protein folding pathways at more relevant temperatures. In some cases, high-temperature transition pathways are qualitatively different from those at lower temperatures. It may then be necessary to carry out a systematic cooling procedure, in which the ensemble of pathways is brought to lower temperatures in small steps.

Finally, a presumed notion of a reaction coordinate can sometimes be used to acquire an reactive pathway. By controlling such a variable, it may be possible to drive the system from one stable state to another. If this procedure is successful, shooting and shifting moves can be applied to the resulting chain of states.

In our experience, no general recipe for generating initial trajectories is ideally suited to all situations. It seems best to proceed by experimenting with various schemes. In this process, one often gains useful insight about the relevant basins of attraction and the coordinates that characterize them.

H. Transition Path Sampling with an Existing Molecular Dynamics Program

Transition path sampling based on the shooting and shifting algorithms can be easily implemented using existing molecular dynamics (MD) programs. One might, for example, desire to conduct transition path sampling studies using a commercial molecular dynamics package for which the source code is unavailable. Such a study is most conveniently done by developing a separate path sampling module and interfacing it with the existing MD program through system calls and input–output to and from disk.

In such combination of an existing MD program with a new path sampling module, a deterministic shooting move, for example, can be carried out as follows. In the path sampling module, a time slice is first selected from the existing path that is stored in memory or on disk. Next, the corresponding momentum is changed according to the rules described in Section III.A. and the modified phase-space point is written to a file from which it can later be read by the MD program. Then, a forward shot is carried out by first writing the desired number of integration steps into the input file of the MD program and then starting the MD program via a system call. The MD program integrates the equations of motion and stores the resulting trajectory segment on disk. Note that it is not necessary to store every time step of the new trajectory. Rather, it is sufficient to store the new trajectory only at a smaller number of regularly spaced points. When the trajectory is completed, control returns to the path sampling module that reads the forward trajectory from disk and verifies whether the last state of the new path, $x_{\mathcal{F}}^{(n)}$, lies within region B . If it does not, the move is rejected. Otherwise, one proceeds with the backward part of the shooting move. For this purpose, one inverts the momenta of the time slice from which the forward shot was initiated and the modified time slice is again written to a file. Then, the number of backward integration steps is specified in the appropriate input file and the MD program is again started with a system call. After that the newly generated backward trajectory segment is read by the path

1. Randomly select time slice $x_r^{(0)}$ for shooting.
2. Modify momenta and write time slice $x_r^{(n)}$ to disk.
3. Write number of integration time steps for forward shot into input file of MD-code.
4. Start MD-code with system call.
5. Read newly generated forward trajectory segment from disk.
6. Reject move if new endpoint $x_s^{(n)}$ not in B and proceed with next point otherwise.
7. Invert momenta of time slice $x_r^{(n)}$ and write modified time slice to disk.
8. Write number of integration time steps for backward shot into input file of MD-code.
9. Start MD-code with system call.
10. Read newly generated backward trajectory segment from disk and invert momenta.
11. If new initial point $x_0^{(n)}$ is in A , the path is accepted, and rejected otherwise.

Scheme 1.7. Shooting algorithm with existing molecular dynamics code.

sampling module and the momenta along this segment of the trajectory are inverted such that they point forward in the whole path. The shooting move is completed by accepting it if the new initial point, $x_0^{(n)}$, is in region A and rejecting it otherwise. A summary of the algorithm for shooting moves in a path sampling interface is given in Scheme 1.7.

In a shifting move, the path sampling code is interfaced with the existing MD program in an analogous way. A Fortran code originally developed as a path sampling interface to the Car–Parrinello molecular dynamics code CPMD [62] can be downloaded from the web site <http://www.pathsampling.org>. Modifications necessary to use this interface in combination with other molecular dynamics packages are straightforward.

IV. COMPUTING RATES AND REVERSIBLE WORK

Because trajectories harvested by transition path sampling are true dynamical pathways, they can be used to compute dynamical observables. For transitions between two stable states A and B , the most important observables are rate constants k_{AB} and k_{BA} for forward and backward transitions, respectively. In this section, we show how rate constants can be calculated by, in effect, reversibly changing ensembles of trajectories. We accomplish this transformation by extending the importance sampling techniques discussed in previous sections. The calculation of rate constants with transition path sampling thus does not require an understanding of reaction mechanism, offering a significant advantage over conventional methods.

The standard prescription for computing a rate constant k_{AB} is a two-step procedure, based on the perspective of transition state theory. It requires the

choice of a putative reaction coordinate, q , and its success depends on the validity of that choice. In the first step, one computes the reversible work $W(q)$ to bring the system from stable state A to a surface, $q = q^*$, dividing A and B . Here, q^* is the value of q for which the free energy $W(q)$ is locally maximal, so that $q = q^*$ approximates the transition state surface. From the computed reversible work, the transition state theory estimate of the rate constant, k_{TST} , can be determined for the particular choice of q ,

$$k_{\text{TST}} = \frac{\langle |\dot{q}| \rangle}{2\langle \theta(q^* - q) \rangle} \langle \delta(q - q^*) \rangle = \frac{1}{2} \langle |\dot{q}| \rangle \frac{\exp\{-\beta W(q^*)\}}{\int dq' e^{-\beta W(q')} \theta(q^* - q')} \quad (1.75)$$

Here, \dot{q} is the time derivative of the reaction coordinate q and $\theta(x)$ is the Heaviside step function. The estimate k_{TST} is accurate to the extent that reactive trajectories cross the surface $q = q^*$ only once during a transition. In the second step of the procedure, corrections to transition state theory are computed by initiating many fleeting trajectories from the $q = q^*$ surface. The fates of these trajectories determine the time-dependent transmission coefficient,

$$\kappa(t) = \frac{2}{\langle |\dot{q}| \rangle} \frac{\langle \dot{q}_0 \delta(q_0 - q^*) \theta(q_t - q^*) \rangle}{\langle \delta[q_0 - q^*] \rangle} \quad (1.76)$$

where q_t is the reaction coordinate at time t . Since after a short transient time τ_{mol} fleeting trajectories started from the dividing surface are committed to one of the stable states, $\kappa(t)$ becomes constant in the time range $\tau_{\text{mol}} < t \ll \tau_{\text{rxn}}$, where $\tau_{\text{rxn}} \equiv (k_{AB} + k_{BA})^{-1}$ is the reaction time. This plateau value of $\kappa(t)$ is the transmission coefficient κ . Note that in this formalism regions A and B are adjacent. The plateau value of

$$k(t) = \kappa(t) k_{\text{TST}} \quad (1.77)$$

is a formally exact expression for the (classical) rate constant k_{AB} . The practical utility of this ‘‘reactive flux’’ method depends on the size of κ . For the optimal choice of q , that is, the true reaction coordinate, κ is maximized and is typically near unity if the crossing of the dynamical bottleneck is not diffusive. In this case, hundreds of fleeting trajectories are sufficient to determine κ to within, say, 10%. For a poor choice of q , on the other hand, $\kappa \ll 1$, leading to severe numerical problems. In this case, even thousands of fleeting trajectories are insufficient to distinguish κ from zero. This failure reflects the fact that states with $q = q^*$ are poor approximations of true

transition states. Most trajectories passing through these states are therefore not reactive.

Preceding sections described how reactive trajectories may be harvested without a priori knowledge of a reaction coordinate. In a similar way, it is possible to compute rate constants accurately without this knowledge. In Section IV.A, we begin by identifying the central time correlation function of kinetics as a ratio of partition functions for different ensembles of trajectories. As such, it may be considered as the exponential of a work, specifically the work to reversibly confine the final state of a trajectory to the product state B . The rarity of transitions ensures that this work is large, so it cannot be computed directly. Instead, we borrow and extend well-known techniques from the statistical mechanics of configurations, namely umbrella sampling and thermodynamic integration. With these techniques, and a convenient factorization reminiscent of that in the reactive flux method, rate constants can be calculated with reasonable computational cost even when the underlying mechanism is not well understood. We will conclude this section by discussing the question of how to find the path length t necessary to allow sampling of all important transition pathways. Note that all formulas and algorithms presented in this section are equally valid for stochastic and deterministic dynamics.

A. Population Fluctuations

In a system at equilibrium, the populations of stable states A and B fluctuate in time due to spontaneous transitions between them. The dynamics of transitions are therefore characterized by the correlation of state populations in time

$$C(t) \equiv \frac{\langle h_A(x_0)h_B(x_t) \rangle}{\langle h_A(x_0) \rangle} \quad (1.78)$$

Here, $\langle \dots \rangle$ denotes an average over the equilibrium ensemble of initial conditions. $C(t)$ is the conditional probability to find the system in state B at time t provided it was in state A at time 0. According to the fluctuation-dissipation theorem [63], dynamics of equilibrium fluctuations are equivalent to the relaxation from a nonequilibrium state in which only state A is populated. At long timescales, these nonequilibrium dynamics are described by the phenomenology of macroscopic kinetics. Thus, the asymptotic behavior of $C(t)$ is determined by rate constants k_{AB} and k_{BA} . At long times, and provided that a single dynamical bottleneck separating A from B causes simple two-state kinetics,

$$C(t) \approx \langle h_B \rangle (1 - \exp\{-t/\tau_{\text{rxn}}\}) \quad (1.79)$$

At short times, $C(t)$ will reflect microscopic motions in the transition state region, which are correlated over a timescale τ_{mol} . τ_{mol} is essentially the time required to cross the dynamical bottleneck separating the stable regions and commit to one of the stable states. Equation (1.79) links the reaction time τ_{rxn} measured experimentally with the microscopic correlation function $C(t)$. At long times,

$$C(t) \approx k_{AB}t \quad (1.80)$$

provided the kinetics of transitions from A to B are two state. The slope of $C(t)$ in this region is the forward reaction rate constant k_{AB} . If metastable states mediate these transitions, the corresponding kinetics will be more complicated. In such cases, sampling trajectories that linger in metastable states may be difficult. In general, sequences of transitions between stable and metastable states should be dissected. Considering component transitions separately facilitates both sampling and analysis of reactive trajectories.

In general, $C(t)$ contains all the information needed to determine kinetic coefficients measured experimentally. In the following sections, we describe the calculation of $C(t)$ using transition path sampling. Once $C(t)$ is determined for times greater than τ_{mol} , rate constants can be extracted by analysis of $C(t)$.

B. Reversible Work

In the context of transition path sampling, the time correlation function in Eq. (1.78) is naturally written in terms of sums over trajectories

$$C(t) = \frac{\int \mathcal{D}x(t) h_A(x_0) \mathcal{P}[x(t)] h_B(x_t)}{\int \mathcal{D}x(t) h_A(x_0) \mathcal{P}[x(t)]} = \frac{Z_{AB}(t)}{Z_A} \quad (1.81)$$

The second equality in Eq. (1.81) follows from the definition of partition functions for ensembles of trajectories [see Eq. (1.3)]. Specifically, $Z_{AB}(t)$ is the partition function for the ensemble of trajectories that begin in region A and end in region B a time t later. Similarly, Z_A is the partition function for the ensemble of all trajectories that begin in region A without any restriction on the endpoint at time t . Z_A is written without time argument because the dynamical weight $\mathcal{P}[x(t)]$ is normalized, so that $Z_A = \int dx_0 \rho(x_0) h_A(x_0)$ is just the equilibrium probability to find the system in state A . While Z_A “counts” all trajectories with initial points in A , Z_{AB} counts only those trajectories which start in A and end in B .

The partition function Z_A transforms into the partition function Z_{AB} if the constraint $h_B(x_t)$ on the endpoint of the path is introduced. We therefore

interpret the ratio of partition functions in Eq. (1.81) as the exponential of a reversible work $W_{AB}(t)$ to change between these two ensembles,

$$W_{AB} \equiv -\ln \frac{Z_{AB}(t)}{Z_A} \quad (1.82)$$

$W_{AB}(t)$ is an effective change in free energy, describing the confinement of trajectory endpoints x_t to region B , while preserving the constraint that initial points x_0 lie in region A . Provided this confinement process is carried out reversibly, the corresponding work is independent of the specific procedure used. In Section IV.C, we describe an advantageous choice for this process that is closely related to the method of thermodynamic integration.

C. Umbrella Sampling

For ensembles of configurations, changes in free energy that are large compared to $k_B T$ are often computed by introducing an artificial bias potential, $U(x)$ [41, 44, 63]. In the simplest cases, $U(x)$ is chosen so that the appropriate importance sampling visits rare but interesting states with the same frequency as typical equilibrium states. Differences in free energy may then be computed directly, and correcting for the presence of $U(x)$ is straightforward. When the relevant phase space is very large, it is generally more efficient to implement umbrella sampling by dividing space into a series of overlapping regions, or “windows” [63]. In this case, the i th bias potential confines the system to the i th window, and ensures uniform sampling within the window. The distribution of states in the entire space is then constructed by requiring that distributions within windows are consistent in overlapping regions. These methods are readily generalized to the sampling of trajectories.

In order to make explicit the correspondence between umbrella sampling of configurations and umbrella sampling of trajectories, we consider a schematic example. Imagine that an order parameter $\lambda(x)$ successfully distinguishes between stable states A and B . For concreteness, suppose that these states are two different conformers of a single molecule. In the first conformational state, λ has a certain range of values, $\lambda_{\min}^A < \lambda < \lambda_{\max}^A$. A distinct range of values characterizes the second conformer, $\lambda_{\min}^B < \lambda < \lambda_{\max}^B$, with $\lambda_{\min}^B > \lambda_{\max}^A$. We could compute the reversible work, w_B , to confine the molecule to the latter conformational state using umbrella sampling in the conventional way. The total range of λ is first divided into narrow windows, $\lambda^{(i)} - \Delta/2 < \lambda < \lambda^{(i+1)} + \Delta/2$. Here, Δ is a small, positive quantity ensuring that adjacent windows overlap. The corresponding bias potential for the i th

window, $U^{(i)}(x)$, is infinite for $\lambda < \lambda^{(i)} - \Delta/2$, as well as for $\lambda > \lambda^{(i+1)} + \Delta/2$ and vanishes inside the window. For a wise choice of the window boundaries $\lambda^{(i)}$, the distribution of λ , $P(\lambda)$, will not vary considerably within each window. Within a window, then, $P(\lambda)$ may be accurately computed up to a constant of proportionality with straightforward importance sampling. These proportionality constants are obtained by demanding continuity in the intervals of overlap between adjacent windows. Finally, w_B is calculated by appropriate integration over the order parameter distribution,

$$w_B = -k_B T \ln \int_{\lambda_{\min}^B}^{\lambda_{\max}^B} d\lambda P(\lambda) \quad (1.83)$$

Umbrella sampling for ensembles of trajectories can be carried out in close analogy to the procedure described above. For the purpose of computing rate constants (in this case the rate of conformational transitions), we focus on computing the reversible work to confine trajectories' endpoints to state B , given that these trajectories begin in state A . $P_A(\lambda, t)$ is defined to be the distribution of λ at the endpoints of trajectories initiated in A ,

$$P_A(\tilde{\lambda}, t) = \frac{\int \mathcal{D}x(t) h_A(x_0) \mathcal{P}[x(t)] \delta[\tilde{\lambda} - \lambda(x_t)]}{Z_A} = \langle \delta[\tilde{\lambda} - \lambda(x_t)] \rangle_A \quad (1.84)$$

Here, $\langle \cdots \rangle_A$ denotes an average over all pathways beginning in state A . As before, we divide the range of λ into narrow, overlapping windows. Each window corresponds to a region $\mathcal{W}[i]$ in phase space. Within each window, $P_A(\lambda, t)$ can be computed accurately up to a constant of proportionality using the methods described in Section III. Specifically, importance sampling is used to harvest trajectories according to the path weight functional

$$\mathcal{P}_{A \mathcal{W}[i]}[x(t)] = \rho(x_0) \mathcal{P}[x(t)] h_A(x_0) h_{\mathcal{W}[i]}(x_t) \quad (1.85)$$

The path ensemble $\mathcal{P}_{A \mathcal{W}[i]}[x(t)]$, consists of pathways starting in A and ending in the window region $\mathcal{W}[i]$.

The distribution of the order parameter λ in region $\mathcal{W}[i]$ is,

$$\begin{aligned} P_{A \mathcal{W}[i]}(\tilde{\lambda}, t) &= \frac{\int \mathcal{D}x(t) h_A(x_0) \mathcal{P}[x(t)] h_{\mathcal{W}[i]}(x_t) \delta[\tilde{\lambda} - \lambda(x_t)]}{\int \mathcal{D}x(t) h_A(x_0) h_{\mathcal{W}[i]}(x_t)} \\ &= \langle \delta[\tilde{\lambda} - \lambda(x_t)] \rangle_{A \mathcal{W}[i]} \end{aligned} \quad (1.86)$$

Since $P_{A \mathcal{W}[i]}$ is the ensemble distribution function for pathways with

endpoints in region $\mathcal{W}[i]$ the distribution $P_{A\mathcal{W}[i]}(\lambda, t)$ can have nonvanishing values only in the order parameter range corresponding to $\mathcal{W}[i]$ and must vanish for all other values of λ . As can be seen by comparing Eq. (1.84) with Eq. (1.86), inside window i the distribution $P_{A\mathcal{W}[i]}(\lambda, t)$ with path endpoints restricted to region $\mathcal{W}[i]$ is proportional to the complete order parameter distribution $P_A(\lambda, t)$ obtained as an average over pathways with unrestricted endpoints. By matching the distributions $P_{A\mathcal{W}[i]}(\lambda, t)$ in the overlapping regions, one can thus obtain the complete distribution $P_A(\lambda, t)$. To reduce the number of required windows it is sometimes convenient to introduce a bias $U(x_t)$ in the sampling. Of course, such a bias must be appropriately corrected for in the calculation of path averages.

By integrating, this distribution function over those values of the order parameter λ belonging to region B one can obtain the probability for a path starting in A to reside in B at a time t later

$$C(t) = \frac{\langle h_A(x_0)h_B(x_t) \rangle}{\langle h_A \rangle} = \int_{\lambda_{\min}^B}^{\lambda_{\max}^B} d\lambda P_A(\lambda, t) \quad (1.87)$$

Accordingly, the reversible work required to confine the endpoint of paths starting in A to region B is,

$$W_{AB}(t) = -\ln \int_{\lambda_{\min}^B}^{\lambda_{\max}^B} d\lambda P_A(\lambda, t) \quad (1.88)$$

where we have used the fact that $P_A(\lambda, t)$ is normalized.

Figure 1.12 shows the results of an umbrella sampling path sampling simulation carried out to determine $C(t)$ for proton transfer in the protonated water trimer described in Section I. The curves in (a) are the distributions $P_{A\mathcal{W}[i]}(\lambda, t)$ restricted to the window regions $\mathcal{W}[i]$. In this case, the order parameter $\lambda(x)$ was defined as the angular difference $\Delta\phi \equiv \phi_2 - \phi_1$, where the angles ϕ_1 and ϕ_2 were defined as indicated in Figure 1.2. Requiring continuity of $P_A(\lambda, t)$, where the regions $\mathcal{W}[i]$ overlap, one obtains the distribution $P_A(\lambda, t)$ shown on a logarithmic scale in (b). The correlation function $C(t)$ is then determined by integration of $P_A(\lambda, t)$ over values of λ corresponding to the final region B .

D. A Convenient Factorization

In principle, the time correlation function $C(t)$ may be calculated by repeating the procedure described above at each time t . But this umbrella sampling of trajectories is laborious, requiring independent sampling of pathways for each window of the order parameter. Indeed, it mimics a

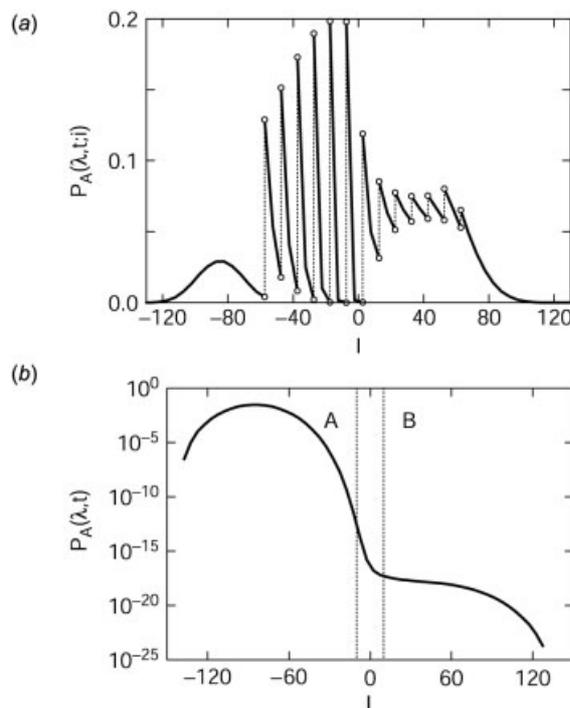


Figure 1.12. (a) Distributions of the order parameter $\lambda(x_t)$ in the window regions $\mathcal{W}[i]$ for proton transfer in the protonated water trimer. In this calculation, the energetics of the trimer were described by an empirical valence bond model [64]. The solid lines denote the individual distributions and the open circles mark points where adjacent order parameter windows overlap. Each distribution was obtained by sampling Newtonian transition pathways of length $t = 147$ fs with shooting and shifting moves at an energy corresponding to a temperature of $T \approx 290$ K. The order parameter λ used in this illustrative calculation was defined as the angle difference $\Delta\varphi \equiv \varphi_2 - \varphi_1$, where φ_1 and φ_2 are angles describing the geometry of the OOO-ring of the cluster (see Fig. 1.2). Region A was defined by $\Delta\varphi < -10^\circ$ and region B by $\Delta\varphi > 10^\circ$. (b) The distribution $P_A(\lambda, t)$ shown on a logarithmic scale is obtained by requiring continuity at the overlap of adjacent windows. In (b) regions A and B are indicated by vertical dotted lines. Integration of $P_A(\lambda, t)$ over the values of λ corresponding to region B yields $C(t) = 9.8 \times 10^{-17}$.

quasistatic process in which the constraint on trajectory endpoints is introduced reversibly. Repeating this process many times, and especially for long trajectories, would incur significant computational cost. In this section, we describe a far less expensive scheme for computing $C(t)$ in a time interval $0 < t < \mathcal{T}$, requiring that umbrella sampling be performed only once and only for relatively short trajectories.

In Section IV.C, we took advantage of the relationship between $C(t)$ and an effective reversible work, $W_{AB}(t)$, to constrain the state of a system at time t . In effect, we gradually applied this constraint to the ensemble of trajectories of length t that begin in state A , obtaining a new ensemble of reactive trajectories. Here, we consider a different process connecting the same two ensembles of trajectories. Because this process is also reversible, it involves the same work $W_{AB}(t)$. In the new process, trajectories of length $\mathcal{T} > t$ that begin in state A are first constrained to visit state B at a time t' . The behavior of these trajectories at times later than t' is unconstrained. The work associated with this first step is thus $W_{AB}(t')$, and it may be computed using umbrella sampling. In the second step, the time at which trajectories are constrained to visit state B is shifted from t' to t . We denote the work associated with the second step as $\Delta W_{AB}(t, t')$. We will show that $\Delta W_{AB}(t, t')$ can be computed for all t at once with little cost, by sampling trajectories that visit state B at any time prior to \mathcal{T} . As a result, computing the total work,

$$W_{AB}(t) = W_{AB}(t') + \Delta W_{AB}(t, t') \quad (1.89)$$

for all times $t < \mathcal{T}$ requires umbrella sampling only for a single time t' . In addition, the time t' may be chosen to be small, further reducing the computational cost.

The quantity $\Delta W_{AB}(t, t')$ is the effective reversible work to shift the time at which trajectories are required to reach state B . It can be written in terms of correlation functions $C(t)$ and $C(t')$:

$$\exp[-\Delta W_{AB}(t, t')] = \frac{C(t)}{C(t')} = \frac{\langle h_A(x_0)h_B(x_t) \rangle}{\langle h_A(x_0)h_B(x_{t'}) \rangle} \quad (1.90)$$

For simplicity, we define a function

$$R(t, t') = \exp[-\Delta W_{AB}(t, t')] \quad (1.91)$$

such that $R(t, t') = C(t)/C(t')$.

In order to calculate $R(t, t')$ efficiently, we introduce a path functional $H_B[x(\mathcal{T})]$ that is unity if at least one state along the trajectory $x(\mathcal{T})$ lies in B and vanishes otherwise,

$$H_B[x(\mathcal{T})] \equiv \max_{0 \leq i \leq \mathcal{T}} h_B(x_i) \quad (1.92)$$

Because $h_B(x_t) = h_B(x_t)H_B[x(\mathcal{T})]$ for all $t < \mathcal{T}$, we may rewrite Eq. (1.91)

as

$$R(t, t') = \frac{\langle h_A(x_0)h_B(x_t)H_B[x(\mathcal{T})] \rangle}{\langle h_A(x_0)H_B[x(\mathcal{T})] \rangle} \times \frac{\langle h_A(x_0)H_B[x(\mathcal{T})] \rangle}{\langle h_A(x_0)h_B(x_{t'})H_B[x(\mathcal{T})] \rangle} \quad (1.93)$$

Each of the factors in this equation can be written as an average over pathways beginning in region A and visiting region B at some time in the interval $[0, \mathcal{T}]$. These pathways are reactive, but are not required to reside in B at any specific time. Their weight functional is

$$\mathcal{P}_{AB}^*[x(\mathcal{T})] \equiv h_A(x_0)\mathcal{P}[x(\mathcal{T})]H_B[x(\mathcal{T})]/Z_{AB}^*(\mathcal{T}) \quad (1.94)$$

The partition function normalizing $\mathcal{P}_{AB}^*[x(\mathcal{T})]$ is given by

$$Z_{AB}^*(\mathcal{T}) = \int \mathcal{D}x(\mathcal{T}) h_A(x_0)\mathcal{P}[x(\mathcal{T})]H_B[x(\mathcal{T})] \quad (1.95)$$

We denote averages in this ensemble by the subscript AB and an asterisk, for example,

$$\langle h_B(x_t) \rangle_{AB}^* \equiv \frac{\int \mathcal{D}x(\mathcal{T}) h_A(x_0)\mathcal{P}[x(\mathcal{T})]h_B(x_t)H_B[x(\mathcal{T})]}{\int \mathcal{D}x(\mathcal{T}) h_A(x_0)\mathcal{P}[x(\mathcal{T})]H_B[x(\mathcal{T})]} \quad (1.96)$$

The path distribution $\mathcal{P}_{AB}^*[x(\mathcal{T})]$ can be sampled with algorithms nearly identical to those described in Section III. We simply replace the characteristic function $h_B(x_t)$ with the functional $H_B[x(\mathcal{T})]$ in all expressions. In sampling this distribution, one may determine $R(t, t')$ for all times $t, t' < \mathcal{T}$ at once. Also, in this case efficiency can be increased by growing the backward segments of a shooting move first, because only in this case early rejection can be exploited to reduce the computational cost of shooting moves.

The population correlation function $C(t)$ can be efficiently determined over the entire interval $0 < t < \mathcal{T}$ as follows: First, $\langle h_B(x_t) \rangle_{AB}^*$ is computed by sampling pathways of a single length \mathcal{T} . In order to extract a rate constant from $C(t)$, \mathcal{T} must be longer than the transient time τ_{mol} , as indicated by linear scaling of $\langle h_B(x_t) \rangle_{AB}^*$. In general, it is best to experiment with the value of \mathcal{T} , in order to access the asymptotic regime without wasting computational effort on unnecessarily long pathways. After computing $\langle h_B(x_t) \rangle_{AB}^*$, $C(t)$ is determined at a single time t' using the umbrella sampling methods described in Section IV.C. To reduce the cost of this calculation, t' may be chosen to be much shorter than \mathcal{T} . Combining results

1. Calculate the path average $\langle h_B(x_i) \rangle_{AB}^*$ using the transition path ensemble $\mathcal{P}_{AB}^*[x(\mathcal{T})]$.
2. Determine the time derivative $d\langle h_B(x_i) \rangle_{AB}^*/dt$.
3. If $d\langle h_B(x_i) \rangle_{AB}^*/dt$ displays a plateau, proceed with step 4, otherwise repeat procedure with a longer time \mathcal{T} starting from step 1.
4. Calculate the time correlation function $C(t')$ at a specific time t' in the interval $[0, \mathcal{T}]$ with umbrella sampling.
5. Determine $C(t)$ in the entire interval $[0, \mathcal{T}]$ using $C(t) = C(t') \langle h_B(t) \rangle_{AB}^* / \langle h_B(t') \rangle_{AB}^*$.
6. Take the derivative of $C(t)$ and extract the rate constant k_{AB} from the plateau value of $\dot{C}(t)$.

Scheme 1.8. Algorithm for the calculation of transition rate constants.

of these two steps yields $C(t)$ in the whole interval of interest,

$$C(t) = \frac{\langle h_B(x_t) \rangle_{AB}^*}{\langle h_B(x_{t'}) \rangle_{AB}^*} \times C(t') \quad (1.97)$$

A rate constant can then be obtained from the asymptotic slope of $C(t)$. This algorithm for calculating rate constants is summarized in Scheme 1.8. We emphasize that $C(t)$ is calculated exactly in this prescription. Specifically, we do not assume an underlying separation of timescales. Rather, the behavior of $C(t)$ reveals whether such a separation exists.

The results of the first step of a rate constant calculation in the transition path sampling method are shown in Figure 1.13 for proton transfer in the protonated water trimer. Panel (a) depicts the path average $\langle h_B(t) \rangle_{AB}^*$ as a function of time. Initially $\langle h_B(t) \rangle_{AB}^*$ is zero because pathways must cross a gap separating region A from region B . At least 50 fs are necessary to cross this gap. After this time, $\langle h_B(t) \rangle_{AB}^*$ starts to grow and reaches a linear regime after ~ 200 fs. Figure 1.13(b) shows the time derivative of $\langle h_B(x_t) \rangle_{AB}^*$, which displays a plateau after 200 fs. A plateau value of $\langle \dot{h}_B(t) \rangle_{AB, \text{plateau}}^* = 0.34 \times 10^{-2} \text{ fs}^{-1}$ can be inferred from the figure. To calculate the rate constant k_{AB} , we now use the time correlation function $C(t') = 9.83 \times 10^{-17}$ obtained by umbrella sampling for a particular time $t' = 147$ fs (see Fig. 1.12). Note that t' can be smaller than \mathcal{T} and does not even need to be in the plateau regime. At t' the path average $\langle h_B(t') \rangle_{AB}^* = 0.16$. By combining these values, we obtain a rate constant of $k_{AB} = 2.1 \times 10^{-3} \text{ s}^{-1}$.

E. Correspondence with Reactive Flux Theory

The factorization of $C(t)$ introduced in Section IV.D is reminiscent of the factorization used in conventional rate constant calculations [3]. In both cases, a reversible work calculation is needed to compute the correlation of

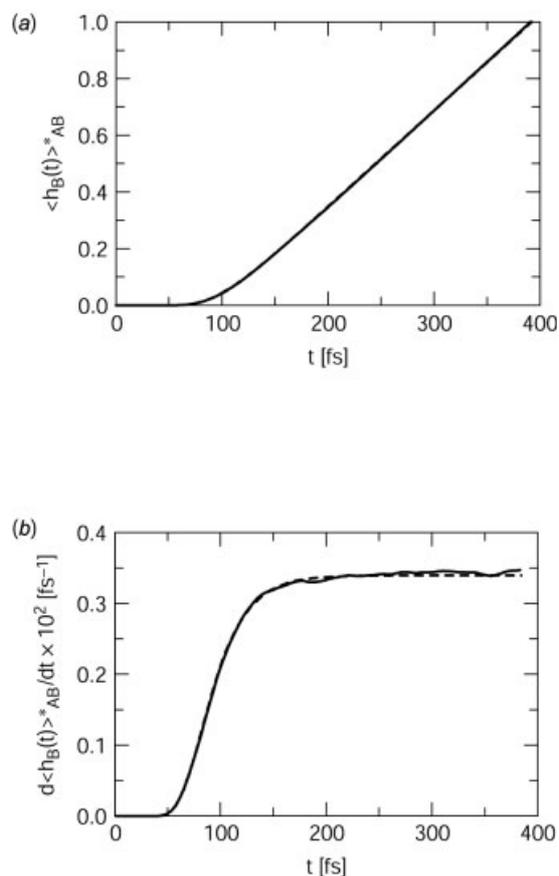


Figure 1.13. (a) Path average $\langle h_B(t) \rangle_{AB}^*$ for proton transfer in the protonated water trimer computed directly in a transition path simulation (solid line) and from the distribution of transition times (dashed line) as discussed in Section IV.F. The average was obtained by sampling Newtonian transition pathways of length $t \sim 390$ fs at a total energy equivalent to a temperature of $T \approx 290$ K. Region A was defined by $\Delta\varphi < -10^\circ$ and region B by $\Delta\varphi > 10^\circ$. (b) Time derivative of the $\langle h_B(t) \rangle_{AB}^*$ curves shown in panel (a). The time derivative reaches a plateau only for times larger than the longest transition time τ , that is, for times larger than ~ 200 fs.

population fluctuations at a specific time. The remaining time dependence of correlations can be computed independently of the absolute scale, and at much less cost. For the reactive flux method, a reversible work calculation at time 0^+ corresponds to transition state theory. (Trajectories of infinitesimal length do not recross the dividing surface $q = q^*$.) In the method based

on transition path sampling, the specific time at which $C(t)$ is calculated remains arbitrary. In fact, this resemblance is more than superficial. For a particular choice of the boundaries defining states A and B , the two factorization schemes are intimately related.

If a single dividing surface, $q = q^*$, is chosen as the boundary of both states A and B , then the two states are adjoining: $h_A(x) + h_B(x) = \theta(q - q^*) + \theta(q^* - q) = 1$. For such a choice, the time derivative of the population correlation function in Eq. (1.78) exactly corresponds to the reactive flux through the $q = q^*$ surface, $k(t) = dC(t)/dt$. With the factorizations we have described,

$$k_{\text{TST}}\kappa(t) = C(t') \frac{dR(t, t')}{dt} \quad (1.98)$$

The factors $\kappa(t)$ and $dR(t, t')/dt$ both describe the full time dependence of this flux. They therefore differ only by a proportionality constant, $C(t')/k_{\text{TST}}$. For a choice of t' that is vanishingly small, this constant is just $\dot{C}(t')t'/k_{\text{TST}} = t'$, since $C(0^+) = 0$. For adjoining states and small t' , our factorization of Eq. (1.78) is thus related to that of reactive flux theory by a simple multiplicative constant.

When states A and B are not adjoining, reactive trajectories must cross a region in which $h_A + h_B = 0$. In this case, the transient, short-time behavior of $\dot{C}(t)$ will differ from that of $k(t)$. In particular, $\dot{C}(0^+) = 0 \neq k_{\text{TST}}$, because infinitesimally short trajectories cannot bridge the gap between A and B . At long times, however, $\dot{C}(t)$ approaches $k(t)$ asymptotically, provided the region of phase space included in the gap has nearly vanishing weight. In this case, $\dot{C}(t)$ and $k(t)$ are described by the same phenomenological kinetics as $t \rightarrow \infty$. We have verified this asymptotic correspondence for several systems [6, 12].

F. How Long Should Pathways Be?

In harvesting transition pathways, it is important that trajectories are long enough to exhibit typical barrier crossing behavior. Trajectories that are constrained to connect stable states in a very short time may not be representative of the natural ensemble of reactive trajectories. In this section, we discuss quantitative criteria for the appropriate minimum length (in time) of trajectories. These criteria are based on the condition that correlations of population fluctuations have reached their asymptotic behavior. They are thus related to the corresponding condition for the reactive flux method, namely, that the flux is essentially constant on a molecular timescale. Depending on the chosen boundaries of stable regions, however,

the appropriate duration of trajectories can be quantitatively different for transition path sampling.

In general, regions A and B defined in transition path sampling are not adjacent and the system must cross a gap between these regions during a transition. For example, the results depicted in Figure 1.13 were obtained from calculations in which region A was defined by $\Delta\varphi < -10^\circ$ while region B was characterized by $\Delta\varphi > 10^\circ$. Since in general the reaction coordinate describing the course of the transition is unknown, it is important that regions A and B do not overlap and do not contain states belonging to the basin of attraction of the other stable region. Due to the associated finite width of the gap the system needs a certain finite transition time τ to cross the gap. Therefore, the pathways in transition path sampling must be long enough to accommodate the transition times τ of the most important pathways. In the following paragraph, we will justify and analyze this requirement in greater detail.

Consider, again, a system in which transitions from A to B occur rarely. Along each possible pathway $x(\mathcal{T})$ of length \mathcal{T} the system spends a certain time $\tau_A[x(\mathcal{T})]$ in A , then stays between A and B for a time $\tau[x(\mathcal{T})]$, and finally arrives in B spending the rest of the time, $\tau_B[x(\mathcal{T})]$, there. For simplicity, we assume that there are no multiply entries and exits to and from regions A and B . Hence, every pathway $x(\mathcal{T})$ has a single well defined transition time $\tau[x(\mathcal{T})]$. Since different pathways have different transition times it is convenient to introduce a distribution of transition times,

$$p(\tilde{\tau}; \mathcal{T}) \equiv \int \mathcal{D}x(\mathcal{T}) \mathcal{P}_{AB}[x(\mathcal{T})] \delta(\tilde{\tau} - \tau[x(\mathcal{T})]) \quad (1.99)$$

where the argument \mathcal{T} in $p(\tau; \mathcal{T})$ indicates that the distribution of transition times depends on the total path length \mathcal{T} . As an example for such a distribution Figure 1.14 shows $p(\tau; \mathcal{T})$ for proton-transfer pathways in the protonated water trimer described in Section I. As can be inferred from the figure, the system needs at least ~ 50 fs to cross the gap and some trajectories spend up to 200 fs in the region between A and B .

Now, the path average $\langle h_B(x_t) \rangle_{AB}^*$ appearing in the expressions for the rate constant in Eq. (1.97) can be expressed as an integral over the distribution of transition times $p(\tau; \mathcal{T})$. Analysis of that expression then yields a criterion for the path length \mathcal{T} .

To demonstrate this we rewrite Eq. (1.96) as

$$\langle h_B(x_t) \rangle_{AB}^* = \int \mathcal{D}x(\mathcal{T}) \mathcal{P}_{AB}^*[x(\mathcal{T})] \theta(t - \tau[x(\mathcal{T})] - \tau_A[x(\mathcal{T})]) \quad (1.100)$$

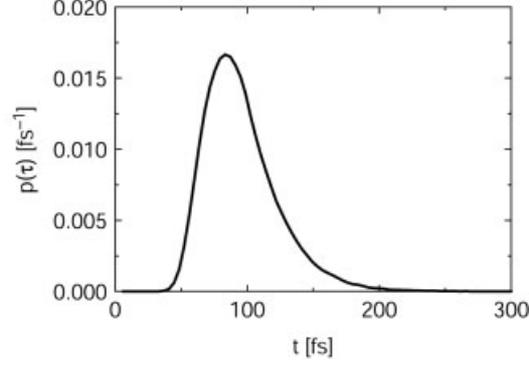


Figure 1.14. Distribution of transition times τ for proton transfer in the protonated water trimer for the same system as in Figure 1.13.

where $h_B(x_t)$ for the specific pathway $x(\mathcal{T})$ has been expressed using the Heaviside step function $\theta(t)$. In doing so we have assumed that trajectory $x(\mathcal{T})$ leaves region A and enters region B exactly once. We next insert two Dirac delta functions into the right-hand side of Eq. (1.100) and integrate over their arguments:

$$\begin{aligned} \langle h_B(x_t) \rangle_{AB}^* &= \int d\tilde{\tau} \int d\tilde{\tau}_A \int \mathcal{D}x(\mathcal{T}) \mathcal{P}_{AB}^*[x(\mathcal{T})] \theta(t - \tau[x(\mathcal{T})] - \tau_A[x(\mathcal{T})]) \\ &\quad \times \delta(\tilde{\tau} - \tau[x(\mathcal{T})]) \delta(\tilde{\tau}_A - \tau_A[x(\mathcal{T})]) \end{aligned} \quad (1.101)$$

The delta functions allow us to replace the functionals $\tau[x(\mathcal{T})]$ and $\tau_A[x(\mathcal{T})]$ appearing in the argument of the Heaviside step function by the numbers $\tilde{\tau}$ and $\tilde{\tau}_A$, respectively. We therefore obtain

$$\langle h_B(x_t) \rangle_{AB}^* = \int d\tau \int d\tau_A \theta(t - \tau - \tau_A) p(\tau, \tau_A; \mathcal{T}) \quad (1.102)$$

where we have renamed the integration variables. The distribution function $p(\tau, \tau_A; \mathcal{T})$ is the probability to observe both τ and τ_A in a path with length \mathcal{T} ,

$$p(\tau, \tau_A; \mathcal{T}) \equiv \int \mathcal{D}x(\mathcal{T}) \mathcal{P}_{AB}^*[x(\mathcal{T})] \delta(\tilde{\tau} - \tau[x(\mathcal{T})]) \delta(\tilde{\tau}_A - \tau_A[x(\mathcal{T})]) \quad (1.103)$$

For a given transition time τ , the time τ_A spent in region A can have values

ranging from 0 to $T - \tau$. Assuming that τ_A is uniformly distributed in this interval and that it is furthermore statistically independent from τ , that is,

$$p(\tau, \tau_A; \mathcal{F}) = \frac{\theta(\mathcal{F} - \tau - \tau_A)\theta(\tau_A)}{\mathcal{F} - \tau} p(\tau; \mathcal{F}) \quad (1.104)$$

One can carry out the integration over τ_A in Eq. (1.102) obtaining

$$\langle h_B(x_t) \rangle_{AB}^* = \int_0^{\mathcal{F}} d\tau p(\tau; \mathcal{F}) \theta(t - \tau) \frac{t - \tau}{\mathcal{F} - \tau} \quad (1.105)$$

Thus, the contribution to $\langle h_B(x_t) \rangle_{AB}^*$ of pathways with transition time τ is zero for $0 \leq t \leq \tau$ and grows linearly from 0 to 1 in the interval $\tau < t \leq \mathcal{F}$. The path average $\langle h_B(x_t) \rangle_{AB}^*$ computed from Eq. (1.105) for proton-transfer pathways in the protonated water trimer is shown along with its time derivative in Figure 1.13 as dashed lines. The excellent agreement of these results with the results obtained in a path simulation directly indicates that the assumptions made in Eq. (1.104) are correct for the example system.

Equation (1.105) can be used to obtain a criterion on the path length required to sample all important pathways. According to the reasoning of the preceding sections, pathways are sufficiently long if the time derivative $dC(t)/dt$ reaches a plateau. But $dC(t)/dt$ can reach a plateau only if $\langle h_B(x_t) \rangle_{AB}^*$ does so. In this plateau regime, the first derivative of $\langle h_B(x_t) \rangle_{AB}^*$ is constant and its second derivative must vanish. By calculating the second time derivative of $\langle h_B(x_t) \rangle_{AB}^*$ from Eq. (1.105), we obtain

$$\frac{d^2}{dt^2} \langle h_B(x_t) \rangle_{AB}^* = \frac{p(t; \mathcal{F})}{\mathcal{F} - t} \quad (1.106)$$

Consequently, $dC(t)/dt$ can reach a plateau value only if the path length \mathcal{F} is larger than the maximum transition time τ_{\max} after which the distribution of transition times $p(\tau; \mathcal{F})$ vanishes. More exactly, the existence of a plateau requires that $p(\tau)/(\mathcal{F} - \tau)$ vanishes in the whole plateau regime. The maximum transition time τ_{\max} can be considerably larger than the molecular transient τ_{mol} from reactive flux theory. For example, $\tau_{\max} \approx 250$ fs for proton transfer in the protonated water trimer as can be inferred from Figure 1.14 (and from Fig. 1.13). In contrast, the time-dependent transmission coefficient $\kappa(t)$ shown in Fig. 1.15 indicates that a trajectory initiated from the dividing surface defined by $\Delta\phi = 0$ commits to one of the stable states in less than 100 fs. In general, the path length should therefore be chosen to be larger than the maximum transition time or the time necessary to commit into one of the stable states, whichever is larger.

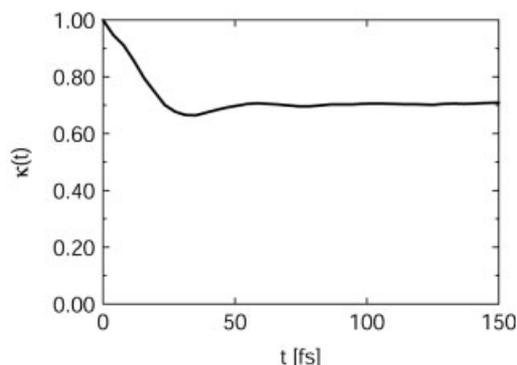


Figure 1.15. Transmission coefficient $\kappa(t)$ as a function of time for the same system as in Figure 1.14. After a molecular transient time of less than 100 fs $\kappa(t)$ reaches a plateau.

V. ANALYZING TRANSITION PATHWAYS

Thus far, we have described efficient methods for harvesting pathways that exhibit transitions of interest. We have also detailed how these trajectories may be used to characterize reactive dynamics at a macroscopic level by computing rate constants. In this section, we focus on gleaning mechanistic information from harvested pathways, in order to understand transitions at a microscopic level. For complex systems, this task is made challenging by the huge number of irrelevant degrees of freedom. It is generally not possible to recognize the coordinate that describes a transition's progress simply by visualizing trajectories (e.g., using computer graphics) or by following the behavior of preconceived order parameters. Rather, identifying such a reaction coordinate requires a thorough statistical analysis of pathways and transition states. Here, we describe several concepts and diagnostics that facilitate this analysis.

A. Reaction Coordinates and Order Parameters

A coordinate that successfully distinguishes between two basins of attraction does not necessarily characterize the dynamical bottleneck between them. We emphasize this fact by referring to discriminating coordinates as *order parameters*, and reserving the term *reaction coordinate* for the specific coordinate that describes a transition's dynamical mechanism. This distinction is illustrated in Figure 1.16. For the free energy landscapes in both panels of the figure, the coordinate q serves as a reasonable order parameter. Specifically, the range of values of q in the basin of attraction of state A does

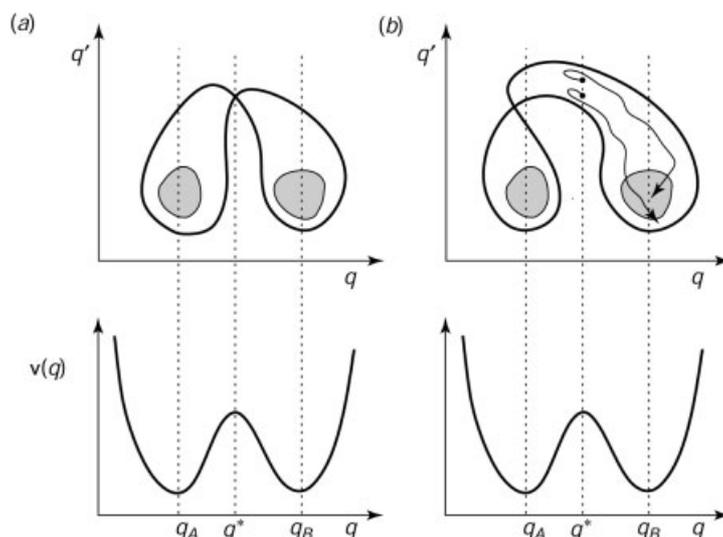


Figure 1.16. Two illustrative two dimensional free energies $w(q, q')$ depending on two variables q and q' , and their corresponding reduced free energy functions $w(q) \equiv -k_B T \ln \int dq' \exp[-\beta w(q, q')]$. In both cases, $w(q)$ has the same bistable form, but in (a) the coordinate q is a reasonable reaction coordinate, as the transition state surface coincides with $q = q^*$. In (b), on the other hand, q is not at a reasonable reaction coordinate. The orthogonal variable q' is crucial for the mechanism of $A \rightarrow B$ transitions, and the maximum in $w(q)$ at $q = q^*$ does not coincide with the transition state surface. The trajectories initiated at configurations with $q(r) = q^*$ and all ending in B illustrate this.

not significantly overlap the range of values in the basin of state B . As a result, the reduced free energy $w(q)$, obtained by integrating out the orthogonal coordinate q' , has a pronounced maximum, located at $q = q^*$. But the utility of q as a reaction coordinate differs greatly in the two scenarios. In (a), the ensemble of configurations with $q = q^*$ is a good approximation to the transition state surface dividing the basins of attraction. Here, q' is essentially irrelevant to the dynamical bottleneck, and q may be accurately called the reaction coordinate. In (b), however, barrier crossing occurs primarily in the direction of q' . Typical configurations with $q = q^*$ are poor approximations to true transition states, lying well within the basin of attraction of state B . In this case, the order parameter q does not coincide with the reaction coordinate. This latter situation is in fact common in physical systems. For example, symmetry and density successfully characterize fluid and solid phases of a simple material, but the bottleneck for crystallization of a supercooled liquid involves the size and structure of a crystal nucleus [65,66]. One should thus not expect order parameters in general to function as reaction coordinates.

B. The Separatrix and the Transition State Ensemble

In a simple system, saddle points of the potential energy surface coincide with transition states. At these stationary points, forces vanish, and the system is not driven to either of the corresponding stable states. In other words, the stable states are equally accessible from the transition state. This concept of equal accessibility is readily extended for complex systems, in which individual saddle points are no longer relevant for transitions of interest. Specifically, transition states are defined to be configurations from which the system relaxes into one or the other stable state with equal probability [27, 67]. Transition states defined in this way are generally unrelated to the local topography of the potential energy surface, and may be strongly influenced by entropy. The set of all such transition states forms the *separatrix*, a high-dimensional surface dividing basins of attraction.

In order to formalize this definition of the separatrix, we introduce a function $p_B(r, t_s)$ called the *committor*:

$$p_B(r, t_s) \equiv \frac{\int \mathcal{D}x(t_s) \mathcal{P}[x(t_s)] \delta(r_0 - r) h_B(x_{t_s})}{\int \mathcal{D}x(t_s) \mathcal{P}[x(t_s)] \delta(r_0 - r)} \quad (1.107)$$

Since transition probabilities are normalized, the denominator on the right-hand side of this equation is just the equilibrium probability distribution for configuration r . The parameter $p_B(r, t_s)$ is the probability that a system with initial configuration r will reside in state B at time t_s . The analogous committor for state A , $p_A(r, t_s)$ is similarly defined. If the timescale of molecular fluctuations, τ_{mol} , is well separated from the typical time between spontaneous transitions, τ_{rxn} , then p_A and p_B will be nearly independent of t_s for $\tau_{\text{mol}} < t_s \ll \tau_{\text{rxn}}$. In this time regime, a trajectory will have committed to one or the other basin of attraction, even if it originated in the transition state region. Once such a trajectory has committed, subsequent spontaneous transitions are unlikely for $t_s \ll \tau_{\text{rxn}}$. In the discussions that follow, we assume that t_s has been chosen in this asymptotic regime, and omit the explicit dependence of committors on time. With this choice, p_A and p_B are essentially functions only of a configuration r , quantifying the propensity of that configuration to relax into a particular basin of attraction under the system's intrinsic dynamics. States A and B are thus equally accessible when $p_A(r) = p_B(r)$. This equation defines the separatrix.

The committor p_B is a direct statistical indicator for the progress of transitions from A to B . In this sense, it is an ideal reaction coordinate. But interpreting this highly nonlinear function of atomic coordinates in terms of molecular motions and intuitively meaningful fields is not straightforward. Understanding a transition's mechanism basically amounts to identifying

the simple, low-dimensional coordinates that determine p_B . The following sections are concerned with this task. We first describe an efficient scheme for locating the separatrix, and then introduce unambiguous diagnostics for verifying the determinants of p_B . These tools greatly aid the interpretation of trajectories harvested by transition path sampling, but are not unique to that methodology. Indeed, any proposed reaction coordinate must satisfy the statistical criteria we describe in Section V.D.

C. Computing Committors

Locating the separatrix by screening typical configurations for $p_A \approx p_B$ would be an extraordinary computational challenge. Not only are such configurations rare at equilibrium, but evaluating committors requires the generation of many fleeting trajectories for each tested configuration. Fortunately, the pathways harvested by transition path sampling are guaranteed to include examples of the separatrix. In collecting a properly weighted ensemble of transition states, it is only necessary to screen the configurations along these relatively short, reactive pathways. Transition paths may even cross the separatrix more than once. In this case, each crossing provides a valid example of the transition state ensemble. Such multiple crossings are not uncommon in complex systems. For example, isomerization pathways of an alanine dipeptide cross the separatrix up to seven times [17].

The computational cost of identifying transition states is further reduced by the fact that configurations far from the separatrix may be excluded rather quickly. In practice, $p_B(r)$ is estimated for a particular configuration by initiating a finite number N of fleeting trajectories, $x^{(i)}(t_s)$, from that configuration

$$p_B^{(N)}(r) \approx \frac{1}{N} \sum_{i=1}^N h_B(x_{t_s}^{(i)}) \quad (1.108)$$

where we imagine that initial momenta are drawn at random from the appropriate distribution. This procedure is graphically illustrated in Figure 1.17. For large N , $p_B^{(N)}(r)$ is a Gaussian random variable according to the central limit theorem, with fluctuations of size

$$\sigma = \sqrt{\langle (p_B^{(N)} - p_B)^2 \rangle} = \sqrt{\frac{p_B(1 - p_B)}{N}} \quad (1.109)$$

Here, $\langle \dots \rangle$ indicates an average over many independent calculations of $p_B^{(N)}$. These fluctuations describe the typical errors in our estimate of p_B using a

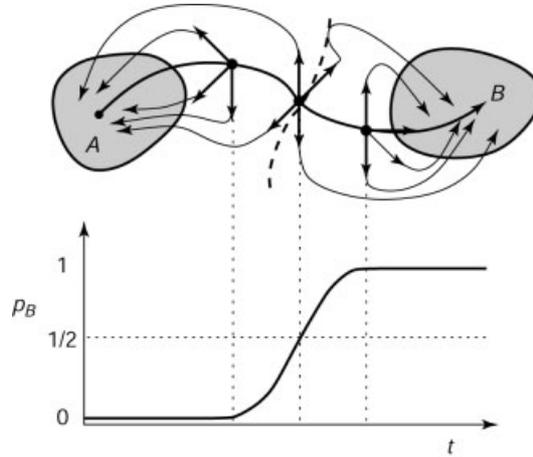


Figure 1.17. The committor p_B for a time slice at time t along a transition path (thick solid line, top panel) is computed by determining the fraction of fleeting trial trajectories (thin solid lines) initiated with randomized momenta that reach region B in a time t_s . Typically 10–100 of these fleeting trajectories are needed to obtain p_A and p_B with sufficient accuracy. While $p_B \approx 0$ for the left time slice in the top panel, because all trajectories started from that time slice end in A , $p_B \approx 1$ for the time slice on the right, because all the corresponding fleeting trajectories end in B . The configurations for which $p_B \approx \frac{1}{2} \approx p_A$ are part of the separatrix (thick dashed line) and are called transition states.

finite number of fleeting trajectories. The same fleeting trajectories $x^{(i)}(t_s)$ may be used to estimate p_A . For simplicity, we assume here that $p_A + p_B = 1$, that is, there are no additional basins of attraction. In this case, the error in our estimate of p_A is identical to that for p_B . Typical errors are only of size $\sigma \sim \sqrt{p_B/N} \ll 1$ when $p_A \approx 1$. Similarly, when $p_B \approx 1$, $\sigma \sim \sqrt{p_A/N} \ll 1$. We may thus determine that a configuration lies deep within the basin of attraction of A or B with a relatively small number of trajectories. (Of course, N must be large enough that $p_A^{(N)}$ and $p_B^{(N)}$ may be considered Gaussian random variables. $N \approx 10$ should be sufficient for this purpose.) By contrast, when $p_B \approx p_A \approx 1/2$, typical errors are only small for large N . Determining that a configuration lies near the separatrix thus requires a large number of fleeting trajectories.

The dependence of fluctuations in $p_A^{(N)}$ and $p_B^{(N)}$ on p_A , p_B , and N suggests a general scheme for estimating committor values. A desired level of statistical accuracy, σ_{des} , is chosen in advance. This choice determines the number of fleeting trajectories required to screen a given configuration. A minimum number, N_{min} , of trajectories is first generated to ensure Gaussian

1. Start with a time slice just outside state A .
2. Generate $N_{\min} \sim 10$ short fleeting trajectories from that time slice with momenta selected from the appropriate distribution.
3. Determine p_B from the fraction of paths that end in state B .
4. If the interval $[p_B^{(N)} - \alpha\sigma^{(N)}, p_B^{(N)} + \alpha\sigma^{(N)}]$ does not include $\frac{1}{2}$ the time slice is rejected as a member of the transition state ensemble, otherwise more fleeting trajectories are generated until either $\frac{1}{2}$ is in $[p_B^{(N)} - \alpha\sigma^{(N)}, p_B^{(N)} + \alpha\sigma^{(N)}]$ or an upper limit N_{\max} of the number of fleeting trajectories is reached.
5. If the latter is the case, the time slice is accepted as a member of the transition state ensemble. Otherwise it is rejected.
6. Move to the next time slice on the path and repeat steps 2–4 until region B is reached.

Scheme 1.9. Algorithm for determining transition states on a transition path.

statistics. Additional trajectories are subsequently generated until the error estimate

$$\sigma^{(N)} = \sqrt{\frac{p_B^{(N)}[1 - p_B^{(N)}]}{N}} \quad (1.110)$$

is smaller than σ_{des} . At this stage, the configuration is excluded from the transition state ensemble if the interval $[p_B^{(N)} - \alpha\sigma^{(N)}, p_B^{(N)} + \alpha\sigma^{(N)}]$ does not include the value $\frac{1}{2}$. The constant α depends on the desired confidence level. For example, $\alpha = 2$ is necessary to obtain a confidence level of 95%. If $\frac{1}{2}$ is inside the interval, we continue generating trajectories until the value $\frac{1}{2}$ falls outside the confidence interval. If a maximum number, N_{\max} , of trajectories is reached, and the confidence interval still includes $\frac{1}{2}$, then the configuration is accepted as a member of the transition state ensemble. This procedure, summarized in Scheme 1.9, minimizes the effort of excluding configurations that are far from the separatrix, and yields a consistent level of statistical error. As a rule of thumb, an error of 5% requires that $N_{\max} \sim 100$ trajectories are generated for configurations in the transition state region.

Our definition of a committor in Eq. (1.107) is applicable to both stochastic and deterministic dynamics. In the case of deterministic dynamics, care must be taken that fleeting trajectories are initiated with momenta drawn from the appropriate distribution. As discussed in Section III.A.2, global constraints on the system may complicate this distribution considerably. The techniques described in Section III.A.2 and in the Appendix of [10] for shooting moves may be simply generalized to draw initial momenta at random from the proper equilibrium distribution.

D. Committed Distributions

By screening transition pathways for examples of the separatrix, as described in Section V.D, an ensemble of transition states may be collected. For complex systems, this ensemble is expected to be structurally diverse, reflecting the many ways a particular transition can occur. But the reduced dimensionality of the separatrix ensures that even seemingly dissimilar transition states will share certain patterns. One may begin to search for these patterns by examining the distribution of various order parameters within the transition state ensemble. Coordinates that are important to the dynamical bottleneck will be narrowly distributed about values distinct from those characterizing stable states. In other words, they will be correlated with the transition. Such a correlation does not, however, guarantee that an order parameter is a useful reaction coordinate. Indeed, it is common for many coordinates to follow a transition adiabatically. These coordinates undergo a discernible change on the course from reactants to products, but do not play a significant role in driving the transition.

Distributions of committor values are a powerful diagnostic for differentiating coordinates that drive a transition from those that are simply correlated with it. Consider an order parameter q whose potential of mean force $w(q)$ has a maximum at $q = q^*$. If q serves as a reaction coordinate, then the ensemble of configurations with $q = q^*$ coincides with the separatrix [see Fig. 1.18(a)]. The committor distribution for this ensemble,

$$P(\tilde{p}_B) = \frac{\langle \delta[\tilde{p}_B - p_B(r, t_s)] \delta[q^* - q(r)] \rangle}{\langle \delta[q^* - q(r)] \rangle} \quad (1.111)$$

is thus sharply peaked at $p_B = \frac{1}{2}$. If, on the other hand, q does not characterize the dynamical bottleneck, most configurations in the ensemble with $q = q^*$ lie within the basins of attraction of states A and B [see Fig. 1.18(b)]. In this case, the distribution of p_B is dominated by peaks at $p_B = 0$ and $p_B = 1$. These two scenarios are distinguished by markedly different committor distributions. Figure 1.18 illustrates scenarios for three different schematic free energy surfaces.

In practice, determining informative committor distributions is a two-step process. The first step consists in generating configurations with the constraint $q = q^*$. Various Monte Carlo and molecular dynamics methods have been designed for this purpose [44,68]. In the second step, the committor p_B is determined for each of the collected configurations, using the methods described in Section V.C. From these committor values a histogram $P(p_B)$ can then be constructed. The algorithm for determining committor distributions is summarized in Scheme 1.10.

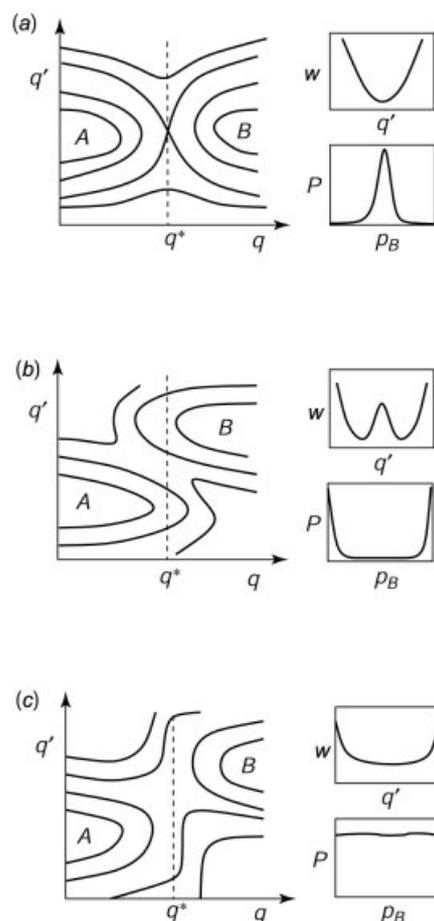


Figure 1.18. Three different free energy landscapes $w(q, q')$, the free energy $w(q^*, q')$ for $q = q^*$ and its corresponding committor distribution $P(p_B)$. (a) The reaction is correctly described by q and the committor distribution of the constrained ensemble with $q = q^*$ peaks at $p_B = 0.5$. (b) q' plays a significant additional role as a reaction coordinate, indicated by the additional barrier in $w(q^*, q')$ and the bimodal shape of $P(p_B)$. (c) Similar to case (b), but now the committor distribution is flat, suggesting diffusive barrier crossing along q' .

As an example, consider again proton transfer in the protonated water trimer as depicted in Figure 1.2. Imagine that the difference $\Delta r \equiv r_1 - r_2$ in the distances of the transferring proton to the oxygens of the donating and accepting water molecule has been postulated as the reaction coordinate for the proton transfer (r_1 and r_2 are defined as indicated in Fig. 1.2). Calculation of the reversible work necessary to control Δr yields a free

1. Compute the free energy $w(q)$ as a function of the proposed reaction coordinate q .
2. Determine the position q^* of the free energy maximum.
3. Collect a series of independent configurations at the top of the barrier, that is, $q = q^*$.
4. Calculate the committor p_B for every configuration.
5. Construct a histogram $P(p_B)$.

Scheme 1.10. Algorithm for determining a committor distribution.

energy profile with two minima separated by a barrier centered at $\Delta r = 0$. The minima of this free energy are located at values of Δr characteristic for the two stable states in each of which a different oxygen atom holds the excess proton. The committor distribution for configurations with $\Delta r = 0$, that is, configurations atop the barrier, is sharply peaked at 0 and 1 indicating that essentially all of these configurations clearly belong to the basin of attraction of one of the two stable states. Thus, the variable Δr is unable to capture the essential dynamics of the transition as the system passes through the dynamical bottleneck. Indeed, as discussed earlier, we found that a functioning reaction coordinate, such as the angular difference $\Delta\varphi$ defined earlier, must include degrees of freedom capable of describing the reorganization of the oxygen ring. The variable Δr merely follows this reorganization adiabatically without having any significant correlation with the transition dynamics. Committor distributions computed for configurations with $\Delta\varphi = 0$ peak at $p_B = \frac{1}{2}$ confirming the efficacy of $\Delta\varphi$ as a reaction coordinate.

This methodology has been also used to extract subtle mechanistic details of the dissociation mechanism of an NaCl ion pair in liquid water [12]. In this case, committor distributions revealed that rearrangement of solvent molecules around the dissociating ion pair significantly contributes to the free energy barrier separating the contact state from the dissociated state.

E. Path Quenching

Scrutinizing path ensembles for common factors can be a difficult task for several reasons. First, the paths are all different on a molecular scale due to thermal motion. Second, there might be different global pathways that are hard to distinguish. It would be helpful if the path ensemble could be somehow smoothed by removing all irrelevant dynamics orthogonal to the reaction coordinate, and reveal the essential dynamics of the optimal pathway. This smoothing can be done by path quenching. The idea behind this technique is that the path sampling method performs a random walk in trajectory space and is usually confined to a basin of attraction within this trajectory space. (This basin of attraction must be distinguished from the

basis of attraction in configuration space, that is, the stable regions A and B .) Each basin of attraction in path space can be characterized by a minimum value of the action as defined in Section III.E. Quenching a path then amounts to minimizing this path action. In this way, different global pathways then correspond to different local minima in the action. As defined in Section III.E, the action $S_{AB}[x(\mathcal{T})]$ is related to the probability of finding a path by

$$P_{AB}[x(\mathcal{T})] = \exp(-S_{AB}[x(\mathcal{T})]) \quad (1.112)$$

The analogy to the canonical Boltzmann distribution $\exp[-\beta V(x)]$ allows sampling of the path ensemble using the dynamical algorithms of Section III.E. By exploiting the same principle, one can use this dynamical algorithm to find the minimum action by sliding downhill in trajectory space. This amounts to quenching the “path temperature” to zero, which leads to a local minimum in path space. There might be many local minima, and in that case annealing the paths to find the nearest global minimum is probably a better option. For Newtonian dynamics with a canonical distribution of initial conditions, quenching the path action corresponds to finding the transition path with the lowest energy.

VI. OUTLOOK

Transition path sampling reduces the computational cost of harvesting reactive trajectories to that of generating dynamical paths of microscopic duration. Because its general methodology makes no reference to specific details of a dynamical system, we expect to see it widely applied. Indeed, examples of successful applications already include chemical reactions in the gas phase and in solution, nucleation of phase transitions, and diffusion of hydrocarbons in zeolites. We anticipate many more.

Some classes of applications will be made challenging by subtle difficulties in the implementation of transition path sampling. One such challenge is the identification of order parameters that unequivocally distinguish between basins of attraction. As discussed in Section II.F, failure to discriminate between stable states may lead to sampling of pathways that are not truly reactive. In many cases, physical intuition, combined with trial and error, is sufficient to determine a reasonable order parameter. But in general, characterizing stable states with only a few variables is a significant problem of pattern recognition. A similar problem arises in characterizing transition state surfaces. Even when many examples of the separatrix can be collected with the methods of Section V.B, identifying a successful reaction

coordinate remains a challenge. A systematic approach to pattern recognition is lacking. Generalizations of standard methods for linear analysis [69] may offer progress in this area.

Sampling long trajectories is also a challenge. When the minimum duration of reactive trajectories (described in Section IV.F) is longer than the timescales characterizing chaos in a deterministic system, shooting moves will fail to generate useful trial trajectories. Due to divergence of neighboring phase-space elements, a typical trial path will differ markedly from the original path, even for the smallest possible momentum displacements. Most trial paths will therefore not be reactive, and thus will be rejected. Analogous difficulties hinder the sampling of strongly damped stochastic paths. When inertia is essentially inconsequential, only shooting points near the separatrix will generate reactive trajectories with a reasonable probability. In this case as well, sampling efficiency is severely degraded by the rejection of nearly all trial paths. It seems that a new genre of trial moves is necessary to confront these problems. One might imagine introducing a bias that imposes correlations between an existing path and trial trajectories generated from it.

For systems that are extremely large or complex, generating trajectories of even microscopic duration is computationally unfeasible. Clearly, harvesting true dynamical pathways of these systems is not possible. One might still hope, however, to construct chains of states that qualitatively resemble reactive trajectories. The stochastic path approach of Elber and co-workers [29, 70] focuses on such pathways. This method utilizes a weight functional for large time step trajectories that would normally be grossly unstable. The instability is artificially suppressed by the requirement that paths connect two given stable states. Although the stochastic path approach greatly extends the range of systems that can be considered, the pathways it generates are in some ways fundamentally dissimilar from natural trajectories: The effective dynamical propagation rule violates the fluctuation–dissipation theorem, and does not conserve an equilibrium distribution. It is not clear how to remedy these serious shortcomings while preserving the computational frugality that makes the method appealing. Transition path sampling may offer a more controlled context for this kind of approach.

In the course of future applications of our methods, we expect that many of the above difficulties will be overcome. Others’ experience and insight will be essential for these advances. We also expect to see the perspectives underlying transition path sampling to be exploited in novel ways. An ensemble view of trajectories is especially promising for the studying dynamical structures like those of supercooled liquids [71, 72]. One can imagine that changes in relaxation patterns as a system approaches the glass transition are best captured using a “thermodynamics” of trajectories. While

it is not obvious how to obtain simplifying insight into this thermodynamics (analogous to, say, van der Waals's picture of fluids), the ideas and formalism we have presented here would seem to be a first step.

ACKNOWLEDGEMENTS

We are very grateful to David Chandler, who inspired and directed the work described in this chapter. We acknowledge his guidance and support throughout our stay at UC Berkeley. This work also profited from enlightening and stimulating discussions with Hans Andersen, Christian Bartels, Felix Csajka, Gavin Crooks, Daan Frenkel, Daniel Laria, Jin Lee, Ka Lum, Jordi Marti, Tom McCormick, Vijay Pande, Javier Rodriguez, Dan Rokhsar, Udo Schmitt, Berend Smit, Sean Sun, Pieter-Rein Ten Wolde, and Thijs Vlugt.

REFERENCES

1. C. J. Cerjan and W. H. Miller, *J. Chem. Phys.* **75**, 2800 (1981).
2. J. P. K. Doye and D. J. Wales, *Z. Phys. D* **40**, 194 (1997).
3. J. B. Anderson, *J. Chem. Phys.* **58**, 4684 (1973); C. H. Bennett, in *Algorithms for Chemical Computations*, edited by R. E. Christoffersen, (American Chemical Society, Washington, D.C. 1977) pp. 63–97; D. Chandler, *J. Chem. Phys.* **68**, 2959 (1978).
4. L. P. Kadanoff, *Phys. Today*, **58**(5), 34 (August 2001).
5. C. Dellago, P. G. Bolhuis, F. S. Csajka, and D. Chandler, *J. Chem. Phys.* **108**, 1964 (1998).
6. C. Dellago, P. G. Bolhuis, and D. Chandler, *J. Chem. Phys.* **108**, 9263 (1998).
7. P. G. Bolhuis, C. Dellago, and D. Chandler, *Faraday Discuss.* **110**, 421 (1998).
8. F. S. Csajka and D. Chandler, *J. Chem. Phys.* **109**, 1125 (1998).
9. D. Chandler, in *Classical and Quantum Dynamics in Condensed Phas Simulations*, edited by B. J. Berne, G. Ciccotti, and D. F. Coker (World Scientific, Singapore, 1998), p. 51.
10. P. L. Geissler, C. Dellago, and D. Chandler, *Phys. Chem. Chem. Phys.* **1**, 1317 (1999).
11. C. Dellago, P. G. Bolhuis, and D. Chandler, *J. Chem. Phys.* **110**, 6617 (1999).
12. P. L. Geissler, C. Dellago, and D. Chandler, *J. Phys. Chem. B* **103**, 3706 (1999).
13. T. J. H. Vlugt, C. Dellago, and B. Smit, *J. Chem. Phys.* **113**, 8791 (2000).
14. P. L. Geissler, T. v. Voorhis, and C. Dellago, *Chem. Phys. Lett.* **324**, 149 (2000); E. P. F. Lee and J. M. Dyke, *Mol. Phys.* **73**, 375 (1991);
15. P. L. Geissler, C. Dellago, D. Chandler, J. Hutter, and M. Parrinello, *Chem. Phys. Lett.* **321**, 225 (2000).
16. P. L. Geissler and D. Chandler, *J. Chem. Phys.* **113**, 9759 (2000).
17. P. G. Bolhuis, C. Dellago, and D. Chandler, *Proc. Natl. Acad. USA* **97**, 5877 (2000).
18. P. G. Bolhuis, D. Chandler, C. Dellago, and P. L. Geissler, *Annu. Rev. Phys. Chem.*, in print (2000).
19. P. G. Bolhuis, C. Dellago, P. L. Geissler, and D. Chandler, *J. Phys.: Condens. Matter* **12**, A147 (2000).
20. J. Marti and F. S. Csajka, *J. Chem. Phys.* **113**, 1154 (2000).
21. J. Marti, F. S. Csajka, and D. Chandler, *Chem. Phys. Lett.* **328**, 169 (2000).

22. P. G. Bolhuis and D. Chandler, *J. Chem. Phys.* **113**, 8154 (2000).
23. T. J. H. Vlugt and B. Smit, *Phys. Chem. Comm.* **2**, 1 (2001).
24. P. L. Geissler, C. Dellago, D. Chandler, J. Hutter, and M. Parrinello, *Science* **291**, 2121 (2001).
25. D. Laria, J. Rodriguez, C. Dellago, and D. Chandler, *J. Phys. Chem. A* **105**, 2646 (2001).
26. G. E. Crooks and D. Chandler, *Phys. Rev. E* **64**, 026109 (2001).
27. L. R. Pratt, *J. Chem. Phys.* **85**, 5045 (1986).
28. R. Elber and M. Karplus, *Chem. Phys. Lett.* **139**, 375 (1987).
29. R. Elber, J. Meller, and R. Olender, *J. Phys. Chem. B* **103**, 899 (1999).
30. G. Henkelman, G. Johannesson, and H. Jonsson, in *Progress on Theoretical Chemistry and Physics*, edited by S. D. Schwartz, (Kluwer Academic Publishers, Dordrecht, 2000).
31. D. Passerone and M. Parrinello, *Phys. Rev. Lett.* **87**, 108302 (2001).
32. E. M. Sevick, A. T. Bell, and D. N. Theodorou, *J. Chem. Phys.* **98**, 3196 (1993).
33. A. F. Voter, *Phys. Rev. Lett.* **78**, 3908 (1997).
34. R. E. Gillilan and K. R. Wilson, *J. Chem. Phys.* **105**, 9299 (1996).
35. P. Eastman, N. Gronbech-Jensen, and S. Doniach, *J. Chem. Phys.* **114**, 3823 (2001).
36. R. Car and M. Parrinello, *Phys. Rev. Lett.* **55**, 2471 (1985).
37. W. G. Hoover, *Computational Statistical Mechanics* (Elsevier, Amsterdam, 1991).
38. S. Chandrasekhar, *Hydrodynamics and Hydromagnetic Stability* (Oxford University Press, Oxford, 1961).
39. M. Allen and D. J. Tildesley, *Computer Simulation of Liquids* (Oxford University Press, New York, 1987).
40. S. Chandrasekhar, *Rev. Mod. Phys.* **15**, 1 (1943).
41. D. P. Landau and K. Binder, *A Guide to Monte Carlo Simulation in Statistical Physics* (Cambridge University Press, Cambridge, 2000).
42. K. Lum and A. Luzar, *Phys. Rev. E* **56**, R6283 (1997).
43. E. I. Shakhnovich, *Curr. Opin. Struct. Biol.* **7**, 29 (1997); J. Shimada, E. L. Kussell, and E. I. Shakhnovich, *J. Mol. Bio.* **308**, 79 (2001).
44. D. Frenkel and B. Smit, *Understanding Molecular Simulation* (Academic Press, Boston, 1996).
45. N. Metropolis, A. W. Metropolis, M. N. Rosenbluth, A. H. Teller, and E. Teller, *J. Chem. Phys.* **21**, 1087 (1953).
46. H. Goldstein, *Classical Mechanics* (Addison-Wesley, Reading, MA, 1980).
47. H. C. Andersen, *J. Comp. Phys.* **52**, 24 (1983).
48. H. A. Posch and W. G. Hoover, *Phys. Rev. A* **39**, 2175 (1989).
49. P. G. de Gennes, *J. Phys. Chem.* **55**, 572 (1971).
50. M. I. Dykman, D. G. Luchinsky, P. V. E. McClintock, and V. N. Smelyanskiy, *Phys. Rev. Lett.* **77** 5529 (1996).
51. R. S. Maier and D. L. Stein, *Phys. Rev. Lett.* **69**, 3691 (1992).
52. D. M. Ceperley, *Rev. Mod. Phys.* **67**, 279 (1995).
53. H. C. Andersen, *J. Chem. Phys.* **72**, 2384 (1980).
54. M. N. Rosenbluth and A. W. Rosenbluth, *J. Chem. Phys.* **23**, 356 (1955); J. I. Siepmann and D. Frenkel, *Mol. Phys.* **75**, 59 (1992); J. J. de Pablo, M. Laso, and U. W. Suter, *J. Chem. Phys.* **96**, 2395 (1992); D. Frenkel, G. C. A. M. Mooij, and B. Smit, *J. Phys. Condensed Matter* **4**, 3053 (1992).

55. M. Sprik, M. L. Klein, and D. Chandler, *Phys. Rev. B* **31**, 4234 (1985).
56. D. D. Frantz, D. L. Freeman, J. D. Doll, *J. Chem. Phys.* **93**, 2769 (1990).
57. B. A. Berg and T. Neuhaus, *Phys. Rev. Lett.* **68**, 9 (1992).
58. C. J. Geyer and E. A. Thompson, *J. Am. Stat. Assoc.* **90**, 909 (1995).
59. E. Marinari and G. Parisi, *Europhys. Lett.* **19**, 451 (1992).
60. U. H. E. Hansmann, *Chem. Phys. Lett.* **281**, 140 (1997).
61. T. Lazaridis and M. Karplus, *Science* **278**, 1928 (1997).
62. CPMD Version 3.3; developed by J. Hutter, A. Alavi, T. Deutsch, M. Bernasconi, S. Goedecker, D. Marx, M. Tuckerman, and M. Parrinello, Max-Planck-Institut für Festkörperforschung and IBM Zurich Research Laboratory (1995–1999).
63. D. Chandler, *Introduction to Modern Statistical Mechanics* (Oxford University Press, New York, 1987).
64. U. W. Schmitt and G. A. Voth, *J. Chem. Phys.* **111**, 9361 (1999).
65. P. R. ten Wolde, M. J. Ruiz-Montero, and D. Frenkel, *Discuss. Faraday Soc.* **104**, 93(1996).
66. U. Gasser, E. R. Weeks, A. Schofield, P. N. Nursey, and D. A. Weitz, *Science* **292**, 258 (2000).
67. M. M. Klosek, B. J. Matkowsky, and Z. Schuss, *Ber. Bunsenges. Phys. Chem.* **95**, 331 (1991); V. Pande, A. Y. Grosberg, T. Tanaka, and E. I. Shakhnovich, *J. Chem. Phys.* **108**, 334 (1998).
68. M. Sprik and G. Ciccotti, *J. Chem. Phys.* **109**, 7737 (1998).
69. J. B. Tenenbaum, V. de Silva, and J. C. Langford, *Science* **290**, 2319 (2000); S. T. Roweis and L. K. Saul, *Science* **290**, 2323 (2000).
70. V. Zaloj and R. Elber, *Comp. Phys. Comm.* **128**, 118 (2000).
71. D. N. Perera and P. Harrowell, *J. Non-Cryst. Solids* **235**, 314 (1998).
72. C. Donati, S. C. Glotzer, P. H. Poole, W. Kob, and S. J. Plimpton, *Phys. Rev. E* **60**, 3107 (1999).